CellPress
OPEN ACCESS

Trends in
Neurosciences

Review

Retinal receptive-field substructure: scaffolding for coding and computation

Sören J. Zapp [1,2,4], Steffen Nitsche [1,2,4] and Tim Gollisch [1,2,3,*,@]

The center-surround receptive field of retinal ganglion cells represents a fundamental concept for how the retina processes and encodes visual information. Yet, traditional approaches of using the receptive field as a linear filter to integrate light intensity over space often do not capture the responses of a ganglion cell to complex visual stimuli. Thus, models with local nonlinearities in subunits of the receptive field or with local temporal dynamics are emerging to better reflect relevant aspects of retinal circuitry and capture stimulus encoding. Here, we review recent efforts to identify such receptive-field substructure and evaluate its role in visual stimulus encoding. The concomitant development of new computational tools may pave the way toward a model-based, functional approach to retinal circuit analysis.

Highlights

Visual stimulus encoding by the retina is not fully captured by the center-surround receptive fields of retinal ganglion cells. Mounting evidence of nonlinear spatial signal integration under natural stimuli and of specific visual functions solved by distinct ganglion cell types indicates the need to better understand receptive-field substructure.

Nonlinear spatial integration can be captured by subdividing ganglion cell receptive fields into subunits, which are thought to correspond to presynaptic bipolar cells. Several statistical and model-based methods have recently been developed to identify the subunit layout from spiking responses of ganglion cells to visual stimuli.

The subunits provide scaffolding for retinal computations, which may act through local adaptation and inhibition to shape responses to dynamic stimulation and to extract specific visual features.

Beyond receptive-field filtering: aligning circuit complexity and functional diversity

The sense of vision in vertebrates relies on a range of different computations. These distill relevant information from the incoming light patterns while discarding irrelevant components and dynamically adapting to the recently encountered structure of inputs. It is now well established that the underlying computations begin in the retina [1]. In this neural network, a multitude of different neuron types (more than 100 in the mouse, [2]), form dozens of parallel information channels [3–5] to extract various visual features, such as lateral [6,7], looming [8], and local motion [9,10] as well as local contrast [11] and color [12].

Understanding what the diversity of neural hardware is for, and how it implements different computations, requires going beyond the often-held view that the main function of the retina is to filter incoming images by the center-surround receptive fields of ganglion cells. Instead, nonlinear interactions, local dynamics, and signal gating within the receptive field have become a focus of interest, requiring new approaches and techniques to dissect the substructure of receptive fields and the signal processing within [13]. These developments occur in a tight interaction of experimental investigations with advances in computational data analysis and retinal circuit models. Among the primary overarching goals are to: (i) determine what is needed to describe and predict retinal responses to natural stimuli; and (ii) infer the structure and operations of the presynaptic circuitry that shape the responses of the different output channels of the retina.

In this review, we highlight recent progress in analyzing the substructure of retinal ganglion cell receptive fields. We focus on the characteristics and functional roles of nonlinear signal integration over space and of local temporal dynamics within the receptive field. In doing so, we emphasize new computational techniques to infer retinal circuit structure from recordings of the spiking activity of ganglion cells. The new methodology resonates well with model-based analyses in other sensory areas and with current developments in the field of artificial neural networks. Given the preserved general structure of the retina across vertebrate species and the many similar

[1]Department of Ophthalmology, University Medical Center Göttingen, 37073 Göttingen, Germany
[2]Bernstein Center for Computational Neuroscience Göttingen, 37073 Göttingen, Germany
[3]Cluster of Excellence 'Multiscale Bioimaging: from Molecular Machines to Networks of Excitable Cells' (MBExC), University of Göttingen, 37075 Göttingen, Germany
[4]These authors contributed equally

*Correspondence: tim.gollisch@med.uni-goettingen.de (T. Gollisch).
@Twitter: @TimGollisch (T. Gollisch).

basic operations, the presented examples are drawn freely from retinas of different vertebrates, mostly monkey, mouse, and salamander.

The center-surround receptive field in the retina and its challenges

The retina transduces incoming light into neuronal signals in its photoreceptors and processes these signals through its network of interneurons (bipolar, horizontal, and amacrine cells) before the retinal ganglion cells send the messages from the retina about the visual world as spike patterns to various brain regions (Figure 1A). The function of this network is often summarized by referring to the famous center-surround structure of ganglion cell receptive fields (Figure 1B), in particular in the simplified views of general neuroscience textbooks and in research that considers processing in downstream brain regions. This is perhaps not surprising, given the long and successful history of research on receptive fields, from their application as spatial filters to predict responses to moving light spots [14] to their utility for efficient stimulus coding [15].

However, many aspects of retinal stimulus encoding cannot be explained by stimulus filtering with center-surround receptive fields. The perhaps best-known example are the so-called 'Y-type

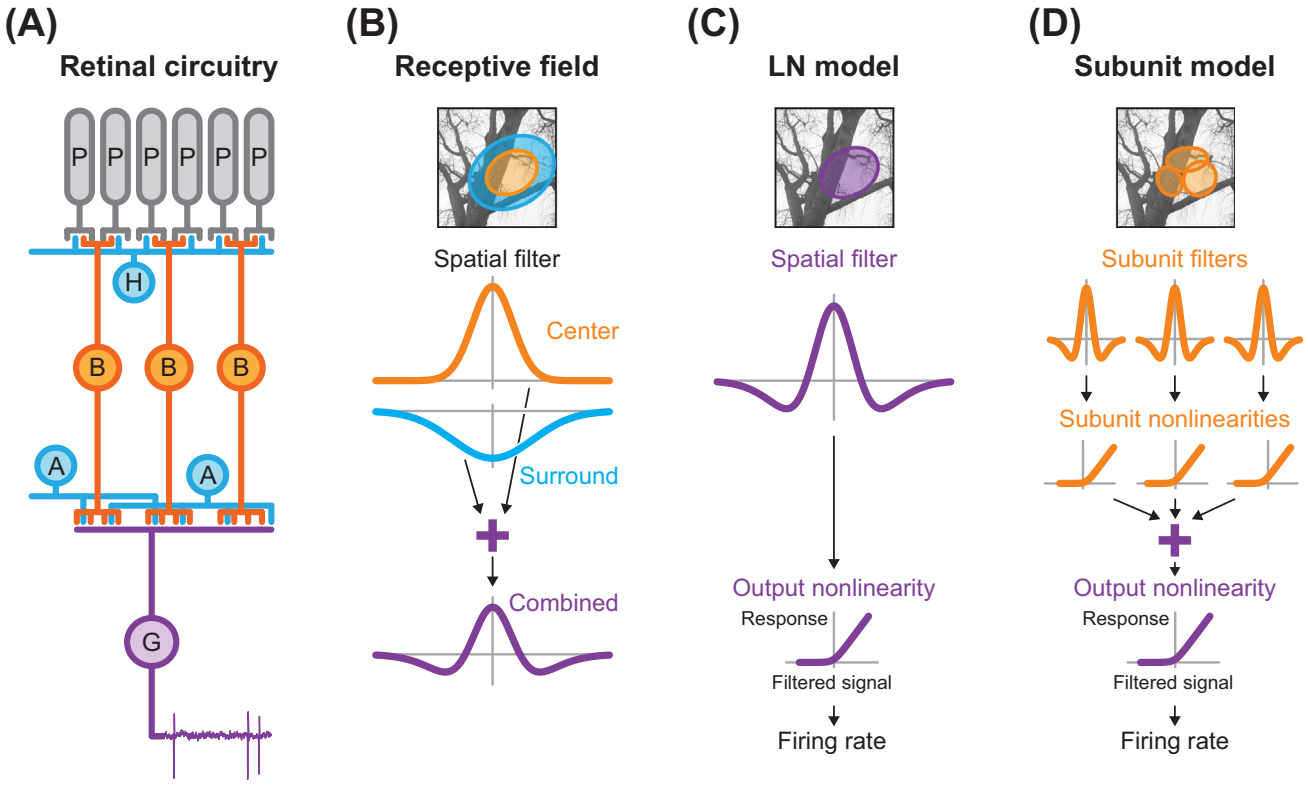

Figure 1. Retinal circuitry and standard models of spatial processing. (A) Schematic retinal circuitry, comprising photoreceptors (P), bipolar cells (B), two types of inhibitory interneurons [horizontal (H) and amacrine (A) cells], and a ganglion cell (G). Visual information propagates from top to bottom. The output is the spiking response of the ganglion cell, as indicated here by an extracellularly recorded voltage trace. (B) Ganglion cell center-surround receptive field, depicted here as a difference-of-Gaussians model. The spatial filter that comprises the receptive field contains a positive 2D Gaussian function for the center (orange, shown here as a 1D Gaussian cross-section) and a concentric, but broader negative Gaussian function for the surround (blue). The filter signals, obtained as weighted sums of pixel values, are added, yielding the combined receptive-field activation. Center and surround are attributed to excitatory bipolar cells and inhibitory interneurons, respectively, as indicated by the corresponding colors. Top: illustration of center and surround as 2D outlines overlaid on a visual stimulus. (C) Spatial linear–nonlinear (LN) model. First, a spatial filter, representing the receptive field of the ganglion cell, is applied to the visual stimulus, corresponding to a weighted summation of stimulus pixel values. The filtered signal is subsequently passed through a nonlinear transformation to produce the response of the ganglion cell, measured, for example, as the evoked firing rate. (D) Subunit model. Multiple subunits act as spatial filters with subsequent nonlinearities. Their responses are summed in a weighted manner and passed through a final nonlinearity to produce the response of the ganglion cell (e.g., firing rate).

ganglion cells', first found in the cat retina [16], which respond strongly (and with frequency doubling) to contrast-reversing spatial gratings with no net change in illumination over the receptive-field center or surround. This demonstrates that activation and deactivation from brightening and darkening do not cancel out for these cells but are combined nonlinearly over space.

More general analyses of whether receptive-field substructure beyond the center-surround shape is relevant under different stimulus scenarios typically require a model-based approach. The receptive field, often extended to contain a temporal dimension, is then used to filter the applied visual stimulus to predict responses for comparison with actual recorded data. In its most prominent model version, receptive field-based filtering is followed by a nonlinear transformation of the filtered signal, which can prevent negative predicted firing rates as well as implement thresholding and saturation of responses. This is known as the linear–nonlinear (LN) model (Figure 1C), which has become, in some respects, a standard model for retinal ganglion cell responses and forms the backbone of many data analysis approaches and extended computational models.

Yet, the receptive field-based LN model is often not a good predictor of responses for many ganglion cells, in particular when natural stimuli are considered. Under flashed natural images, responses of many ganglion cells in both mouse [17] and salamander [18] deviate from LN model predictions. Furthermore, these deviations systematically depend on the spatial contrast of the images inside the receptive field. Similarly, in the macaque retina, models based on receptive-field filtering fail to accurately predict responses of parasol cells to natural movies [19]. OFF parasol cells, in particular, are sensitive to the fine spatial structure of natural stimuli, which is not captured by the LN model [20].

Receptive-field subunits and their functional relevance

The primary model extension to go beyond receptive-field filtering and incorporate sensitivity to spatial structure within the receptive-field center is to subdivide the center into smaller subunits, giving rise to a 'subunit model'. Each subunit acts as an independent spatial filter, the signal of which is nonlinearly transformed (e.g., rectified) before summation into the integrated ganglion cell activation (Figure 1D). The subunits are generally believed to correspond to the bipolar cells that provide the excitatory input to ganglion cells, because this excitation alone already displays nonlinear spatial integration [21]. Signal transmission between bipolar and certain ganglion cells in the mouse retina has indeed been found to be nonlinear [22,23]. However, this does not exclude that, for other types of ganglion cells or under different illumination conditions, subunits might reflect other (e.g., amacrine) cells or correspond to larger groups of bipolar cells, perhaps electrically coupled [24]. In addition, nonlinearities at the input stage to bipolar cells [25] and in photoreceptor signaling [26,27] may contribute to nonlinear spatial integration upstream of bipolar-cell subunits.

Sensitivity to finely structured stimuli and motion signals

The subunit structure of receptive fields helps explain the different functional properties and specific computations of retinal ganglion cells [28]. Most prominently perhaps, rectified subunit signals mediate sensitivity to high spatial frequencies and to small objects below the scale of ganglion cell receptive-field centers [16,23,29]. This is achieved by communicating local changes in illumination (Figure 2) even without net luminance changes across the receptive field. The sensitivity to fine spatial patterns also leads to a joint encoding of luminance and spatial-contrast information [17,18,30] and may contribute to monitoring the proper focusing of images onto the retina [31].

Furthermore, subunits are particularly important for motion processing. Subtle object or texture motion that changes the illumination pattern over individual subunits can then trigger responses even when a linear receptive field would not be activated. This may contribute to the perception

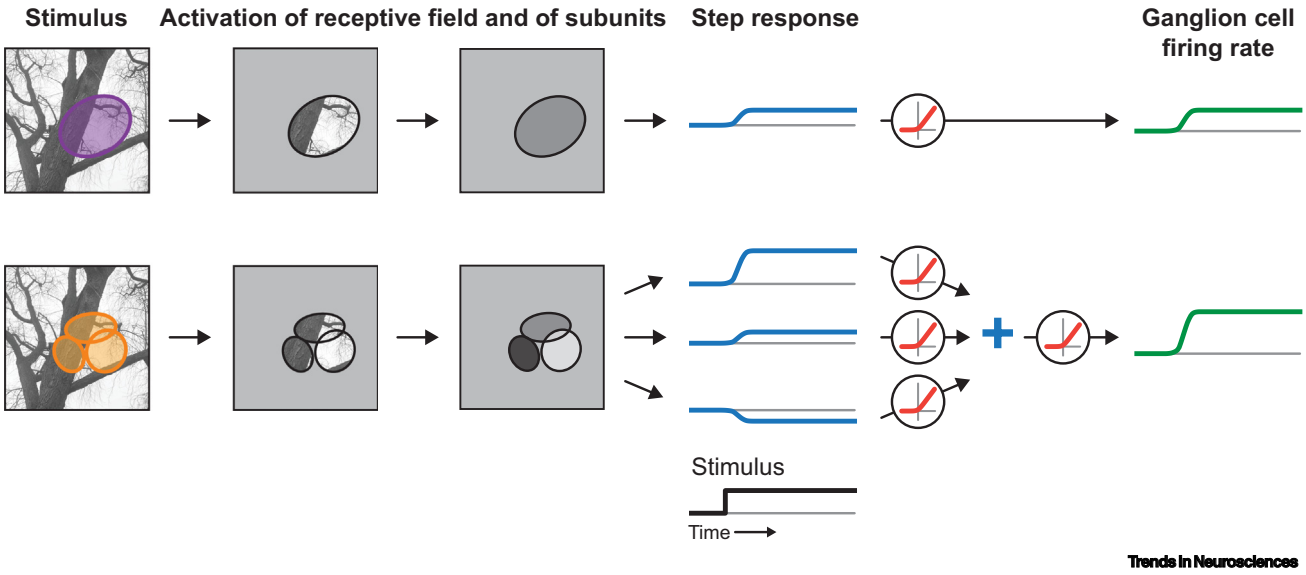

Figure 2. Schematic of ganglion cell processing with and without subunits. Top: linear–nonlinear (LN) model, with receptive field as a spatial filter. Bottom: subunit model, with multiple subunits as filters. The receptive field and the subunits evaluate an image in confined regions (left) to extract mean light-intensity signals. In the subunit model, the small size of the subunits can resolve spatial structure of the stimulus within the receptive field of the ganglion cell. The filtered signals generate activation signals, depicted here as step-response curves (blue, center), corresponding to image presentation after homogeneous background light intensity. Positive activation here corresponds to darkening in the spatial-filter region (OFF-cell responses). The subunit responses are rectified individually before summation, which here leads to a stronger ganglion cell response (green, right) than in the LN model.

of contrast-mediated motion signals ('second-order motion') [32] and to detecting subtle fixational eye movements [33,34]. Combining rectification of bipolar-cell subunit signals with gap-junction coupling between bipolar cells can further enhance sensitivity to specific motion signals [24,35], which is hypothesized to enhance the information about future locations of continuously moving objects [36]. Even for ON–OFF direction-selective ganglion cells, where much of the computation of directionality occurs in the presynaptic starburst amacrine cells, the crucial integration of excitatory and inhibitory inputs occurs locally along the dendritic tree of the ganglion cell, providing a subunit architecture to the motion processing [37,38].

Spatial nonlinearities of inhibitory signals

For some ganglion cells, the subunit-mediated sensitivity to texture motion also pertains to inhibitory signals from the receptive-field periphery. This can lead to response suppression under global image motion and thereby specific sensitivity to local object motion in object motion-sensitive ganglion cells described in the salamander and rabbit retina [9] and in the so-called 'W3 ganglion cells' of the mouse [10]. For other cells, the motion sensitivity in the receptive-field periphery results in disinhibition and increased responsiveness after a global image shift [39]. This exemplifies that nonlinear spatial integration and subunit models also relate to inhibitory signaling in the retina and thereby have a role in gating visual information.

Spatial integration in the suppressive surround of salamander ganglion cells appears to be similarly nonlinear as in the receptive-field center, albeit on a larger spatial scale [40], suggesting that amacrine cells might provide relevant nonlinearities. In mouse retina, the possibility to optogenetically stimulate specific amacrine cell types was recently used to directly probe signal transmission to ganglion cells. This revealed strong rectification of the functionally crucial inhibition that direction-selective ganglion cells receive from starburst amacrine cells [41,42], whereas other amacrine cells were found to transmit signals linearly [43,44].

With inhibition being prevalent not only in the receptive-field surround, but also in the center, one may wonder why the subunit models discussed above are often considered without inhibitory elements. One answer is that local inhibitory effects can be part of the subunits, so that the latter correspond to bipolar cells together with local modulations by inhibitory amacrine and horizontal cells. These may add a surround component to the subunits, modulate their temporal signaling, affect the nonlinearity of signal transmission to the ganglion cell, or shape the adaptation characteristics. The view of subunits as 'effective bipolar cells', the processing of which also comprises inhibitory effects, is reinforced by the finding that many local inhibitory interactions occur directly at the bipolar cell synaptic terminal [45,46], where they fundamentally shape the temporal and spatial characteristics of the subunit signal received by the ganglion cell [47].

An example of local inhibition fundamentally shaping the functional characteristics of a ganglion cell is provided by the widely studied transient OFF alpha ganglion cell of the mouse retina. Even though flashed natural images expose little deviation from linear spatial integration for these cells [17], stimuli with more complex spatiotemporal structure can reveal nonlinear local processing by crossover inhibition, which couples the ON and OFF signaling pathways in the retina [48,49]. Here, it supplies ON-pathway inhibition under light increments to an OFF-type ganglion cell. This is thought to let the cell differentiate between looming and translatory motion of a dark object, given that responses are suppressed for the latter because of inhibition triggered by the brightening at the trailing edge of the object [8]. For saccade-like image transitions, the local crossover inhibition, in conjunction with serial inhibition, has also been suggested to underlie a particular sensitivity to recurring image patterns, which has been observed in these cells [50].

Accessing the subunits for studying retinal computations

The examples discussed above show how inhibition can suppress or gate subunit signals in a selective manner, implementing computations that are not captured by the shape of the center-surround receptive field. This suggests viewing the layout of bipolar-cell subunits as scaffolding for local signal processing, setting up the spatial layout for operating on the visual input. The types and dynamics of the operations, from simple nonlinear transformations to complex inhibitory interactions, then determine the computation performed by the ganglion cell and, thus, its potential function for visual processing. Therefore, access to the subunit layout could be an important step for analyzing the details of retinal computations either through model-based data analysis or by targeting individual subunits for stimulation in experiments.

However, experimental complications arise from the relative inaccessibility of bipolar cells. These cells are buried in the middle layers of the retina; they are mostly nonspiking and receive much of their functionally relevant inputs at their synaptic terminals. Some progress has been made by applying fluorescent reporters to monitor vesicle fusion [51,52] or glutamate release [53] at the bipolar cell terminal. This has allowed observations of rectified transmitter release from certain bipolar cells [22] and characterizations of the functional diversity of bipolar cell types and of the role of inhibition in shaping this diversity [47,54].

Nonetheless, for analyzing how ganglion cells integrate bipolar cell input, alternatives to direct measurements of bipolar cell signals are desirable. A promising approach arises from currently developed methods for model-based inference of bipolar-cell subunits via the relatively easy-to-record ganglion cell activity. By aiming at identifying bipolar cell receptive fields through their effects on the output of the retina, this approach may provide a 'virtual microscope' for studying the layout of bipolar cells and their functional connections to ganglion cells.

Inference of subunits with multilayered models

LNLN models

Perhaps the most straightforward approach to identifying and characterizing receptive-field subunits is to set up a subunit model and fit the parameters to recorded data. The model typically comprises two stacked layers of LN model-like components, one for the subunits and one for the subsequent signal summation and final transformation (Figure 1D), and, therefore, is also called 'LNLN model'. Although the general idea of subunit models dates back to the original observation of nonlinear spatial integration [16], methods for identifying concrete subunit layouts from experimental data emerged only relatively recently. One approach is to use anatomical knowledge to structure the model, for example, by determining subunit positions and weights according to the reconstructed dendritic tree of a ganglion cell. This has been used to demonstrate that the concrete subunit layout matters for accurate response predictions [23]. Similarly, prior assessment of the spatial layout of the relevant photoreceptors, as is possible in the peripheral primate retina [55], has been exploited to constrain LNLN models [56].

However, without knowledge of the anatomical details for a given cell, the difficulty of obtaining a fitted LNLN model lies in finding a workable parameterization, particularly since the number and locations of subunits are unknown. To reduce model complexity, some studies have focused on a single spatial dimension by applying stripe-like visual stimuli and thereby successfully fitted LNLN models to ganglion cells recorded from salamander retina [57,58]. The extracted subunits (Figure 3A) displayed response characteristics and receptive fields similar to bipolar cells that had been independently recorded [57,58]. These findings support the notion that fitting subunit models can be used to infer bipolar cell properties. If the number of parameters is further reduced by assuming regular spacing and identical shapes of subunits, the LNLN model can even be extended to include adaptation-like feedback components [58]. A different application of LNLN model fits is to separate the temporal characteristics of ON-type and OFF-type inputs or of excitatory and suppressive receptive-field components, as has been shown for mouse retinal ganglion cells [59]. Combining the modeling of such different input types with a spatial layout of subunits will be an exciting, yet challenging future direction.

Convolutional neural networks

Fitting complex parameterized models received a push with the advent of deep learning-inspired artificial neural networks. In particular, convolutional neural networks enjoy current popularity, owing to their versatility and computational efficiency. The convolutional operations in these networks (applying the same filtering at all locations of the input space) are reminiscent of how an individual bipolar cell type (or another retinal neuron type) covers visual space with its receptive fields, applying nearly the same signal processing at different spatial locations. Thus, not surprisingly, retinal ganglion cells were among the first applications of convolutional neural networks to perform data-driven circuit modeling [60–62].

Using two convolutional layers, these networks can outperform standard receptive-field models, such as the LN model, in predicting responses of salamander ganglion cells to new stimuli, because the multilayered structure captures essential nonlinear substructure of ganglion cell receptive fields. Moreover, hidden units can match certain properties of bipolar and amacrine cells, such as localized center-surround receptive fields [60,61] (Figure 3B). When trained on salamander ganglion cell responses to natural stimuli, the models are also able to reproduce specific response features beyond what can be explained by spatiotemporal filtering [60,62]. However, despite these agreements, the derived networks have so far remained somewhat abstract, with no direct correspondence to circuit structure or subunit layout.

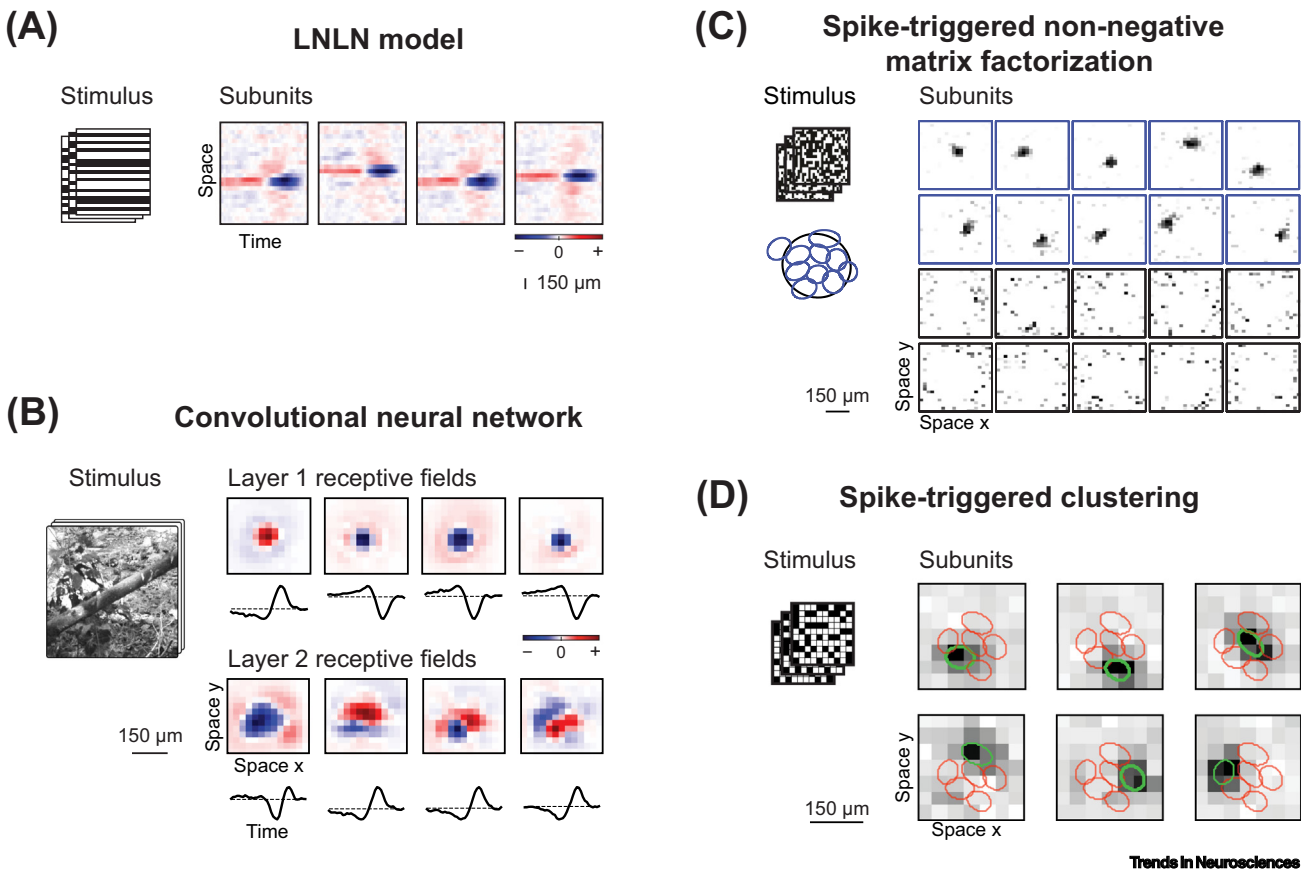

Figure 3. Examples of computationally inferred subunits from different methods. (A) Spatiotemporal subunits with one spatial dimension obtained from fitting a linear–nonlinear–linear–nonlinear (LNLN) model to recorded salamander ganglion cell responses under stimulation with flickering stripes. Subunits in the example are OFF-type, with localized preference for negative contrast (blue), preceded and flanked by preference for positive contrast (red), corresponding to biphasic temporal filtering and a spatial antagonistic surround, respectively. (B) Spatiotemporal subunits, corresponding to receptive fields of convolutional filters from a two-layer convolutional neural network, fitted to recorded salamander ganglion cell responses under natural stimulation. Receptive fields are separated into a spatial component (red/blue: positive/negative contrast preference) and a temporal sensitivity trace below. (C) Spatial subunits obtained by analyzing salamander ganglion cell responses to spatiotemporal white noise with spike-triggered non-negative matrix factorization. The obtained spatial modules (white indicates zero; darker shades indicate more positive entries) can be separated into true subunits (blue module frames) and noise modules (black frames) according to the spatial autocorrelation of the modules. The recovered subunits are fitted by 2D Gaussian profiles (blue ellipses, corresponding to 1.5-sigma contours), which roughly tile the Gaussian fit to the receptive field (black ellipse). (D) Spatial subunits obtained by spike-triggered clustering, here with six clusters, with data recorded from a macaque ganglion cell under spatiotemporal white-noise stimulation. Adapted from [57] (A), [60] (B), [68] (C), and [70] (D).

Inference of subunits by statistical analyses of successful stimuli

Structure of spike-triggered stimulus ensembles

An alternative to fitting a complete subunit model is to extract subunits from statistical analyses of the stimulus–response relationship of a ganglion cell. The basic idea is to record ganglion cell responses to spatially structured stimuli and investigate which spatial patterns make the cell fire. The set of these successful stimuli is often called the spike-triggered stimulus ensemble (STE), and the search for subunits then amounts to identifying the underlying structure of the STE (Figure 4). This often takes the form of a dimensionality-reduction analysis [63], for which a wide arsenal of general techniques is already available. However, to avoid confounding structure in the STE from prior correlations in the applied visual stimuli, these techniques typically rely on stimuli with white-noise statistics, which may not drive cells strongly if a finely structured layout is used, thus requiring long recordings.

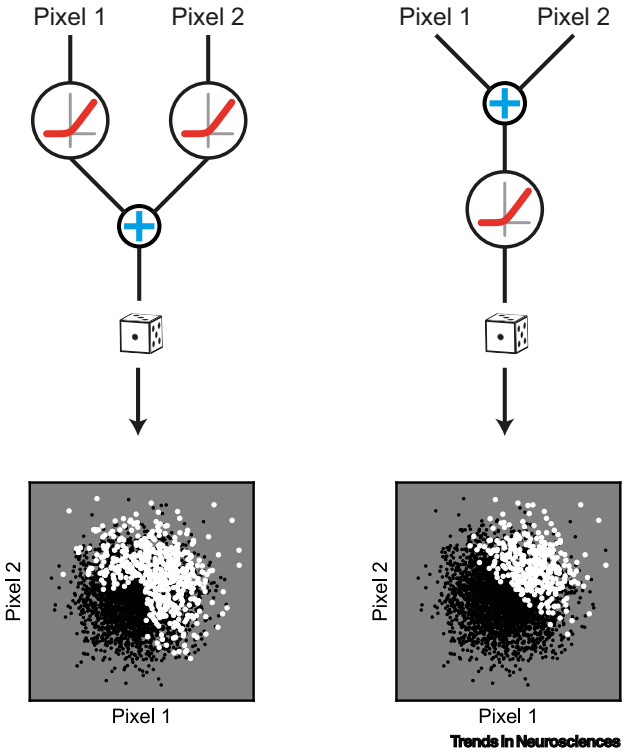

Figure 4. Effect of subunits on the spike-triggered stimulus ensemble (STE). Whether two stimulus pixels are in the same subunit influences the joint distribution of their values in the STE, as exemplified here by two simple models (top). Stimuli are normally distributed contrast values of two pixels. In the first model (left), the values are individually rectified and then summed into a spike probability, representing pixels in distinct subunits. In the second model (right), the two pixel values are summed before rectification, representing pixels within a subunit. Spikes are generated by a binary random process. For each model, scatter plots at the bottom display the set of applied stimuli (black dots) and stimuli that generated spikes, that is, the STE (white dots). The two STEs exhibit different distributions, which forms the basis for the inference of subunits with statistical methods. In particular, the STE of the first model contains a wider range of stimuli with opposing contrast values.

One common technique for extracting relevant dimensions of the STE is spike-triggered covariance analysis, which is essentially a form of principal component analysis of the STE [64,65]. However, when applied to data recorded from the salamander retina under spatiotemporal white noise, the resulting filters are generally not localized subunits, but rather span the entire receptive field, reminiscent of a Fourier decomposition [57,66–68]. Appropriate mixing of these filters might help recover localized components [67], but may be limited by noise in the originally identified filters and, to our knowledge, this has not yet been thoroughly tested on experimental data.

Spike-triggered non-negative matrix factorization and clustering

A more direct identification of subunits can be achieved by analyzing the STE with spike-triggered non-negative matrix factorization (STNMF) [68]. Here, multiple spatial filters are extracted under the constraint of having no negative elements. Non-negative matrix factorization is known to provide parts-based representations of complex data [69], which helps here to uncover local structure in receptive fields. In fact, subunit filters obtained for salamander ganglion cells were found to be localized (Figure 3C) and to match the receptive fields of simultaneously recorded bipolar cells [68], supporting the idea that bipolar cell layouts can be inferred from ganglion cell recordings. Note, however, that subunits are derived with no negative entries, and any structure with opposing sign, such as an antagonistic surround, is not recovered by this method.

A related approach is to find the elementary stimulus patterns that contribute to generating spikes by clustering analyses of the STE. The idea is that activation of one or potentially a few subunits must have been sufficiently strong to trigger a spike, while other subunits remained subthreshold. The centroids of a soft clustering of the STE from macaque OFF parasol cells indeed display

spatially localized structure as expected for receptive-field subunits [70] (Figure 3D). Moreover, the spike clustering can be viewed as a step in the maximum likelihood fit of an LNLN model and, thus, directly yields a subunit model for predicting responses to spatially structured stimuli.

Regardless of the method, inferred subunit layouts so far typically show subunits that more or less tile the receptive field with little overlap. This is consistent with excitatory inputs dominated by a single type of bipolar cell, as appears to be the case for some ganglion cells [71]. Yet, it raises the question whether layouts with two (or more) bipolar cell types, each arranged in its own mosaic pattern, could also be handled. Analyses of model simulations suggest that STNMF may, in principle, be able to extract subunits with considerable overlap [68], but evidence for this from experimental data is still lacking. One possibility is that extracted subunits reflect the bipolar cell type that provides the strongest input. It might then be interesting to explore whether stimuli can be modified (e.g., via ambient light intensity or spatial, temporal, or chromatic characteristics) to switch the dominant input. This might allow reconstructing different bipolar cell mosaics in different stimulation contexts.

Temporal dynamics of subunits

Functions and mechanisms of local and global adaptation

The fact that ganglion cells integrate over many bipolar cells not only breaks up spatial integration into subunits, but is also important for understanding the functional consequences of adaptation. For example, retinal ganglion cells adapt their sensitivity to ambient light intensity and to visual contrast. Conceptually, adaptation could occur locally, at the level of individual subunits before their signals are integrated by the ganglion cell, or globally after stimulus integration, at the level of the entire ganglion cell receptive field (Figure 5). The functional relevance becomes clear, for example, when considering a small object somewhere in the receptive field. Global, but not local, adaptation then reduces sensitivity for a second object irrespective of its location. Furthermore, for a moving object, global adaptation is already triggered when entering the receptive field and thereby reduces responses to the subsequent path through the receptive field, a mechanism that has been suggested to counteract temporal delays of signal transduction for moving objects ('motion anticipation'; Figure 5A) [72]. By contrast, local adaptation of individual subunits maintains sensitivity at other locations within the receptive field or for novel objects moving along different paths (Figure 5B,C). Thus, whether adaptation occurs globally on the spatial scale of the entire receptive field or locally for individual subunits is crucial for understanding the function of a ganglion cell in response to dynamic spatial stimuli.

For retinal contrast adaptation, mechanisms that could support local or global adaptation have both been identified. A primary contribution to contrast adaptation is thought to occur at the bipolar-to-ganglion cell synapse, where synaptic gain is altered by synaptic depression [73,74] and inhibition onto the synaptic terminal [75,76]. Computational models of vesicle depletion or presynaptic inhibition have successfully captured much of the adaptive dynamics of excitation in certain ganglion cells [76,77]. Owing to the local origin at the synaptic terminal, this adaptation might be thought to be local and restricted to individual subunits. Yet, presynaptic inhibition need not be locally confined and could act on multiple subunits. In addition, even local synaptic depression may influence sensitivity throughout the receptive field of the ganglion cell if the reduction in vesicle release leads to a postsynaptic hyperpolarization that affects the entire ganglion cell. Thus, it remains an open question whether gain changes at the bipolar cell terminal contribute to local or global adaptation. A clearer contribution to local adaptation comes from contrast-induced gain changes in the membrane potential of bipolar cells themselves. This has been observed in the salamander retina [78], although not for all bipolar cells [79]. Conversely, unambiguous global adaptation may follow from mechanisms triggered by the activity of the

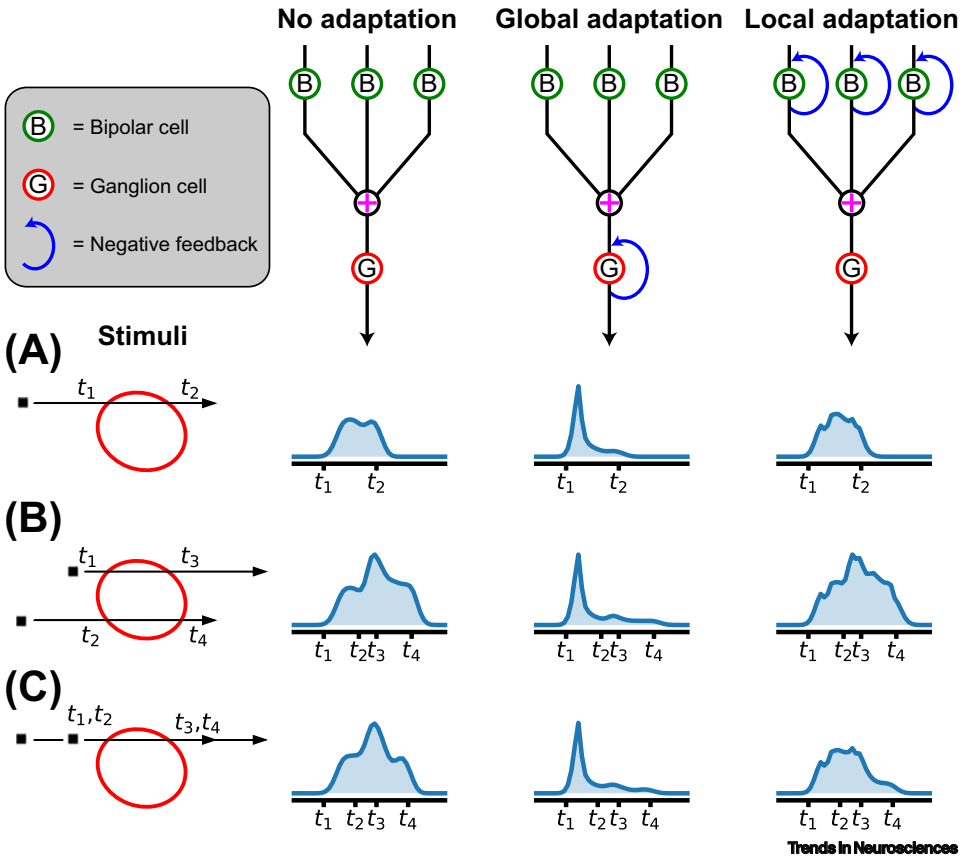

Figure 5. Functional consequences of local and global adaptation. Simulated responses to different motion stimuli (left) of three subunit models with no adaptation, global adaptation, and local adaptation (top). All models simulate multiple bipolar cells (only three are depicted) connected to a ganglion cell. Bipolar cells are simulated by applying a spatial Gaussian filter and a monophasic temporal filter, followed by a rectification. Ganglion cells sum the results and apply another rectification yielding a firing rate as output. Global adaptation is implemented as multiplicative feedback by applying an exponentially decaying filter to the model output, followed by a sigmoidal nonlinearity. Local adaptation occurs through analogous multiplicative feedback on the bipolar-cell level. Simulated firing rates (blue) are shown in response to objects moving through the receptive field, with the actual configurations shown on the left. Black squares and arrows show objects and their trajectories. Red ellipses correspond to the 1.5-sigma contour of the receptive field. Times t_1 to t_4 indicate when the objects enter the first or leave the last subunit, respectively, as indicated in the plots. When a single object moves through the receptive field (A), global adaptation leads to a rapid response decay, causing the peak firing rate to occur well before the object is halfway through the receptive field. This has been associated with the phenomenon of motion anticipation [72] and is not apparent for the models with no adaptation or with local adaptation. By contrast, when two objects successively move through the receptive field, local adaptation leads to distinct responses, depending on whether the two objects move along different trajectories (B) or along the same trajectory (C), yielding reduced response amplitude for the latter. By contrast, the models with no or global adaptation each display similar responses for these two scenarios.

ganglion cell itself, such as inactivation of sodium channels [80] or recruitment of potassium currents [81]. Which mechanisms dominate for a given cell and whether this depends on stimulus context, such as light level, is largely unknown.

Experimental tests of local and global contrast adaptation

Functionally, only a few studies have probed the spatial scope of contrast adaptation so far. A study of cross-adaptation in rabbit retina, which used switching stimulation from one small location to another within the receptive field of a recorded ganglion cell, showed transient increases in

firing rate after each switch [82]. This can be interpreted as being caused by local adaptation, since the new stimulus location after the switch is not yet adapted. However, the fast timescale of the transient firing rate increases might also be explained without any adaptation. Instead, nonlinear spatial integration can boost activity during the brief transition time when activation from one location is ramping up while activation from the other has not yet fully decayed [83].

Local adaptation can also be observed in the object motion-sensitive cells discussed above, which adapt to continuous object motion, but transiently regain sensitivity when the object motion is detected by different regions inside the receptive-field center [84]. Other studies of ganglion cells in salamander [83] and mouse retina [85] have observed that changes in temporal filtering and sensitivity following locally restricted switches in visual contrast can, for different cells, occur either locally, confined to the region of the contrast switch, or globally over the entire receptive field (Figure 6). Thus, as with linear and nonlinear spatial integration, the occurrence of local and global contrast adaptation likely depends on cell type.

A different example of spatially structured adaptation has been found in the context of contrast sensitization [52,86,87], which describes the transient sensitivity increase after exposure to high contrast (as opposed to the more classical adaptive sensitivity decrease). For some sensitizing ganglion cells, the occurrence of either sensitization or classical adaptation depends on stimulus location. Stimuli near the midpoint of the receptive-field center cause adaptation and more peripheral stimuli elicit sensitization, which is thought to result from adaptation in excitatory as well as inhibitory subunits within the receptive-field center [88].

Concluding remarks

Retinal function goes well beyond stimulus filtering by the center-surround receptive fields of ganglion cells. Nonlinear stimulus integration within the receptive field, local signal gating by inhibitory interactions, and local adaptation endow the receptive field with a substructure that shapes visual stimulus encoding. The receptive-field substructure also provides the substrate for different specific computations in the retina, which rely on local analyses of light patterns. Thus, investigating the substructure of ganglion cell receptive fields is important for capturing how the retina encodes natural stimuli and for understanding the circuit implementation of specific retinal functions.

Reliable inference of ganglion cell subunits may prove a valuable tool for circuit analyses in the retina. It provides indirect access to the properties of retinal bipolar cells, supplying information about their receptive-field attributes and nonlinear transformations as well as about their functional connectivity to ganglion cells. For example, obtaining the same subunit contour from two or more ganglion cells indicates shared input from the same bipolar cells [68]. Furthermore, identified subunit layouts will be useful in experiments to optimally place probe stimuli to further study the nonlinearities and dynamics of stimulus encoding and in computational analyses to construct and constrain encoding models.

As methods for analyzing receptive-field substructure mature, it will be interesting to compare their applicability to retinas from different animal models and to evaluate whether they reveal differences in retinal circuit organization. Given that many aspects of the basic retinal layout are similar across vertebrates, one should expect that the general methodology discussed here is not restricted to particular species. For example, nonlinear spatial integration is a general observation, and evidence that receptive-field subunits correspond to bipolar cells has also come from various vertebrates, including guinea pig [21], salamander [68], mouse [23], and macaque [56]. Nonetheless, details differ. For example, in the primate retina, the fovea is a region of specialized connectivity and signal processing characteristics [89,90]. Beyond the fovea, nonlinear spatial

Outstanding questions

How do the different methods developed for identifying subunits of retinal ganglion cells compare? Can different methods be combined to arrive at a consensus of inferred subunit layouts? How robust are the different methods across different ganglion cell types, different animal species, and different illumination conditions?

What is missing in subunit models with nonlinear spatial integration to fully capture ganglion cell encoding of flashed natural images, for which temporal dynamics may be ignored?

The implementation of specific computations by retinal circuits, such as the detection of different types of motion signals, is often explained through conceptual, proof-of-principle models. How can such conceptual models be combined with concrete subunit layouts of individual cells to explore the circuit–function connection in more detail?

Nonlinear signal transformations may already occur in photoreceptors or at the inputs to bipolar cells. How important are these early nonlinearities for stimulus encoding by ganglion cells? How could they be incorporated into stimulus–response models?

Are subunits a relevant concept for the ganglion cell receptive-field surround? How could one extract information about surround subunits from experimental data? Would these subunits correspond to bipolar cells or to amacrine cells?

What makes some ganglion cells adapt locally to visual contrast and others globally? What mechanisms are involved and what specific functional consequences arise?

How can inhibitory interactions be integrated into subunit models and what would be the ways to obtain interaction parameters from experimental data?

After fitting abstract network models, such as convolutional neural networks, to retinal data, can the obtained network structure be related to actual circuit components of the retina?

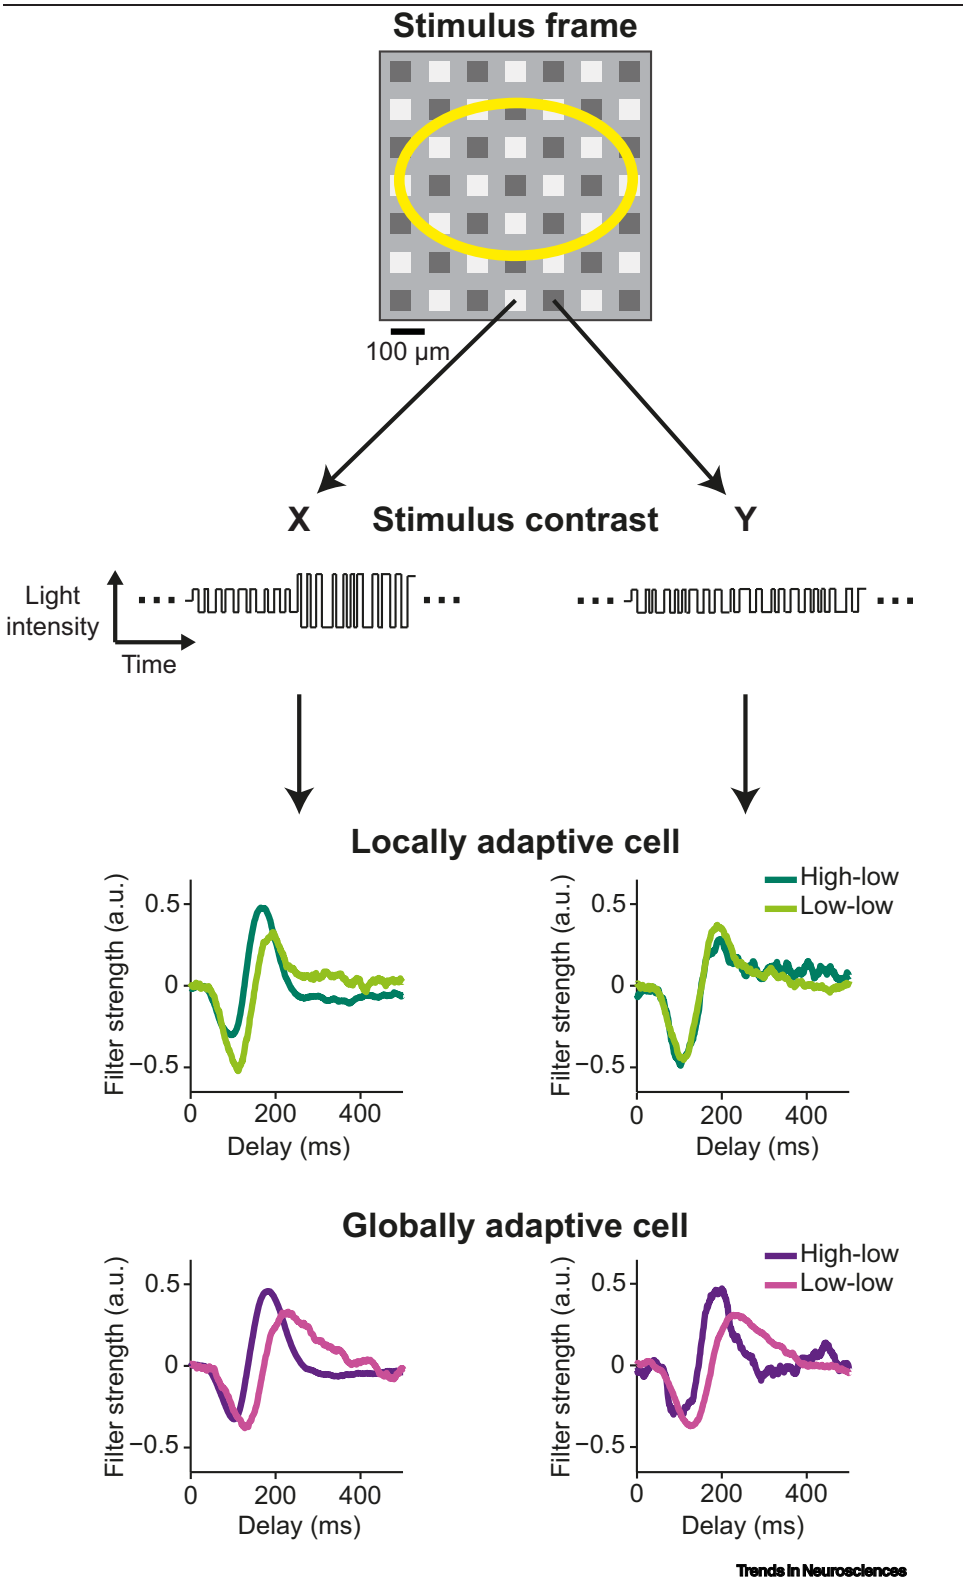

(See figure legend at the bottom of the next page.)

integration appears to be the norm in the salamander retina [91], whereas cell types with more linear characteristics can be found in mammalian retinas, as shown in cat [16], mouse [17], and macaque [20,92]. Similarly, the spatial scope of contrast adaptation displays greater diversity in mouse [85] than in salamander [83]. Furthermore, it seems likely that the basic scaffolding of bipolar-cell subunits supports different species-specific functions through different nonlinear transformations and inhibitory interactions.

Moreover, questions of receptive-field substructure also pertain to other visual areas or sensory systems. For example, subunit models have been applied successfully to primary visual cortex [93–95] and to motion-processing areas of the primate visual pathway [96,97]. In the auditory system, complex dependencies of neural responses on the spectrotemporal structure of acoustic stimuli have led to related models with nonlinear substructure in the spectrotemporal receptive fields of the neurons [98–100]. Subunit identification here might be a promising endeavor to relate this substructure to neural circuitry.

Despite the recent advances discussed here, the challenge of inferring receptive-field substructure in the retina is still far from solved (see Outstanding questions). Future developments may profit from the conceptual overlap that exists with current efforts to analyze large neural ensembles. Finding structure in high-dimensional population activity mirrors finding structure in spike-generating visual stimuli. Thus, similar methods for identifying a manageable number of latent variables may apply, such as dimensionality-reduction techniques [63,101]. For example, matrix factorization approaches are used equivalently for both types of problem [68,96,102–104].

In addition, the rapid developments in machine learning and artificial neural networks promise new approaches to study neural signal processing [105,106]. An enticing question here is how these approaches can be harnessed to go beyond subunits of ganglion cells and study inhibitory interactions in the retina. This could help tackle one of the biggest remaining mysteries about the retina, namely the functions and interactions of the vast set of inhibitory amacrine cells [2,107]. Incorporating such complexity into generic ganglion cell models, amenable for data-driven parameter optimization, will be a considerable challenge. However, recent computational advances may help deal with the resulting model intricacies.

Incorporating large populations of recorded ganglion cells into a single model with shared presynaptic circuitry, for example, will help constrain parameter estimation [61,108]. Further constraints can come from detailed anatomical knowledge, such as the morphology of individual cells [23] or the statistics of neuronal connectivity [109–111]. Derived neural network models can then be used for *in silico* analyses, for example, by extracting maximally effective stimuli [112] or by generating mechanistic hypotheses for specific computations through model reduction techniques [62]. For models of arbitrary structure, combining model simulations with machine-learning algorithms can help with parameter optimization [113]. These and other ongoing developments of computational approaches to infer structural and functional components of neural networks promise new insights into the nonlinear and dynamic substructure of retinal receptive fields.

Figure 6. Measurements of local and global contrast adaptation. Shown are the temporal stimulus filters of a locally adaptive and a globally adaptive ganglion cell, recorded in mouse retina under local switches in temporal contrast. The stimulus comprised locations X and Y, each stimulated homogeneously and independently by binary white noise (top: sample stimulus frame, yellow ellipse sketches an exemplary 3-sigma receptive-field contour of a ganglion cell). At locations Y, contrast remained constant, whereas at locations X, contrast changed every 40 s between a low (low-low condition) and a high (high-low condition) value. For the sample cell with local adaptation, the temporal filter for locations X changed with the contrast at X (left), but the filter for Y remained unaffected (right). By contrast, for the sample cell with global adaptation, the filters for both sets of locations displayed similar adaptive changes. Adapted from [85].

Acknowledgments

This work was supported by the Deutsche Forschungsgemeinschaft [DFG, German Research Foundation; project IDs 154113120 (SFB 889, project C01) and 432680300 (SFB 1456, project B05)] and by the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Programme (grant agreement number 724822).

Declaration of interests

The authors declare no competing interests.

References

1. Schwartz, G. (2021) *Retinal Computation*, Academic Press
2. Yan, W. *et al.* (2020) Mouse retinal cell atlas: molecular identification of over sixty amacrine cell types. *J. Neurosci.* 40, 5177–5195
3. Baden, T. *et al.* (2016) The functional diversity of retinal ganglion cells in the mouse. *Nature* 529, 345–350
4. Grünert, U. and Martin, P.R. (2020) Cell types and cell circuits in human and non-human primate retina. *Prog. Retin. Eye Res.* 78, 100844
5. Goetz, J. *et al.* (2021) Unified classification of mouse retinal ganglion cells using function, morphology, and gene expression. *bioRxiv* Published online June 11, 2021. https://doi.org/10.1101/2021.06.10.447922
6. Barlow, H.B. *et al.* (1964) Retinal ganglion cells responding selectively to direction and speed of image motion in the rabbit. *J. Physiol.* 173, 377–407
7. Wei, W. (2018) Neural mechanisms of motion processing in the mammalian retina. *Annu. Rev. Vis. Sci.* 4, 165–192
8. Münch, T.A. *et al.* (2009) Approach sensitivity in the retina processed by a multifunctional neural circuit. *Nat. Neurosci.* 12, 1308–1316
9. Ölveczky, B.P. *et al.* (2003) Segregation of object and background motion in the retina. *Nature* 423, 401–408
10. Zhang, Y. *et al.* (2012) The most numerous ganglion cell type of the mouse retina is a selective feature detector. *Proc. Natl. Acad. Sci. U. S. A.* 109, E2391–E2398
11. Kuffler, S.W. (1953) Discharge patterns and functional organization of mammalian retina. *J. Neurophysiol.* 16, 37–68
12. Baden, T. and Osorio, D. (2019) The retinal basis of vertebrate color vision. *Annu. Rev. Vis. Sci.* 5, 177–200
13. Demb, J.B. and Singer, J.H. (2015) Functional circuitry of the retina. *Annu. Rev. Vis. Sci.* 1, 263–289
14. Rodieck, R.W. (1965) Quantitative analysis of cat retinal ganglion cell response to visual stimuli. *Vis. Res.* 5, 583–601
15. Atick, J.J. and Redlich, A.N. (1992) What does the retina know about natural scenes? *Neural Comput.* 4, 196–210
16. Enroth-Cugell, C. and Robson, J.G. (1966) The contrast sensitivity of retinal ganglion cells of the cat. *J. Physiol.* 187, 517–552
17. Karamanlis, D. and Gollisch, T. (2021) Nonlinear spatial integration underlies the diversity of retinal ganglion cell responses to natural images. *J. Neurosci.* 41, 3479–3498
18. Liu, J.K. *et al.* (2022) Simple model for encoding natural images by retinal ganglion cells with nonlinear spatial integration. *PLoS Comput. Biol.* 18, e1009925
19. Heitman, A. *et al.* (2016) Testing pseudo-linear models of responses to natural scenes in primate retina. *bioRxiv* Published online December 7, 2016. https://doi.org/10.1101/045336
20. Turner, M.H. and Rieke, F. (2016) Synaptic rectification controls nonlinear spatial integration of natural visual inputs. *Neuron* 90, 1257–1271
21. Demb, J.B. *et al.* (2001) Bipolar cells contribute to nonlinear spatial summation in the brisk-transient (Y) ganglion cell in mammalian retina. *J. Neurosci.* 21, 7447–7454
22. Borghuis, B.G. *et al.* (2013) Two-photon imaging of nonlinear glutamate release dynamics at bipolar cell synapses in the mouse retina. *J. Neurosci.* 33, 10972–10985
23. Schwartz, G.W. *et al.* (2012) The spatial structure of a nonlinear receptive field. *Nat. Neurosci.* 15, 1572–1580
24. Kuo, S.P. *et al.* (2016) Nonlinear spatiotemporal integration by electrical and chemical synapses in the retina. *Neuron* 90, 320–332
25. Schreyer, H.M. and Gollisch, T. (2021) Nonlinear spatial integration in retinal bipolar cells shapes the encoding of artificial and natural stimuli. *Neuron* 109, 1692–1706
26. Endeman, D. and Kamermans, M. (2010) Cones perform a non-linear transformation on natural stimuli. *J. Physiol.* 588, 435–446
27. Angueyra, J.M. *et al.* (2022) Predicting and manipulating cone responses to naturalistic inputs. *J. Neurosci.* 42, 1254–1274
28. Gollisch, T. and Meister, M. (2010) Eye smarter than scientists believed: neural computations in circuits of the retina. *Neuron* 65, 150–164
29. Victor, J.D. and Shapley, R.M. (1979) The nonlinear pathway of Y ganglion cells in the cat retina. *J. Gen. Physiol.* 74, 671–689
30. Yu, Z. *et al.* (2022) Adaptation in cone photoreceptors contributes to an unexpected insensitivity of primate On parasol retinal ganglion cells to spatial structure in natural images. *Elife* 11, e70611
31. Mani, A. and Schwartz, G.W. (2017) Circuit mechanisms of a retinal ganglion cell with stimulus-dependent response latency and activation beyond its dendrites. *Curr. Biol.* 27, 471–482
32. Demb, J.B. *et al.* (2001) Cellular basis for the response to second-order motion cues in Y retinal ganglion cells. *Neuron* 32, 711–721
33. Greene, G. *et al.* (2016) Non-linear retinal processing supports invariance during fixational eye movements. *Vis. Res.* 118, 158–170
34. Hennig, M.H. and Wörgötter, F. (2007) Effects of fixational eye movements on retinal ganglion cell responses: a modelling study. *Front. Comput. Neurosci.* 1, 69–84
35. Manookin, M.B. *et al.* (2018) Neural mechanisms mediating motion sensitivity in parasol ganglion cells of the primate retina. *Neuron* 97, 1327–1340
36. Liu, B. *et al.* (2021) Predictive encoding of motion begins in the primate retina. *Nat. Neurosci.* 24, 1280–1291
37. Jain, V. *et al.* (2020) The functional organization of excitation and inhibition in the dendrites of mouse direction-selective ganglion cells. *Elife* 9, e52949
38. Murphy-Baum, B.L. and Awatramani, G.B. (2022) Parallel processing in active dendrites during periods of intense spiking activity. *Cell Rep.* 38, 110412
39. Jadzinsky, P.D. and Baccus, S.A. (2015) Synchronized amplification of local information transmission by peripheral retinal input. *Elife* 4, e09266
40. Takeshita, D. and Gollisch, T. (2014) Nonlinear spatial integration in the receptive field surround of retinal ganglion cells. *J. Neurosci.* 34, 7548–7561
41. Pottackal, J. *et al.* (2020) Receptoral mechanisms for fast cholinergic transmission in direction-selective retinal circuitry. *Front. Cell. Neurosci.* 14, 604163
42. Pottackal, J. *et al.* (2021) Computational and molecular properties of starburst amacrine cell synapses differ with postsynaptic cell type. *Front. Cell. Neurosci.* 15, 660773
43. Pottackal, J. *et al.* (2021) Photoreceptive ganglion cells drive circuits for local inhibition in the mouse retina. *J. Neurosci.* 41, 1489–1504
44. Park, S.J.H. *et al.* (2018) Convergence and divergence of CRH amacrine cells in mouse retinal circuitry. *J. Neurosci.* 38, 3753–3766
45. Asari, H. and Meister, M. (2012) Divergence of visual channels in the inner retina. *Nat. Neurosci.* 15, 1581–1589
46. Zaghloul, K.A. *et al.* (2007) Functional circuitry for peripheral suppression in Mammalian Y-type retinal ganglion cells. *J. Neurophysiol.* 97, 4327–4340

47. Franke, K. *et al.* (2017) Inhibition decorrelates visual feature representations in the inner retina. *Nature* 542, 439–444
48. Werblin, F.S. (2010) Six different roles for crossover inhibition in the retina: correcting the nonlinearities of synaptic transmission. *Vis. Neurosci.* 27, 1–8
49. Manookin, M.B. *et al.* (2008) Disinhibition combines with excitation to extend the operating range of the OFF visual pathway in daylight. *J. Neurosci.* 28, 4136–4150
50. Krishnamoorthy, V. *et al.* (2017) Sensitivity to image recurrence across eye-movement-like image transitions through local serial inhibition in the retina. *Elife* 6, e22431
51. Odermatt, B. *et al.* (2012) Encoding of luminance and contrast by linear and nonlinear synapses in the retina. *Neuron* 73, 758–773
52. Nikolaev, A. *et al.* (2013) Synaptic mechanisms of adaptation and sensitization in the retina. *Nat. Neurosci.* 16, 934–941
53. Marvin, J.S. *et al.* (2013) An optimized fluorescent probe for visualizing glutamate neurotransmission. *Nat. Methods* 10, 162–170
54. Zhao, Z. *et al.* (2020) The temporal structure of the inner retina at a single glance. *Sci. Rep.* 10, 4399
55. Field, G.D. *et al.* (2010) Functional connectivity in the retina at the resolution of photoreceptors. *Nature* 467, 673–677
56. Freeman, J. *et al.* (2015) Mapping nonlinear receptive field structure in primate retina at single cone resolution. *Elife* 4, e05241
57. Maheswaranathan, N. *et al.* (2018) Inferring hidden structure in multilayered neural circuits. *PLoS Comput. Biol.* 14, e1006291
58. Real, E. *et al.* (2017) Neural circuit inference from function to structure. *Curr. Biol.* 27, 189–198
59. Shi, Q. *et al.* (2019) Functional characterization of retinal ganglion cells using tailored nonlinear modeling. *Sci. Rep.* 9, 8713
60. Maheswaranathan, N. *et al.* (2019) The dynamic neural code of the retina for natural scenes. *bioRxiv* Published online December 17, 2019. https://doi.org/10.1101/340943
61. McIntosh, L.T. *et al.* (2016) Deep learning models of the retinal response to natural scenes. *Adv. Neural Inf. Proces. Syst.* 29, 1369–1377
62. Tanaka, H. *et al.* (2019) From deep learning to mechanistic understanding in neuroscience: the structure of retinal prediction. *Adv. Neural Inf. Proces. Syst.* 32, 8535–8545
63. Pang, R. *et al.* (2016) Dimensionality reduction in neuroscience. *Curr. Biol.* 26, R656–R660
64. Samengo, I. and Gollisch, T. (2013) Spike-triggered covariance: geometric proof, symmetry properties, and extension beyond Gaussian stimuli. *J. Comput. Neurosci.* 34, 137–161
65. Schwartz, O. *et al.* (2006) Spike-triggered neural characterization. *J. Vis.* 6, 484–507
66. Gollisch, T. (2013) Features and functions of nonlinear spatial integration by retinal ganglion cells. *J. Physiol. Paris* 107, 338–348
67. Kaardal, J. *et al.* (2013) Identifying functional bases for multidimensional neural computations. *Neural Comput.* 25, 1870–1890
68. Liu, J.K. *et al.* (2017) Inference of neuronal functional circuitry with spike-triggered non-negative matrix factorization. *Nat. Commun.* 8, 149
69. Lee, D.D. and Seung, H.S. (1999) Learning the parts of objects by non-negative matrix factorization. *Nature* 401, 788–791
70. Shah, N.P. *et al.* (2020) Inference of nonlinear receptive field subunits with spike-triggered clustering. *Elife* 9, e45743
71. Yu, W.-Q. *et al.* (2018) Synaptic convergence patterns onto retinal ganglion cells are preserved despite topographic variation in pre- and postsynaptic territories. *Cell Rep.* 25, 2017–2026
72. Berry, M.J. *et al.* (1999) Anticipation of moving stimuli by the retina. *Nature* 398, 334–338
73. Jarsky, T. *et al.* (2011) A synaptic mechanism for retinal adaptation to luminance and contrast. *J. Neurosci.* 31, 11003–11015
74. Manookin, M.B. and Demb, J.B. (2006) Presynaptic mechanism for slow contrast adaptation in mammalian retinal ganglion cells. *Neuron* 50, 453–464
75. Nagy, J. *et al.* (2021) GABA-A presynaptic inhibition regulates the gain and kinetics of retinal output neurons. *Elife* 10, e60994
76. Cui, Y. *et al.* (2016) Divisive suppression explains high-precision firing and contrast adaptation in retinal ganglion cells. *Elife* 5, e19460
77. Ozuysal, Y. and Baccus, S.A. (2012) Linking the computational structure of variance adaptation to biophysical mechanisms. *Neuron* 73, 1002–1015
78. Rieke, F. (2001) Temporal contrast adaptation in salamander bipolar cells. *J. Neurosci.* 21, 9445–9454
79. Baccus, S.A. and Meister, M. (2002) Fast and slow contrast adaptation in retinal circuitry. *Neuron* 36, 909–919
80. Kim, K.J. and Rieke, F. (2003) Slow Na+ inactivation and variance adaptation in salamander retinal ganglion cells. *J. Neurosci.* 23, 1506–1516
81. Weick, M. and Demb, J.B. (2011) Delayed-rectifier K channels contribute to contrast adaptation in mammalian retinal ganglion cells. *Neuron* 71, 166–179
82. Brown, S.P. and Masland, R.H. (2001) Spatial scale and cellular substrate of contrast adaptation by retinal ganglion cells. *Nat. Neurosci.* 4, 44–51
83. Garvert, M.M. and Gollisch, T. (2013) Local and global contrast adaptation in retinal ganglion cells. *Neuron* 77, 915–928
84. Ölveczky, B.P. *et al.* (2007) Retinal adaptation to object motion. *Neuron* 56, 689–700
85. Khani, M.H. and Gollisch, T. (2017) Diversity in spatial scope of contrast adaptation among mouse retinal ganglion cells. *J. Neurophysiol.* 118, 3024–3043
86. Appleby, T.R. and Manookin, M.B. (2019) Neural sensitization improves encoding fidelity in the primate retina. *Nat. Commun.* 10, 4017
87. Kastner, D.B. and Baccus, S.A. (2011) Coordinated dynamic encoding in the retina using opposing forms of plasticity. *Nat. Neurosci.* 14, 1317–1322
88. Kastner, D.B. and Baccus, S.A. (2013) Spatial segregation of adaptation and predictive sensitization in retinal ganglion cells. *Neuron* 79, 541–554
89. Wässle, H. *et al.* (1989) Cortical magnification factor and the ganglion cell density of the primate retina. *Nature* 341, 643–646
90. Sinha, R. *et al.* (2017) Cellular and circuit mechanisms shaping the perceptual properties of the primate fovea. *Cell* 168, 413–426
91. Bölinger, D. and Gollisch, T. (2012) Closed-loop measurements of iso-response stimuli reveal dynamic nonlinear stimulus integration in the retina. *Neuron* 73, 333–346
92. Petrusca, D. *et al.* (2007) Identification and characterization of a Y-like primate retinal ganglion cell type. *J. Neurosci.* 27, 11019–11027
93. Almasi, A. *et al.* (2020) Mechanisms of feature selectivity and invariance in primary visual cortex. *Cereb. Cortex* 30, 5067–5087
94. Bartsch, F. *et al.* (2021) Model-based characterization of the selectivity of neurons in primary visual cortex. *bioRxiv* Published online September 15, 2021. https://doi.org/10.1101/2021.09.13.460153
95. Vintch, B. *et al.* (2015) A convolutional subunit model for neuronal responses in macaque V1. *J. Neurosci.* 35, 14829–14841
96. Beyeler, M. *et al.* (2016) 3D visual response properties of MSTd emerge from an efficient, sparse population code. *J. Neurosci.* 36, 8399–8415
97. Mineault, P.J. *et al.* (2012) Hierarchical processing of complex motion along the primate dorsal visual pathway. *Proc. Natl. Acad. Sci. U. S. A.* 109, E972–E980
98. Ahrens, M.B. *et al.* (2008) Nonlinearities and contextual influences in auditory cortical responses modeled with multilinear spectrotemporal methods. *J. Neurosci.* 28, 1929–1942
99. McFarland, J.M. *et al.* (2013) Inferring nonlinear neuronal computation based on physiologically plausible inputs. *PLoS Comput. Biol.* 9, e1003143
100. Keshishian, M. *et al.* (2020) Estimating and interpreting nonlinear receptive field of sensory neural responses with deep neural network models. *Elife* 9, e53445
101. Cunningham, J.P. and Yu, B.M. (2014) Dimensionality reduction for large-scale neural recordings. *Nat. Neurosci.* 17, 1500–1509
102. Onken, A. *et al.* (2016) Using matrix and tensor factorizations for the single-trial analysis of population spike trains. *PLoS Comput. Biol.* 12, e1005189
103. Soulat, H. *et al.* (2021) Probabilistic tensor decomposition of neural population spiking activity. *Adv. Neural Inf. Proces. Syst.* 34, 15969–15980

104. Williams, A.H. *et al.* (2018) Unsupervised discovery of demixed, low-dimensional neural dynamics across multiple timescales through tensor component analysis. *Neuron* 98, 1099–1115
105. Barrett, D.G. *et al.* (2019) Analyzing biological and artificial neural networks: challenges with opportunities for synergy? *Curr. Opin. Neurobiol.* 55, 55–64
106. Glaser, J.I. *et al.* (2019) The roles of supervised machine learning in systems neuroscience. *Prog. Neurobiol.* 175, 126–137
107. Diamond, J.S. (2017) Inhibitory interneurons in the retina: types, circuitry, and function. *Annu. Rev. Vis. Sci.* 3, 1–24
108. Klindt, D. *et al.* (2017) Neural system identification for large populations separating 'what' and 'where.'. *Adv. Neural Inf. Proces. Syst.* 30, 3506–3516
109. Helmstaedter, M. *et al.* (2013) Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature* 500, 168–174
110. Marc, R.E. *et al.* (2014) The AII amacrine cell connectome: a dense network hub. *Front. Neural Circuits* 8, 104
111. Park, S.J.H. *et al.* (2020) Connectomic analysis reveals an interneuron with an integral role in the retinal circuit for night vision. *Elife* 9, e56077
112. Walker, E.Y. *et al.* (2019) Inception loops discover what excites neurons most using deep predictive models. *Nat. Neurosci.* 22, 2060–2065
113. Gonçalves, P.J. *et al.* (2020) Training deep neural density estimators to identify mechanistic models of neural dynamics. *Elife* 9, e56261