

Seeing With a Few Photons: Bridging Cellular and Circuit Mechanisms With Perception[☆]

Roozbeh Kiani^a, Petri Ala-Laurila^{b,c}, and Fred Rieke^d, ^aCenter for Neural Science, New York University, New York, NY, United States; ^bDepartment of Neuroscience and Biomedical Engineering, Aalto University, Espoo, Finland; ^cFaculty of Biological and Environmental Sciences, Molecular and Integrative Biosciences Research Programme, University of Helsinki, Helsinki, Finland; and ^dDepartment of Physiology and Biophysics, University of Washington, Seattle, WA, United States

© 2020 Elsevier Inc. All rights reserved.

Abstract	1
Introduction	1
Behavior	2
Frequency-of-Seeing Analyses	2
New and Improved Behavioral Measurements	4
Beyond Photon Detection	4
Summary of Constraints From Behavior	5
Single Photon Detection in Rod Photoreceptors	5
Rods Respond to Single Photons	5
Noise in the Rod Responses	6
Summary	8
Retinal Readout	8
Convergence of Sparse Noisy Signals	8
Convergence of Sparse Noisy Signals, Revisited	11
Cortical Readout of Parallel Retinal Outputs	12
Extraction of Temporal Information	12
Limits on Sensitivity Set by Rod Noise	13
Open Questions	15
References	15

Abstract

What is the relationship between the mechanistic operation of neural circuits and the behaviors that they support? This question has been difficult to answer in part because of the complexity of the behavioral tasks neural circuits subserve and the difficulty in translating this complexity into a clear quantifiable objective of the system. Behavioral tasks and objectives are often simpler for near-threshold sensory tasks. Near sensory threshold, the objective of a system becomes to detect the stimulus feature of interest, and quantification of performance is straightforward. Such a focus on near-threshold tasks neglects some of the richness of behavior, but can allow for tight links between cellular and circuit mechanisms and behavioral performance. The ability of human observers to detect small numbers of photons is a well-studied example. This chapter summarizes progress toward understanding the neural mechanisms that enable and limit vision under these conditions.

Introduction

In starlight, less than 10 out of every 10,000 rod photoreceptors absorb photons during the ~200 ms integration time of rod signals. Yet this weak signal, embedded in noise arising in the remaining rods, can reliably guide visual behavior. Indeed, we have known for nearly a century that dark-adapted human observers detect dim lights with a fidelity limited as much by the division of light into discrete photons as by biological noise or inefficiency (reviewed by [Donner, 1992](#); [Field et al., 2005](#)). This performance places strict constraints on the reliability with which the underlying circuit mechanisms transduce and convey signals. Thus, vision at its sensitivity limits provides an excellent opportunity to bring together biophysical studies of single molecules and synapses, computational studies of signal processing and coding, and behavioral studies of overall system reliability.

[☆] *Change history:* May 2020. R. Kiani and P. Ala-Laurila updated the sections and figures.

This is an update of F. Rieke, 1.18 - Seeing in the Dark: Retinal Processing and Absolute Visual Threshold, Editor(s): Richard H. Masland, Thomas D. Albright, Thomas D. Albright, Richard H. Masland, Peter Dallos, Donata Oertel, Stuart Firestein, Gary K. Beauchamp, M. Catherine Bushnell, Allan I. Basbaum, Jon H. Kaas, Esther P. Gardner, *The Senses: A Comprehensive Reference*, Academic Press, 2008, Pages 393–412, ISBN 9780123708809, <https://doi.org/10.1016/B978-012370880-9.00410-2>.

What type of constraints does photon detection place on neural mechanisms? First, to support vision in starlight, rod photoreceptors must reliably signal the absorption of a single photon and resulting activation of a single rhodopsin molecule. The ability of the rods to detect single photons provides a unique opportunity to study how G-protein coupled receptor signaling cascades work at the level of single active receptor proteins. Second, the neural circuitry in the retina must effectively read out the sparse signals in the rod photoreceptor array to modulate the retinal output. An effective readout must both separate responses of rods absorbing photons from noise generated in other rods and minimize noise added during synaptic and cellular processing; similar issues arise in many other contexts. Third, the handful of additional spikes in the retinal output in response to the absorption of a few photons must be reliably amplified by cortical circuits to generate a behavioral response. Finally, the enormous amplification required for single photon detection in both retina and cortex must be appropriately tempered as light levels increase.

Because of the clarity of the task of detecting a few photons, consideration of constraints like those listed above leads to precise predictions about the operation of the relevant neural mechanisms, and these predictions guide our investigation of the circuit. In this chapter, we summarize the progress and challenges toward understanding the mechanisms supporting photon detection, with an emphasis on how they enhance our general understanding of neural function.

Behavior

A just-detectable flash delivers ~ 50 photons at the cornea. This realization was made only a few years after the realization that light itself was divided into discrete photons (von Kries and Eyster, 1907). Furthermore, a significant fraction of the photons arriving at the cornea will be lost due to absorption and scatter before reaching the retina. These observations together suggested that the quantal nature of light might limit the reliability of rod vision.

The division of light into discrete photons means that nominally constant light sources - e.g. a light bulb or a reflecting object in a scene - will produce visual inputs that vary randomly over time. These variations are described by Poisson statistics. The variance of a Poisson distribution is equal to its mean; hence, if we count photons from a Poisson source for a specific time interval, the variance of the count across intervals is equal to the mean count. The signal to noise ratio - the mean count divided by the standard deviation - increases as the square root of the mean count. This means that we are (blissfully) unaware of the quantal nature of light in typical daylight conditions because the mean photon count is high (i.e. at least 1000) in any time period of relevance for behavior, and variability over time or across nominally identical regions of space is small. But this situation changes dramatically as light levels decrease to those we might encounter in starlight (Fig. 1), and Poisson variations in photon counts substantially degrade the quality of the image reaching the retina. Importantly, these variations are irreducible, and as such set a fundamental limit to the performance of any imaging system, including the visual system. Below, we summarize work testing how close vision comes to this limit.

Frequency-of-Seeing Analyses

Classic "frequency-of-seeing" behavioral experiments in the first half of the 20th century suggested that Poisson fluctuations in photon counts limit the ability of dark-adapted human observers to detect dim lights (Hecht et al., 1942; van der Velden,

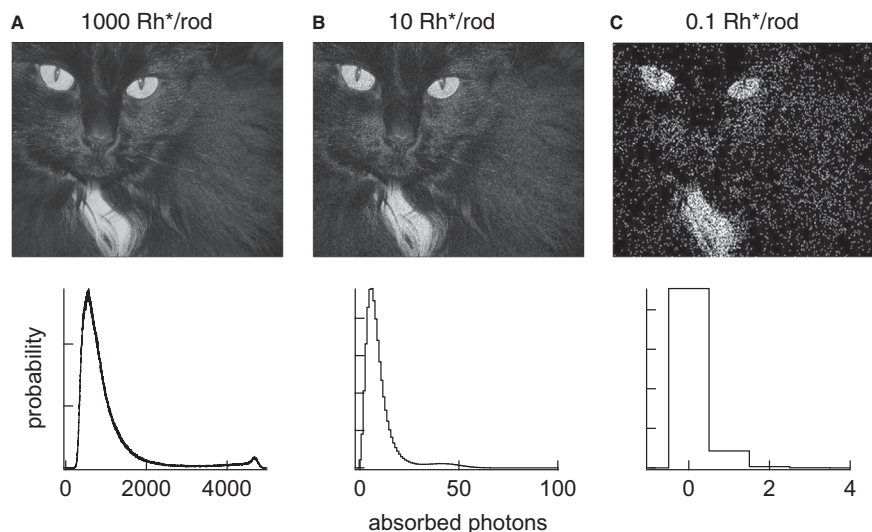


Figure 1 Poisson variations in photon absorption. Images in a-c represent photon counts in the rod array, with each pixel of the image representing a single rod. Lower panels show histograms of the photon counts across rod/pixels. (A) For high mean photon counts, Poisson variability does not noticeably degrade the image. As the mean photon count decreases (B,C), however, Poisson variability becomes more prominent. The lowest light level depicted here, in (C), is a factor of ~ 100 above absolute visual threshold.

1946). The basic experiment is simple: on each trial, a dark-adapted observer is asked to say “yes” if he/she saw a flash, and the probability of seeing is plotted against the number of photons delivered to the cornea (Fig. 2A). Observers are cued at the beginning of each trial and are trained to a certain rate of false-positive responses – that is, “yes” responses when no flash was delivered; this training sets a consistent criterion for detection. The resulting probability or frequency of seeing increases gradually with increasing number of photons at the cornea (Fig. 2A), with reliable detection possible with 50–100 photons. Interpreting these results in terms of the number of photons absorbed in the rod photoreceptors, however, requires analyses that incorporate the unknown probability that a photon at the cornea is absorbed within the rod photoreceptor array.

Hecht et al. (1942) and (independently) van der Velden (1946) hypothesized that the gradual transition between flashes that were rarely seen and those that were nearly always seen was due to a threshold number of photons required for seeing and the Poisson nature of light. According to this hypothesis, a flash that was seen on 50% of the trials, for example, on some trials (randomly) produced less than the threshold number of absorbed photons and was not seen, and on the other trials produced more than the threshold number of absorptions and was seen. The assumption that Poisson fluctuations dominate visual behavior is bold, as it leaves no room for biological noise. We will describe their results first and then return to the issue of how strongly they support this assumption.

The hypothesis that Poisson noise dominates which flashes are seen and unseen can be formalized by fitting the frequency-of-seeing curves with cumulative Poisson distributions (Fig. 2B). Key to this analysis is that the width of the transition from unseen to seen flashes depends on threshold. Conveniently (especially in the 1940s), when the probability of seeing is plotted against the logarithm of the number of photons at the cornea, the unknown probability that a photon at the cornea produces a response within the rod array is a x -axis translation of the theoretical curves in Fig. 2B without a change in shape. Hence, fitting the experimental data is simply a matter of determining which of the curves in Fig. 2B, slid along the x -axis, fits the data in Fig. 2A best. This procedure led to estimated thresholds of 5–7 (Hecht et al., 1942) and 2 (van der Velden, 1946) absorbed photons spread across the ~ 500 rods within the region covered by the stimulus. The x -axis position of the fits corresponded to 5%–10% of the photons at the cornea producing an electrical response in the rods.

These early frequency-of-seeing experiments indicate clearly that dark-adapted humans can detect a small number of absorbed photons. Further, since these few (<10) photons were spread over many (~ 500) rod photoreceptors, these experiments require that rods reliably signal absorption of single photons and that downstream neurons discount noise in the activity of rods that did not absorb photons (issues that we return to in the next section). At the same time, differences in threshold from apparently identical studies (i.e., Hecht et al. and van der Velden) precluded a definitive measurement of sensitivity and highlighted several confounds in the analysis. First, the observer’s criterion for responding was determined somewhat arbitrarily. Indeed, differences between thresholds estimated by Hecht et al. and van der Velden arose at least partially from differences in the detection criteria and corresponding rate of false-positive responses adopted by their observers (Marriot et al., 1959). Consistent with this picture, relaxing the criterion by allowing observers to respond “maybe” and “yes” revealed a tradeoff between false-positive responses and threshold (Barlow, 1956; Teich et al., 1982) – with “maybe” responses exhibiting a higher false-positive rate and lower threshold. This tradeoff suggests that the behavioral threshold is not a fundamental quantity describing how the visual system works but instead depends on the leniency of the experimenters in accepting false-positive responses.

A second issue neglected in these early experiments was the noise within the visual system that produced the false-positive responses. False-positive responses occur when no light is present, and hence are unexpected from a purely Poisson-limited detector. The presence of false-positive responses thus indicates a source of noise within the system. It is convenient to express this noise as an equivalent background, or “dark light,” which can be compared directly with true visual inputs (Fig. 2C). When the noise, expressed as noise in the input, exceeds a threshold number of photon-like events, it causes the perception of a flash – that is, a false-positive response. Similarly, these equivalent noise events generate a background against which true photon absorptions must be detected (Fig. 2C). This noise can be incorporated into the frequency-of-seeing analysis by assuming that the number of photon-like events

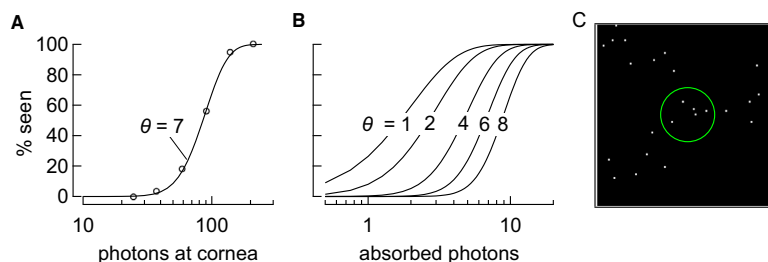


Figure 2 Frequency-of-seeing analysis. (A) Frequency-of-seeing data replotted from Hecht et al. (1942). Open circles are measurements from one observer. The smooth curve is a cumulative Poisson distribution calculated for a threshold of 7 absorbed photons spread across the rods in the area covered by the stimulus. (B) Predicted frequency-of-seeing curves calculated for a range of thresholds. Each curve represents the probability that θ or more photons are absorbed on a single trial plotted against the mean number of absorbed photons. (C) Simulation of dark photon-like noise events in the rod array. Each pixel represents the response of a single simulated rod. The green circle represents the retinotopic location of the flash; photon-like noise events within this circle (represented by bright pixels) will obscure responses to true photon absorptions and will give rise to false-positive responses (i.e., the false perception of light when none is present) when their number exceeds the threshold for detection θ .

consists of contributions from both internal noise and true photon absorptions. Fits to frequency-of-seeing curves with lenient criteria (and hence a non-zero false-positive rate) provided estimates of both threshold and internal noise (Barlow, 1956; Sakitt, 1972; Teich et al., 1982). Such analyses indicate an internal noise equivalent to 0.005–0.02 photon-like events per rod per second (reviewed by Donner, 1992; Field et al., 2005). Importantly, this internal noise represents all the noise generated in the steps between photon absorption and perception, expressed as an equivalent input noise.

With the inclusion of noise, fits to frequency-of-seeing curves can become uncertain. This uncertainty originates because several factors, alone or together, can produce similar shifts in the curves (Fig. 3). Specifically, in the classic analysis, any factor that broadens the frequency-of-seeing curve, including internal noise or fluctuations in an observer's criterion, will decrease the estimated threshold. Indeed, more recent estimates (reviewed in Nelson, 2017) suggest that ~20% of the photons at the cornea generate responses in the rods, higher than the 5%–10% estimated by Hecht et al. (1942). This attenuation reflects a loss of about 50% of the incident photons through scatter and absorption prior to reaching the rods, about 50% of the photons reaching the retina failing to be absorbed within the rods, and 30%–40% of the absorbed photons failing to generate an electrical response (Dartnall, 1972; Baylor, Lamb and Yau, 1979a). Thus, while the early frequency-of-seeing experiments certainly indicate that dark adapted humans can detect a small number of absorbed photons, they do not unambiguously determine whether detection of single photons is possible or provide precise estimates of the internal noise that limits sensitivity.

New and Improved Behavioral Measurements

Two recent studies mitigate some of the confounds associated with false-positive responses and the observers' subjective criteria. These experiments provide tighter bounds on the minimum number of photons required for detection and on the internal noise that limits sensitivity.

Koenig and Hofer (2011) systematically explored the relationship between criterion and threshold. They used this dependence to test two models: one in which every absorbed photon combines with internal noise to determine perception, and another in which a neural threshold somewhere in the visual system dictates that more than one photon is required for seeing. These models predict different dependencies of behavioral threshold on an observer's criterion. They conclude that their data, and that from an earlier related study (Sakitt, 1972), is more consistent with a model in which an internal threshold imposes that more than one absorbed photon is required for detection.

A second recent study (Tinsley et al., 2016) probed absolute threshold using a two-alternative forced choice design in which the observer's task was to identify in which of two time periods a flash was delivered. Because observers respond on every trial, this design removes the effect of variability in criteria across observers or across time within a single observer. These experiments show clearly that human observers can detect flashes resulting in no more than two absorbed photons, and the authors suggest that even single absorbed photons may permit better-than-chance performance. The same experiments revealed an interesting history dependence in the observer's performance. Specifically, following a trial in which a single photon was delivered to the cornea, discrimination performance on a subsequent trial was elevated for 5–10 s. The duration of this priming effect substantially exceeds the integration time of retinal neurons at the same light levels, and hence suggests that a history-dependent modulation in the sensitivity of neurons in later visual areas shapes perceptual sensitivity. This stimulus-induced increase in sensitivity in a given spatial region is in effect a form of attention working at the level of single photons.

Beyond Photon Detection

There is more to vision at low light levels than just detecting the presence or absence of dim lights. One challenge in using vision to guide behavior at low light levels is the slow kinetics of the rod-mediated responses. This slowness originates from a need for longer integration times to capture sparse photons and amplify the resulting molecular activations to produce sizable electrical signals. If uncompensated, these slow responses should cause perception to lag behind actual events in the world. For example, in toads, the




Figure 3 Ambiguities in frequency-of-seeing analysis. Each panel illustrates how the predicted frequency-of-seeing curve shifts when a specific variable in the analysis is altered. Decreases in the observer's threshold (left) shift the frequency of seeing curve to the left and broaden it. Increasing dark noise (middle) similarly shifts the curve to the left and broadens it, while also increasing the false positive rate (the asymptotic value of the curve for 0 photons at the cornea). Increasing the quantum efficiency – the likelihood that a photon at the cornea is absorbed by a rod – shifts the curve to the left but does not change its shape. The similar effects of these changes in variables introduce uncertainty into estimates from fitting experimental frequency-of-seeing curves.

latency to spiking in a ganglion cell following a dim flash is > 1 s. Because these animals hunt moving objects (e.g., crickets), the slow retinal responses threaten to introduce large spatial offsets in the actual and perceived position of the prey (Aho et al., 1993). The animals are able to compensate for these delays (Fig. 4), much like a tennis player must compensate for visual delays to predict the position of an oncoming ball at higher light levels. It is unclear where and how such a prediction is formed, and what limits to the accuracy of prediction are imposed by noise in the stimuli and processing circuitry. More generally, it will be very interesting to explore the range of visually-guided behaviors that are possible near absolute threshold, and to understand what enables and limits the fidelity of these behaviors.

Summary of Constraints From Behavior

To summarize, behavioral measures of absolute sensitivity: (1) pose clear limits on how much noise could be added throughout neural processing; (2) show that the signals resulting from just a few rhodopsin molecules, and perhaps even a single molecule, are amplified sufficiently to produce a motor output; and, (3) show that visually-guided behavior at absolute threshold depends on computations more complex than simple light detection.

Single Photon Detection in Rod Photoreceptors

The behavioral sensitivity to flashes producing much less than one photon absorbed per rod indicates that rod photoreceptors can detect a single absorbed photon. The rods perform this task admirably, with a sensitivity that compares favorably with the best man-made light detectors operating at room temperature. Understanding how this performance is achieved using only proteins and lipids has been a major achievement (reviewed by Pugh and Lamb, 1993; Rieke and Baylor, 1998a; Hamer et al., 2005). The rods meet several functional requirements for reliable photon detection: (1) they efficiently capture incident photons; (2) they amplify the signals resulting from a single rhodopsin molecule activated by absorption of a photon to produce a macroscopic electrical response; (3) they maintain low dark noise; and, (4) they generate near-identical responses to each absorbed photon.

Rods Respond to Single Photons

Initial attempts to test the prediction that rods can detect single absorbed photons were confounded by the spread of voltage changes generated in one rod to neighboring rods via gap junctions (Hagins et al., 1970; Schwartz, 1976; Fain et al., 1976; Copenhagen and Owen, 1976). The development of suction electrode techniques permitted measurement of the current responses of individual rod outer segments without contamination from signals of other rods (Baylor, Lamb and Yau, 1979b). These recordings confirmed that rods can indeed respond to single photons (Fig. 5).

But these recordings also led to broader insights. First, of the many G-protein coupled signaling cascades in biological systems, rod phototransduction provides a unique opportunity to characterize signals originating from individual active receptors, and this has provided a view of how these cascades work that would not be possible from responses averaged across many active receptors or stimulus trials. Indeed, phototransduction in rods is the best understood of the many G-protein cascades in biological systems. This

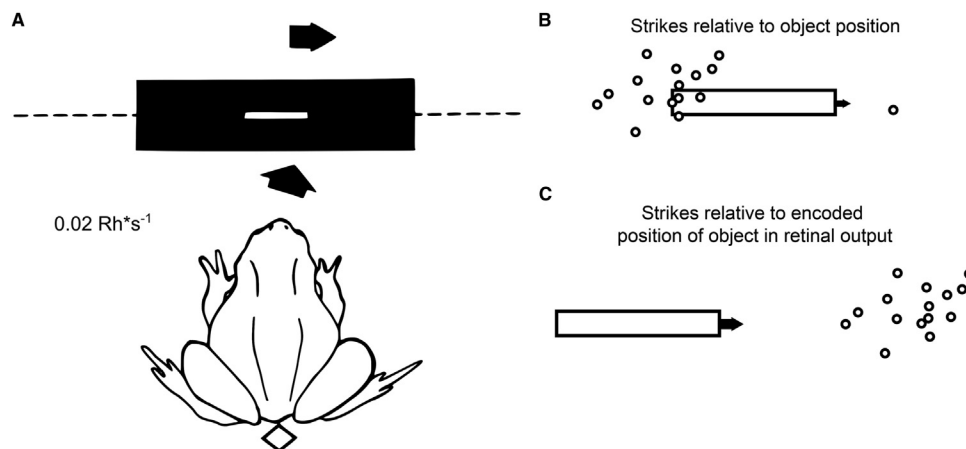


Figure 4 Compensation for visual delays near absolute threshold. (A) Behavioral task. Light objects on a dark background - simulating the appearance of prey - are moved in front of a toad. The toad strikes to attempt to capture the object. (B) Illustration of actual positions of tongue strikes relative to the object. (C) Estimate of strike positions relative to the encoded position of the object in the retinal output signals. The encoded position is calculated from the delay of responses of retinal ganglion cells at this light level and the speed of motion. The strikes are well ahead of the encoded position - indicating a compensation for the delay associated with retinal processing. Adapted from Aho et al. (1993).




Figure 5 Rod photoreceptors respond to single photons. Left panels show a mouse rod before (left) and after (right) being drawn into a suction electrode which measures the outer segment current. Right panel shows the recorded membrane current in response to a series of repeated dim flashes. Discrete responses are generated by absorption of single photons, with the largest response likely resulting from absorption of two photons. Data from Thuy Doan.

understanding has led to direct medical benefits, as we now understand the mechanisms and have potential treatments for several forms of stationary night blindness (Dryja, 2000).

Second, from combined progress in physiology, molecular biology and biochemistry, we now have a good quantitative understanding of how activation of a single rhodopsin molecule is amplified to produce a measurable electrical response. The high amplification provided by the phototransduction process protects the single photon response from noise introduced in downstream processing. This amplification, however, results in a long-lasting single photon response. This is an inevitable consequence of how second messenger cascades work: the underlying reactions require time to produce an amplified response even if they operate at or near the limit set by the rate of diffusional encounters between components (Pugh and Lamb, 1993). Furthermore, packing more proteins into the rod in an attempt to speed diffusion, and with it amplification, would likely hurt rather than help because the protein concentrations are already sufficiently high to hinder diffusion (Calvert et al., 2001). We return to the functional consequences of the slow rod response in the next section in considering the temporal precision of the retinal output signals.

Third, suction electrode recordings allowed measurement of the noise in the rod responses and comparison of this noise with the noise limiting behavioral sensitivity (Baylor et al., 1980; Baylor et al., 1984). We describe this noise and its consequences next.

Noise in the Rod Responses

Noise in the rod responses has three sources (Fig. 6). First, in complete darkness, rods generate occasional discrete noise events due to the spontaneous or thermal activation of rhodopsin molecules (Baylor et al., 1980; Baylor et al., 1984). Thermally- and light-activated rhodopsin molecules activate the downstream elements of the phototransduction cascade identically or near-identically. As a result, the discrete noise events are indistinguishable from events generated by photon absorption. Rhodopsin is quite stable, and the discrete noise events occur only once every 250–300 s in a primate rod at physiological temperature (Field et al., 2019). This means that on average, each of the $\sim 10^8$ rhodopsin molecules in a primate rod activates spontaneously once or twice a millennium. Second, the outer segment current fluctuates continuously due to spontaneous activation of components of the phototransduction cascade downstream of rhodopsin (Baylor et al., 1980; Rieke and Baylor, 1996). Unlike the discrete noise, continuous noise differs in amplitude and frequency content from responses to single photons. Third, the single photon responses themselves vary from one to the next, reflecting variability in the time course of rhodopsin's activity following the absorption of a photon (Rieke and Baylor, 1998b; Whitlock and Lamb, 1999; Doan et al., 2006).

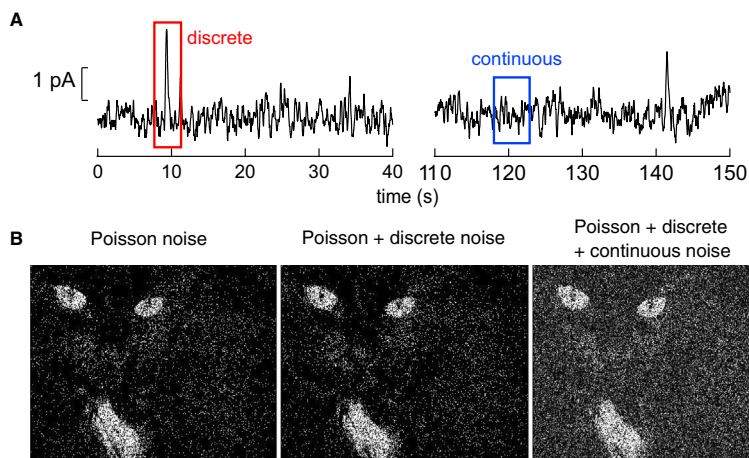


Figure 6 Rod dark noise. (A) Fluctuations in rod outer segment membrane current, measured as in Fig. 5. Discrete photon-like events (red box) and continuous fluctuations (e.g., blue box) are apparent. (B) Simulation of the impact of different noise sources on image quality. Each pixel represents the response of a single rod, with the pixel intensity representing the magnitude of the rod response. The left panel incorporates only noise from Poisson variations in photon absorption, the middle panel also includes discrete photon-like noise events, and the right panel includes both discrete and continuous noise.

The similarity of the discrete noise events to those elicited by absorbed photons led to immediate comparisons of this noise with the noise inferred from the rate of false-positive responses in the frequency-of-seeing behavioral experiments described above. Indeed, the noise inferred from behavior is within a factor of 2–4 of the rate of discrete noise events (reviewed by Donner, 1992; Field et al., 2005). Experimental uncertainty, particularly in the rate inferred from behavior, precludes a more definitive comparison. An alternative approach is to compare the temperature dependence of behavior with that of the rate of thermal activation of rhodopsin. This is not particularly practical in mammals, but it is possible in cold-blooded animals. Indeed, frogs and toads are willing to hunt over a broad temperature range. The threshold light level for their behavior depends inversely on temperature (i.e., frogs and toads can successfully hunt at lower light levels when it is colder). This dependence is highly correlated with the temperature dependence of the rate of thermal activation of rhodopsin (Aho et al., 1988). This similarity in temperature dependence is strong evidence that thermal isomerization of rhodopsin can explain much of the internal noise limiting behavior.

The apparent dominance of noise due to thermal isomerization of rhodopsin on behavioral sensitivity raises a question of why the other sources of rod noise apparently have little impact. Continuous noise produces current fluctuations that are generally smaller than single photon responses (Fig. 6). At the same time, continuous noise is omnipresent while the discrete noise events are rare; hence, continuous noise could take on increased importance when combining signals across rods. The smaller amplitude of the continuous noise means that it can be at least partially removed from the rod signals by appropriate nonlinear processing such as thresholding. Indeed, some such noise removal would appear to be essential to account for behavioral sensitivity. As we will see in the next section, such removal of continuous noise appears to be an important consideration for the operation of the retinal circuitry that reads out the rod signals.

The final noise source is variability in the single photon response. This variability is surprisingly small for a signal produced by a single molecule (Baylor et al., 1979a; Rieke and Baylor, 1998b), and this observation has led to insights into how single G-protein coupled receptors are turned off (Rieke and Baylor, 1998b; Whitlock and Lamb, 1999; Field & Rieke, 2002a; Doan et al., 2006). Because of the low variability of the single photon response, rods allow reliable photon counting. Photon counting, however, is of questionable importance for visual function because of Poisson fluctuations in photon absorption – that is, a visual input producing an average of m absorbed photons will produce $m \pm \sqrt{m}$ on a single trial (mean \pm SD). The rod itself, because of the low variability of the single photon response, is capable of encoding finer gradations in light intensity – i.e., the rod appears over engineered for the task of detecting incident photons. Single photon response variability, however, can have a substantial impact on how effectively single photon responses can be separated from continuous noise (see next section) and on the ability to estimate the time of photon absorption from the rod response. Such temporal information is likely important for tasks such as motion detection and prediction of the location of a moving object, as illustrated in Fig. 4.

Fig. 7 explores the impact of noise on the sensitivity of responses of single rods (Field et al., 2019); we return to how rod noise impacts responses of the collection of rods that guide behavior in the next section. This figure shows the ability to determine at which of two possible times (“early” or “late”) a flash was delivered based on individual rod responses. The left two panels show results for a 200 ms offset between the possible flash times, and the right panels show results for a 20 ms time offset. Each filled circle shows discrimination performance for an individual rod at a single flash strength (each rod was probed at 3–4 flash strengths). The red lines show the performance of an ideal detector that is limited only by Poisson variability in photon absorption. The ideal photon detector will perform the task correctly except for trials in which no photons are absorbed, in which case it is reduced to guessing (i.e., 50% correct).

For long time offsets (left), rod performance (circles) is close to that of the ideal photon detector. In other words, if the rod absorbs a photon, the resulting response can be detected and the time of the flash (early or late) identified with high reliability. For short time offsets (right), discrimination performance based on the rod responses falls short of the ideal photon detector but remains much better than chance even for time offsets much less than the duration of the rod response. This difference reveals the impact of internal noise. The temporal sensitivity is quite good: responses of individual rods permit discrimination of flashes producing only 1 activated rhodopsin molecule (1 Rh^*) for time offsets of only 20–40 ms, about 10% of the duration of the

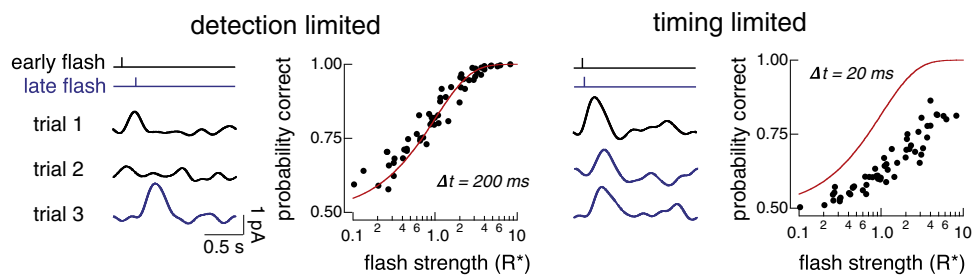


Figure 7 Measuring the sensitivity of the rod responses. Rod sensitivity can be measured from the ability to identify which of two possible flash times (early or late) corresponds to a given rod response. When the temporal offset between the possible flash times is large (left), performance is limited by the ability to detect the response, and only trials (e.g., trial 2) in which no response is apparent lead to identification errors. This conclusion is supported by the similarity of rod performance (closed circles) and the performance of an ideal photon detector (red line). When the time offset is small, trials in which a response is apparent cannot be unambiguously associated with a flash time. Performance in this case falls short of that of an ideal photon detector, indicating a role for internal noise. Adapted from Field et al. (2019).

response itself. Both continuous noise and single photon response variability contribute to limiting the temporal precision of the responses of single rods. We return to this issue below in considering how the collections of rod responses relevant for behavior are read out.

Summary

Behavioral sensitivity requires that rod photoreceptors detect single photons and maintain low noise. The ability to generate reliable responses from single active receptors challenges our understanding of second-messenger signaling. Progress in physiology, molecular biology and biochemistry has led to a clear understanding of how responses initiated by absorption of single photons are amplified by the rod phototransduction cascade and of the sources of noise in the rod responses. The similarity of rod noise and the total noise inferred from behavior indicates that behavioral detection of dim lights is largely limited by a combination of Poisson fluctuations in photon absorption and noise in the rod responses. This dominant role of rod noise in limiting behavior suggests that the retinal and cortical circuits that read out the rod responses operate effectively noiselessly. Explaining how these circuits add so little noise challenges our understanding of how they work.

Retinal Readout

The precision of photon detection by the rods would be wasted if not matched with an equally accurate readout circuit. To enable vision at low light levels, the retinal readout of signals in the rod array must meet three substantial challenges: (1) the circuitry must extract signals from the sparse pattern of photon absorptions created in the rod array while rejecting the noise present in all the rods; (2) the circuitry must extract and represent information about the timing of photon absorptions from the sluggish rod input signals; and, (3) the circuitry must transmit and process rod signals while adding little noise.

Rod signals traverse the retina through several pathways. The rod bipolar pathway (Fig. 8) dominates retinal signaling at or near visual threshold (Kolb and Nelson, 1983; Deans et al., 2002; Grimes et al., 2018). In this pathway, rod signals are first transmitted to rod bipolar cells – a class of ON or depolarizing bipolar cell that receives input almost exclusively from rods (see, however (Pang et al., 2010; Pang et al., 2018; Behrens et al., 2016)). Signals in the rod bipolar cell are transmitted to All amacrine cells and then to ON cone bipolar cells through gap junctions and OFF cone bipolar cells through glycinergic synapses. Signals from the cone bipolar cells are then conveyed to ganglion cells.

Convergence of Sparse Noisy Signals

Two features dominate the retinal processing of rod signals. First, rod signals at or near visual threshold are exceptionally sparse - with visually guided behavior possible when $< 0.1\%$ of the rods in the retinotopic location of the stimulus absorb photons within a 100–200 ms integration time. Second, many rod signals converge within the retinal circuitry to modulate the responses of retinal ganglion cells (Fig. 9A; Sterling et al., 1988). The manner in which these converging signals are combined is a critical factor determining the efficiency of vision under these conditions.




Figure 8 Rod bipolar pathway. This circuit conveys signals generated in the rod photoreceptors through the retina where they can modulate the retinal output signals. Signals from All amacrine cells reach ON cone bipolar cells through gap junctions (+ sign) and OFF cone bipolar cells through glycinergic synapses (- sign). RBC: rod bipolar cell. All: All amacrine cell.




Figure 9 Convergence of rod signals and impact on neural sensitivity. (A) Approximate numbers for convergence of rods to a peripheral primate ON parasol ganglion cell. Exact numbers depend on the ganglion cell type and retinal eccentricity, but the fact that many rod signals converge in successive stages holds for all ganglion cells and across retinal eccentricities, except in the rod-free fovea. (B) Stimulus-response relations for neurons across the rod bipolar pathway from mouse retina. Responses were estimated from families of responses to flashes of increasing strength; examples of such families are shown at the top. Data from A.P. Sampath and Felice Dunn.

The gain of signal transfer between cells in the rod bipolar pathway is well matched to the pattern of convergence. This is apparent in Fig. 9 where the stimulus-response relations for cells later in the pathway progressively shift to the left, indicating an increase in sensitivity. Thus, sparse single-photon responses in the rod array are boosted to create sizable responses in downstream cells. The high gain of signal transfer, together with the high gain of the phototransduction process itself, means that activation of a single rhodopsin molecule can produce a measurable change in the firing rate of a retinal ganglion cell. For example, ganglion cells in both cat and primate can produce one to three extra action potentials when one of the thousands of rods within their receptive field absorbs a photon (Barlow et al., 1971; Mastronarde, 1983; Ala-Laurila and Rieke, 2014). This high gain is essential for vision at absolute threshold, but it also threatens to saturate retinal responses at slightly higher light levels in which 1%–2% of the rods absorb photons during the 100–200 ms integration time of rod signals. Such saturation is prevented by adaptive mechanisms that decrease synaptic gain in the rod bipolar pathway (Dowling, 1963; Frishman and Sieving, 1995; Dunn et al., 2006; Schwartz and Rieke, 2013).

The discussion above focuses on average responses to many stimulus trials (as in Fig. 9) and hence neglects the noise that obscures the light responses of cells in the rod bipolar pathway. How noise is excluded, retained, or generated is central to understanding the relation between neural responses and behavior. Multiple stages of processing in the retina and cortex face the problem of convergence of sparse noisy signals. And this problem is not restricted to vision at low light levels given the commonality of sparse signals in neural circuits for a wide variety of tasks. Averaging across inputs via linear summation is a poor strategy under these conditions. This is because signal and noise are non-uniformly distributed across inputs, and averaging will inextricably mix inputs conveying signal with those conveying only noise. Instead, signaling fidelity can be substantially enhanced if each input is weighted by the likelihood that it conveys a signal - effectively identifying and retaining those inputs most likely carrying the signal of interest while rejecting the remaining inputs that are most likely noise.

What does this mean in the context of photon detection? At visual threshold, a small subset of the rods carries the signals that provide the basis for behavior. Yet all the rods generate noise. The task facing the retina is like standing in the middle of a crowded stadium and trying to determine what a few of the thousands of people in the stands are yelling. Averaging sound across the entire stadium is not effective. Instead, it is more effective to first try to identify the people of interest and retain sound only from them. This selection is most effective if it occurs on a person-by-person basis, based on prior information about the people of interest (e.g., that they are speaking in French). Signal processing strategies that use a similar solution to retain single photon responses while rejecting noise could greatly enhance visual performance (Baylor et al., 1984; van Rossum and Smith, 1998); this is potentially a 10–20-fold improvement in absolute sensitivity given the extreme sparsity of the inputs. To be effective, such a separation of signal and noise must occur before signals from different rods are mixed - i.e., within the rod itself or at the synapse between rods and rod bipolar cells.

The intuition from above can be formalized into a prediction of how rod response of different amplitudes should be weighted in transmission from rods to rod bipolar cells (van Rossum and Smith, 1998; Field and Rieke, 2002b; Bernston et al., 2004). This prediction is based entirely on measured rod signal and noise properties, with no free parameters. This is a quite unique opportunity in studying the relevant computational properties of neural circuits: we have a clear, parameter-free prediction of the relevant properties of a specific synapse, and getting these properties right has a large impact on circuit function and on behavior. Testing this prediction directly assesses our understanding of the functional role of the rod output synapse in photon detection.

Given this long-winded setup, it should not be surprising that the properties of the rod output synapse indeed agree with predictions from optimizing the detection of sparse signals (Fig. 10). Specifically, rod responses grow linearly with increasing flash

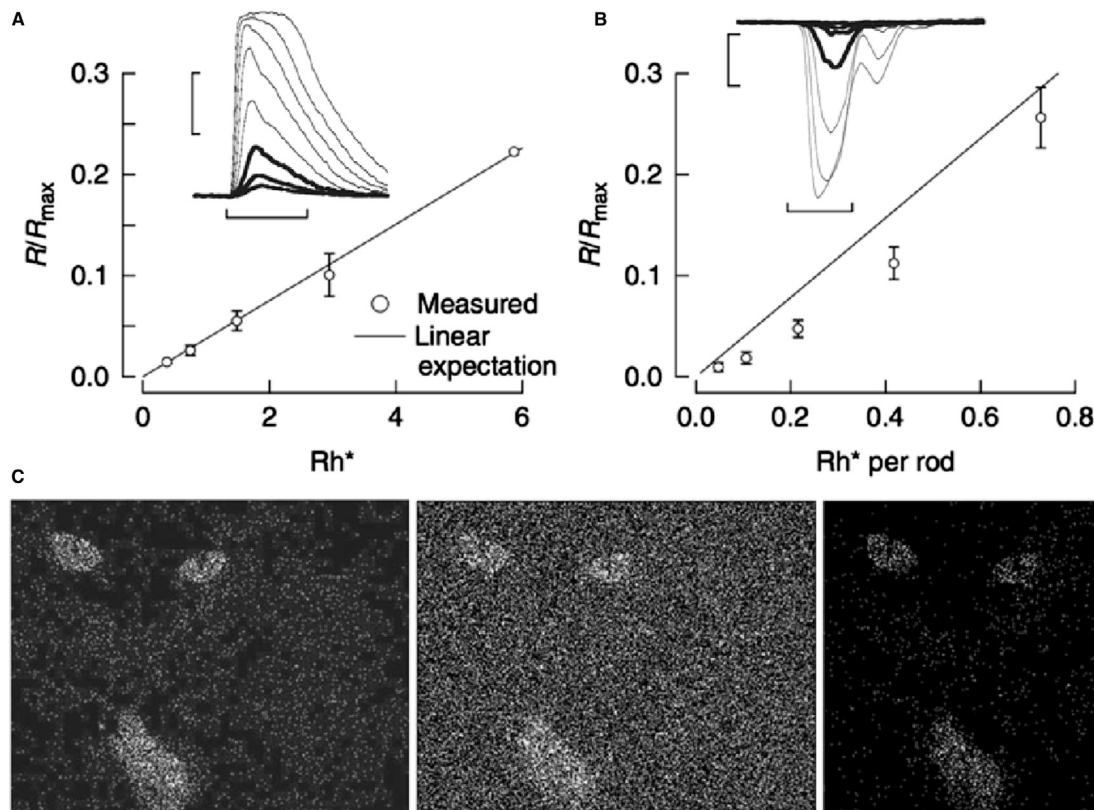


Figure 10 Evidence for nonlinear signal transfer from rods to rod bipolar cells. (A) Summary of the dependence of rod responses on flash strength. Responses to flashes producing $<5\text{--}6 Rh^*$ scale linearly with flash strength. (B) Summary of the dependence of excitatory synaptic inputs to rod bipolar cells on flash strength. Rod bipolar responses scale supralinearly with flash strength. (C) Impact of thresholding nonlinearity on encoding. The left panel shows simulated responses in the rod array with Poisson fluctuations; each pixel in the image represents the response of one rod photoreceptor. The middle panel shows this image with noise added to each pixel reflecting that measured in the rods (see also Fig. 6). The right panel applies a threshold to each pixel of the middle panel, with the height of the threshold set to separate signal (pixels representing absorbed photons) from noise. Modified from Field and Rieke (2002b).

strength (Fig. 10A) – i.e., a doubling in flash strength, within error bars, produces a doubling of the rod response. Responses of rod bipolar cells, however, grow supralinearly (Fig. 10B) – i.e., a doubling of flash strength produces more than a doubling of the rod bipolar response. This supralinearity can be explained if the rod bipolar response elicited when a rod responds to two absorbed photons is more than twice that elicited when the rod responds to a single photon. This is why the responses begin to grow nonlinearly for flashes producing an average of $0.5\text{--}1 Rh^*/\text{rod}$ – i.e., flashes in which the likelihood of $2 Rh^*/\text{rod}$ is sizable. Quantitative analysis of these experiments is consistent with the rod output synapse effectively acting as a threshold that eliminates small rod responses, most of which produced by continuous noise (Field and Rieke, 2002b). The position of this threshold is such that the synapse removes more than half of the rod’s single photon responses as well.

Eliminating single photon responses may seem like a poor strategy for optimizing visual sensitivity, but the cost of retaining these relatively small single photon responses would be a prohibitive increase in noise (Fig. 10C). One key component of this theoretical argument is that the low probability of photon absorption near absolute visual threshold introduces a strong prior probability (99.9% at $0.001 Rh^*/\text{rod}$) that a rod is generating noise and not a single photon response. Hence, if your job is to separate single photon responses from noise, you start with a strong bias that most rods are only producing noise and not responding to absorbed photons; the only single-photon responses that should be retained are those that have an amplitude sufficiently large to overcome this bias. This argument implies that the threshold should be positioned so that the likelihood that it is crossed by noise is less than 0.1% even though such a threshold position also eliminates many single photon responses. A position of the threshold that retains most or all of the single photon responses would greatly increase noise, and this increase in noise would cause sensitivity to suffer much more than the 2-fold effect of eliminating half of the single photon responses.

Thresholding nonlinearities are quite common in biological systems – e.g., in spike generation, or at synapses. The threshold at the rod output synapse operates through saturation of the synapse between rods and rod bipolar cells – specifically saturation of signaling in the rod bipolar cell dendrites. This saturation results in a supralinear dependence of the bipolar response on changes in glutamate release from the rods (Sampath and Rieke, 2004). Rod bipolar cells sense glutamate via a metabotropic receptor (mGluR6); binding of glutamate to this receptor closes ion channels in the rod bipolar dendrites. The decrease in glutamate release during the single photon response then decreases mGluR6 signaling and permits ion channels in the bipolar dendrites to open. The

threshold at the synapse occurs because the rate of glutamate release in the dark is sufficient to close nearly all the channels in the rod bipolar dendrites. Small fluctuations in glutamate release - such as those produced by continuous noise or by small single photon responses - do not produce a sufficient change in mGluR6 receptor activity to be effective in opening channels in the rod bipolar dendrites, and hence are eliminated in transmission. A key factor in the ability of this mechanism to separate signal and noise is its position: if the saturation is too strong, too many single photon responses will be rejected; if it is too weak, too much noise will be retained. How an appropriate position of the threshold is achieved is an interesting open problem.

Convergence of Sparse Noisy Signals, Revisited

The issue of convergence of sparse signals recurs at multiple stages of retinal processing. At visual threshold, only a few of the hundreds or thousands of rods that provide functional input to a ganglion cell absorb photons. Each rod bipolar cell receives input from 20–30 rods, and hence only a few percent of the rod bipolar cells will convey single photon responses through the retina while the remainder will convey only noise. Furthermore, cellular and synaptic mechanisms within the retinal circuitry between the rods and ganglion cells inevitably add additional noise that threatens to obscure the sparse rod signals.

A second thresholding stage, at the ON cone bipolar output synapse, removes much of the noise introduced between rods and ganglion cells (Ala-Laurila and Rieke, 2014). Evidence for this thresholding step comes from a suprathreshold increase in the ganglion cell response with increasing flash strength. For example, in Fig. 11, ON ganglion cell responses more than double for a doubling of flash strength (compare top and bottom panels on the left). This thresholding step is a property of the synapse between ON cone bipolar cells and retinal ganglion cells (see Fig. 8). It operates at light levels ~ 1000 -fold lower than the threshold at the synapse between rods and rod bipolar cells (i.e., near $0.001 \text{ Rh}^+/\text{rod}$ in Fig. 11). It is not present in the responses of OFF ganglion cells (Fig. 11, right) - an asymmetry we return to below.

How do these two nonlinearities together shape retinal output signals? As described above, thresholding at the synapse between rods and rod bipolar cells serves to separate signals from rods absorbing photons from noise generated in other rods. Importantly, this nonlinearity operates on signals from individual rods. The downstream nonlinearity operates on the collected signals from the 500–1000 rods that provide (indirect) input to a single ON cone bipolar cell. Surmounting this threshold requires two or more single photon responses or photon-like noise events in this collection of rods; these coincident rod responses must be separated in time by no more than 100–200 ms. The different light levels at which these two nonlinearities work - the threshold at rod output synapse operating near 1 absorbed photon per rod, and the threshold at the cone bipolar output synapse operating near 0.001 absorbed photons per rod - reflect the very different levels of rod convergence shaping signals at these two locations.

As summarized above, the nonlinearity at the rod output synapse is well positioned to maximize the ability to detect dim lights. Thresholding at the cone bipolar output synapse, however, is not easily explained by a similar argument. Instead, the relatively high threshold means that rod noise, including events from thermal isomerization of rhodopsin, produce little or no spiking in ON

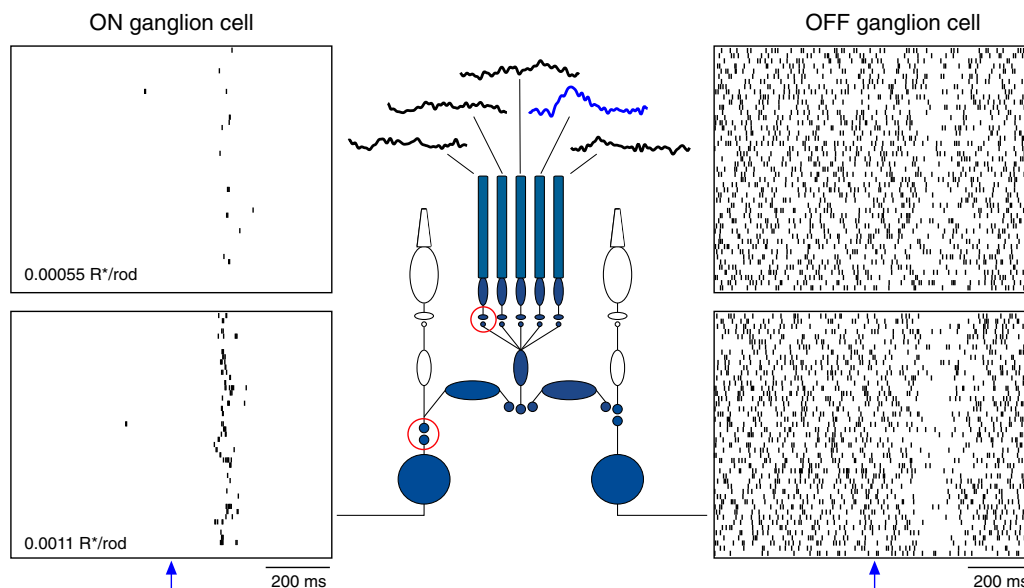


Figure 11 Encoding of near-threshold responses in the retinal output. Two thresholding nonlinearities (red circles in the middle panel) shape how converging rod responses modulate the retinal output. The first is at the synapse between rods and rod bipolar cells, and serves to separate signals from rods absorbing photons (e.g., blue trace at top of middle panel) from noise in the remaining rods (black traces); this is the nonlinearity described in Fig. 10. A second nonlinear step at the ON cone bipolar output synapse removes noise introduced within the retinal circuitry, and makes the responses of ON (but not OFF) ganglion cells near-deterministic reporters of light input - i.e. if the cells respond, light is almost certainly present. This can be seen from the low spontaneous firing rates of ON ganglion cells (left). Modified from Ala-Laurila and Rieke (2014).

ganglion cells; instead, these cells are near-silent in darkness, with an abrupt increase in firing rate for flashes that exceed threshold. This thresholding hinders absolute sensitivity, but makes the readout of the responses of these cells particularly simple: if the cell responds, some light is present with high reliability (e.g., Fig. 11, left).

What are the implications of this non-veridical representation of incident photons on behavior? As summarized above, a persistent issue in behavioral work on the limits of visual perception is whether human observers can in fact detect single photons, or if instead only flashes producing more than a threshold number of absorbed photons are seen. This binary separation into alternatives is in fact too stark, since a single photon response can combine with noise to exceed a thresholding step. The thresholding apparent in the responses of primate ganglion cells, however, should decrease performance in detecting single photons compared to expectations from performance for flashes producing 2–3 absorbed photons, and this failure of a linear model should be measurable in future behavioral experiments.

Cortical Readout of Parallel Retinal Outputs

ON and OFF retinal ganglion cells show an interesting asymmetry in their responses to weak flashes: ON cell responses are strongly shaped by the nonlinearity at the ON cone bipolar output synapse, while OFF cell responses depend near-linearly on flash strength (Fig. 11). This asymmetry is part of the evidence for a threshold located at the ON cone bipolar output synapse, since ON and OFF signals traverse much of the rod bipolar pathway together and a nonlinearity at an earlier location in the circuit would be expected to impact both ON and OFF ganglion cell responses (Fig. 8). The ON/OFF asymmetry is apparent in both mouse and primate retinas (Smeds et al., 2019). This means that OFF cells report photon absorptions more veridically, but with substantial noise, while ON cells report a thresholded version of photon absorption with very little noise. Just-detectable flashes based on OFF cell responses are as low or slightly lower than those based on responses of ON cells.

Surprisingly, mouse behavior under these conditions depends primarily on activity in ON cells (Smeds et al., 2019). Evidence for this conclusion comes from comparing ganglion cell responses and behavior in mice in which the single photon response in the rods was modified, through genetic manipulations, to be ~3-fold smaller than normal (Fig. 12, inset). The decreased rod response had a larger impact on sensitivity of ON ganglion cells than OFF ganglion cells due to the thresholding in the ON circuit. The sensitivity of visually-guided behavior exhibited a shift consistent with that of the ON cells (continuous red line in Fig. 12) but considerably larger than the shift exhibited by the OFF cells (dotted red line in Fig. 12). Thus, for this experimental paradigm, cortical circuits appear to ignore higher sensitivity inputs provided by OFF cells. This result brings up a host of interesting questions about which retinal output cells provide the basis for specific behaviors.

Extraction of Temporal Information

The rod-mediated dim flash responses of both amphibian (Ashmore and Falk 1980; Schnapf and Copenhagen, 1982) and mammalian (Berntson and Taylor, 2000; Euler and Masland, 2000; Field and Rieke, 2002b) bipolar cells are considerably briefer than the

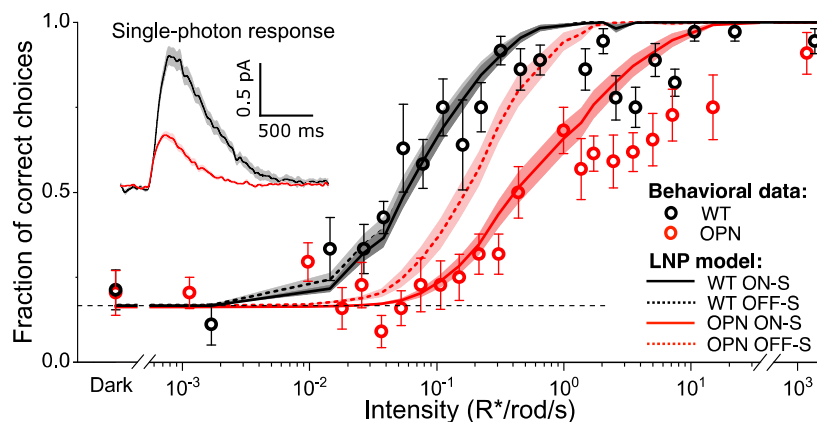


Figure 12 Comparison of mouse behavioral sensitivity with predictions based on responses of ON- and OFF-sustained (ON-S and OFF-S) ganglion cells. Comparisons are shown for wild-type mice and OPN mice, in which rods also express the cone photopigment. The inset shows single photon responses in wild-type (black) and OPN (red) mouse rod photoreceptors. Expression of the cone photopigment in rods decreases their sensitivity by a factor of 2–3. Due to the nonlinearity at the ON cone bipolar output synapse (see Fig. 11, middle), this reduction in rod sensitivity reduces the sensitivity of ON RGCs more than OFF RGCs. The open circles in the main panel show the fraction of trials in which a mouse correctly identified the arm of a water maze with a visual stimulus (with stimulus intensity on the x-axis); black symbols show behavior from wild-type mice and red symbols that of OPN mice. The smooth curves are fits based on an ideal observer analysis of the responses of ON-S and OFF-S ganglion cells. RGC predictions overlap with behavioral responses of wild-type mice (continuous and dotted black for ON-S and OFF-S RGCs, respectively). RGC predictions for OPN mice, however, differ, and behavioral sensitivity aligns more closely with the sensitivity predicted from ON RGCs (continuous red) than that predicted from OFF RGCs (dotted red). Adapted from Smeds et al. (2019).

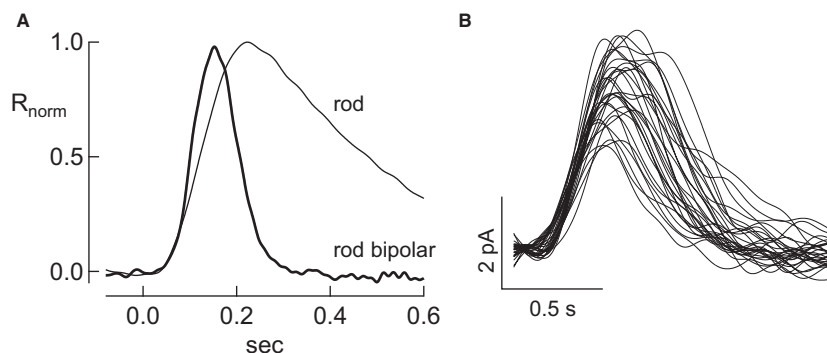


Figure 13 Rising phase of rod response dominates downstream signals. (A) Comparison of the time course of the single photon responses in mouse rods and rod bipolar cells. (B) Superposition of multiple isolated single photon responses from a mouse rod. Data from A.P. Sampath and Thuy Doan.

rod phototransduction currents. For example, Fig. 13A compares dim flash responses of mouse rods and rod bipolar cells. The rod bipolar response reaches peak and recovers much more quickly than the rod photocurrent; in fact, the bipolar response is nearly complete by the time the rod response reaches peak. The synaptic inputs to retinal ganglion cells evoked by dim flashes have kinetics similar to those of the rod bipolar responses (Field et al., 2005). Thus, the transfer of signals between rod outer segments and rod bipolar cells is a key determinant of the kinetics of rod mediated signals in the retina.

The speeding of rod responses in transmission to rod bipolar cells has several consequences for the encoding of single-photon responses. First, the time of photon absorption is difficult to discern from the relatively long-lasting rod response; this time is represented more explicitly in the briefer bipolar response. Second, the synapse transmits the most reliable part of the rod response – the rising phase – while suppressing the more variable recovery phase. Fig. 13B superimposes 20 isolated single-photon responses to illustrate the variability in the rod single-photon responses. Variability is small during the rising phase of the rod response and increases during the response recovery (Hamer et al., 2003; Field & Rieke, 2002a). The selective transmission of the initial part of the rod response should simplify the estimation of the time of photon arrival. We return to this issue in the next section.

It is useful to distinguish two different aspects of how time-varying signals are encoded. One aspect of temporal encoding is the ability to detect time-varying inputs. This aspect is often summarized by a “flicker-fusion” frequency - i.e., the temporal frequency at which a time-varying input appears constant. The flicker fusion frequency for dark-adapted rod vision is quite low (3–5 Hz) compared to that for cone vision (~60 Hz) (Hecht et al., 1933). The low fusion frequency for rod vision corresponds nicely to the duration of rod-mediated responses in retinal ganglion cells (Lee et al., 1997), and suggests that flicker fusion is at least largely determined by the speed of rod-mediated responses in the retina. A second aspect of the encoding of time varying inputs has to do with the precision with which the timing of a visual stimulus is represented. Thus, the onset time of a slow response - such as the rod response to a signal photon - can be determined with high precision even though the response itself is long lasting. We will refer to this latter property as temporal precision.

A direct approach to estimate the temporal precision of neural responses is to measure the similarity of responses to a repeated stimulus. Fig. 14 shows examples of the spike responses of a mouse ganglion cell to a repeated random stimulus with a mean intensity of 2.5 Rh^+ /rod/s (i.e., considerably above absolute threshold) and a contrast of 50%. Each vertical tick in Fig. 14B represents a single action potential, and each row shows the response to a single repetition of the stimulus. Fig. 14A shows an estimate of the average response of a single rod to the same stimulus. Fig. 14C shows a short section of both rod and ganglion cell responses (at the time of the red arrow in Fig. 14B). The ganglion cell responses show strong similarity from one trial to the next; in particular, they show similarity on a timescale much shorter than the timescale for modulations of the rod response, indicating that the temporal precision is much greater than might be expected from the slow rod inputs. Indeed, the standard deviation of the first spike time in bursts such as that in the bottom panel averages ~5 ms (range across events 2–10 ms; see (Murphy and Rieke, 2006)) for these stimuli, which is ~3% of the correlation time of the rod response.

The temporal precision of the rod responses illustrated in Fig. 14 is one of several examples of acuity that exceeds nominal limits set by the sampling properties of the receptors. Other familiar examples are chromatic sensitivity, which exceeds expectations from the coarse wavelength sensitivities of the cone photoreceptors, and spatial acuity, which exceeds expectations from the spacing between foveal cones. In each such instance, the true limit to acuity is set by the signal and noise properties of the receptor responses. We explore this issue - and specifically which rod noise sources limit sensitivity - next.

Limits on Sensitivity Set by Rod Noise

As described in the preceding sections, most studies of the relationship between rod noise and behavioral sensitivity have emphasized the role of thermal isomerization of rhodopsin and neglected the possible impact of other sources of rod noise. We now return to this issue, armed with information about how retinal circuits read out the rod signals. We consider the ability to detect dim lights and to determine their time of occurrence.

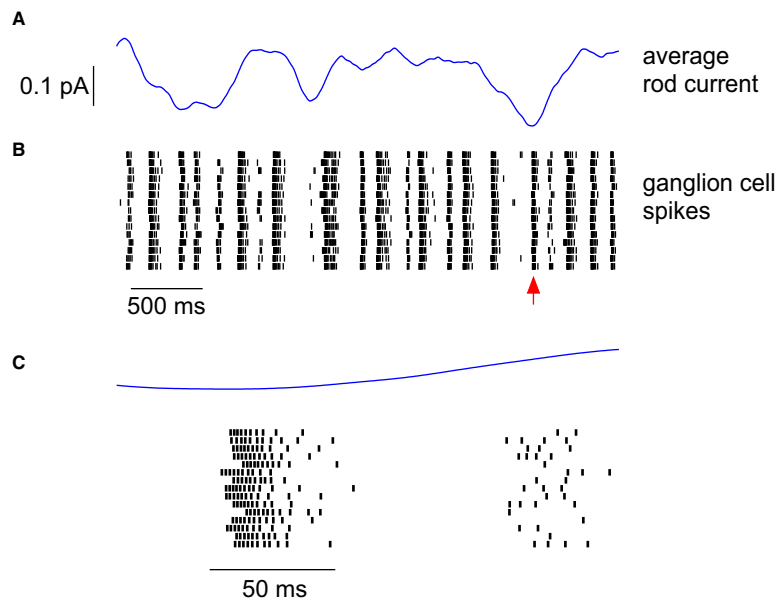


Figure 14 Rod and ganglion cell responses to a repeated random stimulus. (A) Rod responses were estimated by convolving the response to a brief flash with the stimulus. (B) Each row shows the spike response of a mouse ON sustained ganglion cell to one trial of the stimulus. (C) Responses at the time of the red arrow in (B) on an expanded time scale. Data from Gabe Murphy.

The first step in this analysis is to develop models that predict responses in rods and allow evaluation of the impact of each source of rod noise on behavioral sensitivity. These models apply an optimally positioned threshold to rod responses to separate signal and noise and pool the resulting signals over a large collection of rods, such as those that contribute to an observer's responses during a behavioral experiment (Fig. 15A). Such models, inspired by how rod bipolar cells read out rod signals, provide an optimal readout of the rod signals assuming that the only source of noise is that in the rods themselves. We then evaluate the sensitivity of these pooled signals using the same two-alternative forced choice task as that used for the signals from individual rods in Fig. 7. Generative models of the rod signals allow us to repeat this analysis while increasing or decreasing each of the three rod noise sources from its baseline value and hence to determine the impact of each noise source on detection and temporal sensitivity. If sensitivity improves sharply when we decrease a specific noise source, then we infer that noise source at least partially limits sensitivity. If sensitivity changes little when we manipulate a noise source, we infer that noise source is largely irrelevant for sensitivity.

Fig. 15B and C show results from this analysis. Each line represents the impact of manipulating one of the three sources of noise in the rod signals, with a value of 1 meaning that the noise was set to its measured value, less than 1 meaning that the noise was decreased in magnitude and greater than 1 meaning that it was increased. The noise sources that were not manipulated were held fixed at their baseline values (1 on this scale). The open symbol shows sensitivity for a linear readout of the rod signals with each noise source at its baseline value.

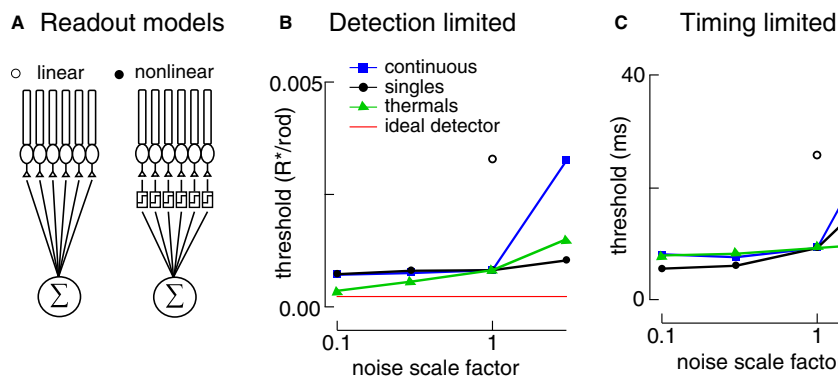


Figure 15 Limits to RGC sensitivity set by rod noise. (A) Linear and nonlinear models for pooling rod signals. Simulations in (B) and (C) are for 3000 rods. (B) Dependence of detection threshold (75% correct in a two alternative forced choice task like that in Fig. 7 for large time offsets of flashes) on each rod noise source. An individual noise source was varied while others were held at their baseline values (i.e., at a value of one on the x-axis). The open circle shows threshold for a linear readout of the rods. (C) Dependence of temporal sensitivity on rod noise sources. Adapted from Field et al. (2019).

Detection sensitivity of the rod pool falls short of the limit set by Poisson fluctuations in photon absorption (i.e., the ideal photon detector, red line in Fig. 15B), indicating a sizable contribution from internal noise. Increasing continuous noise (blue line) sharply reduced sensitivity, as this noise could no longer be effectively separated from single photon responses and removed by the nonlinear threshold. The importance of removal of continuous noise is also apparent from the improved detection sensitivity of nonlinear readout models compared to a linear readout model (open circle). Detection sensitivity improved when the rate of thermal noise events decreased (green line); decreasing continuous noise or single photon response variability had little effect on detection sensitivity, indicating that they did not limit sensitivity for baseline noise values. This analysis indicates that noise from thermal isomerization of rhodopsin limits the detection sensitivity of signals in the pool of rods that form the basis of ganglion cell responses and of the frequency-of-seeing and related behavioral experiments.

Temporal sensitivity had a different dependence on manipulations of rod noise sources (Fig. 15C). It was insensitive to changes in the thermal isomerization rate but decreased when continuous noise was increased. Temporal sensitivity also fell substantially when variability in the single photon response was increased. Temporal sensitivity depended more strongly on single photon response variability when it was made to occur uniformly throughout the response, rather than only during the recovery phase of the response (see Fig. 13). Hence the low variability of the single photon response, particularly the low variability of the initial rising phase of the response, is important for accurate encoding of the time of photon arrival.

The overall temporal sensitivity is also impressive. Even for flashes producing $< 0.01 \text{ Rh}^*/\text{rod}$, responses in the rod pool permit discrimination of temporal separations between the possible flash times as small as 10 ms, substantially less than the 200–400 ms duration of the rod response. Such temporal sensitivity could be important for tasks such as motion detection at low light levels. It will be interesting in future experiments to compare these calculated limits to temporal sensitivity with the behavioral ability to detect moving objects, and to compensate for delays and anticipate their position as in Fig. 4.

Open Questions

The ability of dark-adapted human observers to detect dim lights with a sensitivity approaching physical limits is a classic part of vision science. Indeed, this extreme sensitivity has guided decades of work on the underlying mechanisms. That focus has led to substantial breakthroughs in our understanding of how G-protein coupled signaling cascades can amplify signals from single receptors, and more recently of how neural circuits can separate signals of interest from noise. Yet our work is not done.

Past work has focused almost exclusively on the ability to detect the presence of weak lights, while neglecting other aspects of sensitivity that could be central to visually-guided behavior at low light levels. Temporal sensitivity is a clear example. For example, the ability to detect differences in timing of inputs at different spatial locations is central to motion detection. More generally, it is not clear what aspects of behavior are possible at low light levels, what computations are required to support these behaviors, and what noise sources limit the fidelity of these computations. Looking forward, it will be particularly interesting to explore whether rod noise limits the fidelity of vision at low light levels for tasks other than detecting a few photons.

Another glaring gap in our understanding of the mechanistic basis of absolute visual sensitivity regards the mechanisms that reside between the retina and motor output. Indeed, many of the same issues that have guided investigation of retinal circuits apply to cortex. Cortical circuits must amplify the handful (tens) of extra spikes sent down the optic nerve in response to a just-detectable flash to produce a motor output. And this amplification must occur without adding substantial noise, given that behavioral sensitivity approaches limits set by rod noise. This picture conflicts with views of cortex as noisy (Tolhurst et al., 1983; Deneve et al., 1999; Goris et al., 2014), with a high level of spontaneous activity (Kenet et al., 2003; Ringach, 2009; Kiani et al., 2015) that could mask weak retinal inputs. Understanding how these constraints are met may reveal an unappreciated mode of operation of cortical circuits.

References

- Aho, A.C., Donner, K., Helenius, S., Larsen, L.O., Reuter, T., 1993. Visual performance of the toad (*Bufo bufo*) at low light levels: retinal ganglion cell responses and prey-catching accuracy. *J. Comp. Physiol.* 172 (6), 671–682.
- Aho, A.C., Donner, K., Hyden, C., Larsen, L.O., Reuter, T., 1988. Low retinal noise in animals with low body temperature allows high visual sensitivity. *Nature* 334 (6180), 348–350.
- Ala-Laurila, P., Rieke, F., 2014. Coincidence detection of single-photon responses in the inner retina at the sensitivity limit of vision. *Curr. Biol.* 24 (24), 2888–2898.
- Ashmore, J.F., Falk, G., 1980. Responses of rod bipolar cells in the dark-adapted retina of the dogfish, *Scyliorhinus canicula*. *J. Physiol.* 300, 115–150.
- Barlow, H.B., 1956. Retinal noise and absolute threshold. *J. Opt. Soc. Am.* 46 (8), 634–639.
- Barlow, H.B., Levick, W.R., Yoon, M., 1971. Responses to single quanta of light in retinal ganglion cells of the cat. *Vision Res.* (Suppl 3), 87–101.
- Baylor, D.A., Lamb, T.D., Yau, K.W., 1979a. Responses of retinal rods to single photons. *J. Physiol.* 288, 613–634.
- Baylor, D.A., Lamb, T.D., Yau, K.W., 1979b. The membrane current of single rod outer segments. *J. Physiol.* 288, 589–611.
- Baylor, D.A., Matthews, G., Yau, K.W., 1980. Two components of electrical dark noise in toad retinal rod outer segments. *J. Physiol.* 309, 591–621.
- Baylor, D.A., Nunn, B.J., Schnapf, J.L., 1984. The photocurrent, noise and spectral sensitivity of rods of the monkey *Macaca fascicularis*. *J. Physiol.* 357, 575–607.
- Berrens, C., Schubert, T., Haverkamp, S., Euler, T., Berens, P., 2016. Connectivity map of bipolar cells and photoreceptors in the mouse retina. *Elife* 5.
- Berntson, A., Taylor, W.R., 2000. Response characteristics and receptive field widths of on-bipolar cells in the mouse retina. *J. Physiol.* 524 (Pt 3), 879–889.
- Berntson, A., Smith, R.G., Taylor, W.R., 2004. Transmission of single photon signals through a binary synapse in the mammalian retina. *Vis. Neurosci.* 21 (5), 693–702.
- Calvert, P.D., Govardovskii, V.I., Krasnoperova, N., Anderson, R.E., Lem, J., Makino, C.L., 2001. Membrane protein diffusion sets the speed of rod phototransduction. *Nature* 411 (6833), 90–994.

- Copenhagen, D.R., Owen, W.G., 1976. Coupling between rod photoreceptors in a vertebrate retina. *Nature* 260 (5546), 57–59.
- Dartnall, H.J.A., 1972. Photosensitivity. In: Dartnall, H.J.A. (Ed.), *Handbook of Sensory Physiology, Photochemistry of Vision*, vol. VII/1. Springer-Verlag, pp. 122–145.
- Deans, M.R., Volgyi, B., Goodenough, D.A., Bloomfield, S.A., Paul, D.L., 2002. Connexin36 is essential for transmission of rod-mediated visual signals in the mammalian retina. *Neuron* 36 (4), 703–712.
- Deneve, S., Latham, P.E., Pouget, A., 1999. Reading population codes: a neural implementation of ideal observers. *Nat. Neurosci.* 2 (8), 740–745.
- Doan, T., Mendez, A., Detwiler, P.B., Chen, J., Rieke, F., 2006. Multiple phosphorylation sites confer reproducibility of the rod's single-photon responses. *Science* 313 (5786), 530–533.
- Donner, K., 1992. Noise and the absolute thresholds of cone and rod vision. *Vision Res.* 32 (5), 853–866.
- Dowling, J.E., 1963. Neural and photochemical mechanisms of visual adaptation in the rat. *J. Gen. Physiol.* 46, 1287–1301.
- Dryja, T.P., 2000. Molecular genetics of Oguchi disease, fundus albinopunctatus, and other forms of stationary night blindness: LVII Edward Jackson Memorial Lecture. *Am. J. Ophthalmol.* 130 (5), 547–563.
- Dunn, F.A., Doan, T., Sampath, A.P., Rieke, F., 2006. Controlling the gain of rod-mediated signals in the Mammalian retina. *J. Neurosci.* 26 (15), 3959–3970.
- Euler, T., Masland, R.H., 2000. Light-evoked responses of bipolar cells in a mammalian retina. *J. Neurophysiol.* 83 (4), 1817–1829.
- Fain, G.L., Gold, G.H., Dowling, J.E., 1976. Receptor coupling in the toad retina. *Cold Spring Harbor Symp. Quant. Biol.* 40, 547–561.
- Field, G.D., Rieke, F., 2002a. Nonlinear signal transfer from mouse rods to bipolar cells and implications for visual sensitivity. *Neuron* 34 (5), 773–785.
- Field, G.D., Rieke, F., 2002b. Mechanisms regulating variability of the single photon responses of mammalian rod photoreceptors. *Neuron* 35 (4), 733–747.
- Field, G.D., Sampath, A.P., Rieke, F., 2005. Retinal processing near absolute threshold: from behavior to mechanism. *Annu. Rev. Physiol.* 67, 491–514.
- Field, G.D., Uzzell, V., Chichilnisky, E.J., Rieke, F., 2019. Temporal resolution of single-photon responses in primate rod photoreceptors and limits imposed by cellular noise. *J. Neurophysiol.* 121 (1), 255–268.
- Frishman, L.J., Sieving, P.A., 1995. Evidence for two sites of adaptation affecting the dark-adapted ERG of cats and primates. *Vision Res.* 35 (3), 435–442.
- Goris, R.L., Movshon, J.A., Simoncelli, E.P., 2014. Partitioning neuronal variability. *Nat. Neurosci.* 17 (6), 858–865.
- Grimes, W.N., Baudin, J., Azevedo, A.W., Rieke, F., 2018. Range, routing and kinetics of rod signaling in primate retina. *Elife* 7.
- Hagins, W.A., Penn, R.D., Yoshikami, S., 1970. Dark current and photocurrent in retinal rods. *Biophys. J.* 10 (5), 380–412.
- Hamer, R.D., Nicholas, S.C., Tranchina, D., Lamb, T.D., Jarvinen, J.L., 2005. Toward a unified model of vertebrate rod phototransduction. *Vis. Neurosci.* 22 (4), 417–436.
- Hamer, R.D., Nicholas, S.C., Tranchina, D., Liebman, P.A., Lamb, T.D., 2003. Multiple steps of phosphorylation of activated rhodopsin can account for the reproducibility of vertebrate rod single-photon responses. *J. Gen. Physiol.* 122 (4), 419–444.
- Hecht, S., Schlaer, S., Verrijs, C.D., 1933. Intermittent stimulation by light: II. The measurement of critical fusion frequency for the human eye. *J. Gen. Physiol.* 17 (2), 237–249.
- Hecht, S., Schlaer, S., Pirenne, M.H., 1942. Energy, quanta, and vision. *J. Gen. Physiol.* 25, 819–840.
- Kenet, T., Bibitchkov, D., Tsodyks, M., Grinvald, A., Arieli, A., 2003. Spontaneously emerging cortical representations of visual attributes. *Nature* 425 (6961), 954–956.
- Kiani, R., Cueva, C.J., Reppas, J.B., Peixoto, D., Ryu, S.I., Newsome, W.T., 2015. Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. *Neuron* 85 (6), 1359–1373.
- Koenig, D., Hofer, H., 2011. The absolute threshold of cone vision. *J. Vis.* 11 (1).
- Kolb, H., Nelson, R., 1983. Rod pathways in the retina of the cat. *Vision Res.* 23 (4), 301–312.
- Lee, B.B., Smith, V.C., Pokorny, J., Kremers, J., 1997. Rod inputs to macaque ganglion cells. *Vision Res.* 37 (20), 2813–2823.
- Marriot, F.H., Morris, V.B., Pirenne, M.H., 1959. The absolute visual threshold recorded from the lateral geniculate body of the cat. *J. Physiol.* 146 (1), 179–184.
- Mastrorade, D.N., 1983. Correlated firing of cat retinal ganglion cells. II. Responses of X- and Y-cells to single quantum events. *J. Neurophysiol.* 49 (2), 325–349.
- Murphy, G.J., Rieke, F., 2006. Network variability limits stimulus-evoked spike timing precision in retinal ganglion cells. *Neuron* 52 (3), 511–524.
- Nelson, P., 2017. From Photon to Neuron.
- Pang, J.J., Gao, F., Lem, J., Bramblett, D.E., Paul, D.L., Wu, S.M., 2010. Direct rod input to cone BCs and direct cone input to rod BCs challenge the traditional view of mammalian BC circuitry. *Proc. Natl. Acad. Sci. U. S. A.* 107 (1), 395–400.
- Pang, J.J., Yang, Z., Jacoby, R.A., Wu, S.M., 2018. Cone synapses in mammalian retinal rod bipolar cells. *J. Comp. Neurol.* 526 (12), 1896–1909.
- Pugh, E.N.J., Lamb, T.D., 1993. Amplification and kinetics of the activation steps in phototransduction. *Biochim. Biophys. Acta* 1141 (2–3), 111–149.
- Rieke, F., Baylor, D.A., 1998a. Single-photon detection by rod cells of the retina. *Rev. Mod. Phys.* 70, 1027–1036.
- Rieke, F., Baylor, D.A., 1996. Molecular origin of continuous dark noise in rod photoreceptors. *Biophys. J.* 71 (5), 2553–2572.
- Rieke, F., Baylor, D.A., 1998b. Origin of reproducibility in the responses of retinal rods to single photons. *Biophys. J.* 75 (4), 1836–1857.
- Ringach, D.L., 2009. Spontaneous and driven cortical activity: implications for computation. *Curr. Opin. Neurobiol.* 19 (4), 439–444.
- Sakitt, B., 1972. Counting every quantum. *J. Physiol.* 223 (1), 131–150.
- Sampath, A.P., Rieke, F., 2004. Selective transmission of single photon responses by saturation at the rod-to-rod bipolar synapse. *Neuron* 41 (3), 431–443.
- Schnapf, J.L., Copenhagen, D.R., 1982. Differences in the kinetics of rod and cone synaptic transmission. *Nature* 296 (5860), 862–864.
- Schwartz, E.A., 1976. Electrical properties of the rod syncytium in the retina of the turtle. *J. Physiol.* 257 (2), 379–406.
- Schwartz, G.W., Rieke, F., 2013. Controlling gain one photon at a time. *Elife* 2, e00467.
- Smeds, L., Takeshita, D., Turunen, T., Tiihonen, J., Westö, J., Martyniuk, N., et al., 2019. Paradoxical rules of spike train decoding revealed at the sensitivity limit of vision. *Neuron* 104 (3), 576–587.e11.
- Sterling, P., Freed, M.A., Smith, R.G., 1988. Architecture of rod and cone circuits to the on-beta ganglion cell. *J. Neurosci.* 8 (2), 623–642.
- Teich, M.C., Prucnal, P.R., Vannucci, G., Breton, M.E., McGill, W.J., 1982. Multiplication noise in the human visual system at threshold. 3. The role of non-Poisson quantum fluctuations. *Biol. Cybern.* 44 (3), 157–165.
- Tinsley, J.N., Molodtsov, M.I., Prevedel, R., Wartmann, D., Espigulé-Pons, J., Lauwers, M., et al., 2016. Direct detection of a single photon by humans. *Nat. Commun.* 7, 12172.
- Tolhurst, D.J., Movshon, J.A., Dean, A.F., 1983. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res.* 23 (8), 775–785.
- van der Velden, H.A., 1946. The number of quanta necessary for the perception of light in the human eye. *Ophthalmologica* 111, 321–331.
- van Rossum, M.C., Smith, R.G., 1998. Noise removal at the rod synapse of mammalian retina. *Vis. Neurosci.* 15 (5), 809–821.
- von Kries, J., Eyster, J.A.E., 1907. Über die zur Erregung des Sehorgans erforderlichen Energiemengen. *Z. Sinnesphysiol.* 41, 394.
- Whitlock, G.G., Lamb, T.D., 1999. Variability in the time course of single photon responses from toad rods: termination of rhodopsin's activity. *Neuron* 23 (2), 337–351.