Structural and functional diversity of a dense sample of retinal ganglion cells

J. Alexander Bae[1,2*], Shang Mu[1*], Jinseop S. Kim[1*†], Nicholas L. Turner[1,3*],
Ignacio Tartavull[1], Nico Kemnitz[1], Chris S. Jordan[4], Alex D. Norton[4], William M. Silversmith[4],
Rachel Prentki[4], Marissa Sorek[4], Celia David[4], Devon L. Jones[4], Doug Bland[4], Amy L. R. Sterling[4],
Jungman Park[5], Kevin L. Briggman[6,7†], H. Sebastian Seung[1,3] and the EyeWirers[8]

[1]Neuroscience Institute, [2]Electrical Engineering Dept. and [3]Computer Science Dept.,
Princeton University, Princeton, NJ 08544 USA.
[4]WiredDifferently, Inc., 745 Atlantic Ave, Boston, MA 02111 USA.
[5]Device Business Unit, Marketing Group,
KT Gwanghwamun Bldg. West, 178, Sejong-daero, Jongno-gu, Seoul, Korea 03154
[6]Department of Biomedical Optics, Max Planck Institute for Medical Research, Heidelberg 69120, Germany.
[7]Circuit Dynamics and Connectivity Unit, National Institute of Neurological Disorders and Stroke,
Bethesda, MD 20892, USA.
[8]http://eyewire.org
†Present addresses: Dept. of Structure and Function of Neural Networks,
Korea Brain Research Institute, Daegu 41068, Republic of Korea (JSK);
Center of Advanced European Studies and Research (caesar),
Dept. of Computational Neuroethology, 53175 Bonn, Germany (KLB)
*Co-first authors

August 31, 2017

Abstract

To aid understanding of retinal structure and function, we present as an online resource the dendritic arbors and visual responses of ganglion cells in a single patch of mouse retina. We divide the inner plexiform layer, which contains the dendritic arbors of ganglion cells, into four sublaminae defined by a purely anatomical principle of arbor segregation. The sublaminae serve as the starting point for a hierarchical clustering of our ganglion cells. We propose and apply a quantitative criterion for validating a cluster as a ganglion cell type: the aggregate neurite density of a type should be approximately uniform ("density conservation"). Finally, we find that ganglion cells arborizing in the inner marginal sublamina of the inner plexiform layer exhibit significantly more sustained visual responses on average.

The characterization of neuronal diversity is the first of seven goals listed for the United States BRAIN Initiative (BRAIN Initiative Working Group, 2014). The goal is challenging even in the retina, a relatively simple region of the mammalian central nervous system. Ganglion cells (GCs) are the retina's only projection neurons, with axons that leave the retina and travel to the brain. A recent large-scale physiological survey based on calcium imaging identified over 35 to 50 mouse GC types, each of which carries a specific kind of visual information [Baden et al., 2016]. However, less than 1% of these GCs had their dendritic arbors reconstructed, making it difficult to relate function to structure.

Here we present a large-scale survey of almost 400 GCs in the mouse retina, combining anatomical information from serial block face scanning electron microscopy (Denk and Horstmann, 2004) with physiological information from calcium imaging. The set of reconstructed GCs is dense, meaning that it includes the arbor of every soma inside a 0.3×0.35 mm^2 patch of mouse retina. Visual responses from calcium imaging are available for 82% of the cells. Our survey is based on the e2198 dataset, a small fraction of which was previously used to study retinal circuits for motion computation [Briggman et al., 2011, Kim et al., 2014, Greene et al., 2016]. A previous dense sample of GCs from the e2006 dataset contained $10\times$ fewer cells, lacked physiological information, and the arbors of all but the smallest cells were severely cut off by the borders of the $(0.1$ mm$)^2$ patch of mouse retina [Helmstaedter et al., 2013]. In previous large-scale anatomical surveys using light microscopy, GCs were sparsely sampled from many retinas, and lacked physiological information [Badea and Nathans, 2004, Kong et al., 2005, Coombs et al., 2006, Völgyi et al., 2009, Sümbül et al., 2014].

Our dense sample of GCs reveals new principles of retinal organization. First, we show how to optimally subdivide the inner plexiform layer (IPL) of the retina using the purely anatomical principle that arbors should segregate into distinct sublamina. For GC dendritic arbors, segregation is maximized by subdividing the IPL into two marginal sublamina flanking a central sublamina. For bipolar cell (BC) axonal arbors, in contrast, segregation is maximized by subdividing the IPL into inner and outer sublamina. Using the marginal-central and inner-outer dichotomies as a starting point, we divide our dense sample of GCs into six high-level clusters, which are subdivided further to end up with 47 clusters. This raises the question of how to validate whether clusters are indeed

1

GC types.

To answer this question, we propose a "density conservation" principle: the aggregate arbor density of a GC type should be approximately uniform across the retina. Our principle is meant to supersede the traditional principle that the dendritic arbors of a GC type should "tile" the retina with little overlap [Wässle et al., 1981]. The latter seems inconsistent with observations of substantial overlap between dendritic arbors for mouse GC types [Zhang et al., 2012, Rousso et al., 2016]. Density conservation is shown to be satisfied by 24 of our clusters, which are internally validated as pure types. Six of these types appear novel, in the sense that we have been unable to find any matching reports in the literature.

Using the calcium imaging data, we have the opportunity to relate the above structural analyses to retinal function. It was previously proposed that marginal GCs have more sustained responses to visual stimuli than central GCs [Roska and Werblin, 2001]. We reexamine this conventional wisdom, and confirm that inner marginal GCs are significantly more sustained on average, but other high-level GC clusters do not exhibit statistically significant differences in sustainedness. The analogous finding has recently been reported for mouse BCs [Franke et al., 2017].

Finally, we emphasize that our dense sample of GCs is unprecedented in its combined coverage of anatomy and physiology. Dense sampling was a key enabler for our analyses of arbor segregation, density conservation, and sustainedness, and should enable further discoveries by other researchers. We are making the data publicly available at the EyeWire Museum (http://museum.eyewire.org), where every reconstructed GC can be interactively viewed along with its visual response properties. Due to its cellular and subcellular resolution, the Museum is novel relative to traditional atlases of brain regions.

Dense anatomy and physiology

As previously described in Briggman et al. [2011], the e2198 dataset had a physiological component (time series of visual responses observed via two-photon calcium imaging *ex vivo*) and an anatomical component (3D image stack from serial block face scanning electron microscopy). Both imaging techniques were targeted at the same patch of a single mouse retina (Fig. 1a). The orientation of the patch relative to body axes was inferred from the reconstructed GCs (Methods), as information about the orientation of the retina was not recorded at the time of dissection.

In the anatomical dataset (Fig. 1b), every cell body in the ganglion cell layer (GCL) was examined for evidence of an axon. Displaced GCs, with cell bodies in the inner nuclear layer (INL), were neglected by our survey. The axon emanated directly from the cell body in some cases, and branched from a primary dendrite in others. Through this systematic search, we identified a total of 396 GCs in e2198. Many of the GCs are rendered in Supplementary Fig. 1 and the Supplementary Video. The dendritic arbors of all cells were reconstructed by almost 30,000 members of the online community known as EyeWire (Methods, Supplementary Notes). For subsequent analyses, we also made use of reconstructions of starburst amacrine cells (SACs) and bipolar cells (BCs) from previous studies [Kim et al., 2014, Greene et al., 2016].

Eleven cells were so severely cut off by the borders of the imaged volume that they were discarded. Four more cells were discarded as "weirdos" that may have resulted from developmental abnormalities. Excluding severely cut-off and "weirdo" cells left a sample of 381 that were used for subsequent analysis. All cells, including cut-off and "weirdo" cells, can be examined interactively at the EyeWire Museum.

Previous light microscopic surveys sparsely sampled GCs from many locations in many retinas [Badea and Nathans, 2004, Kong et al., 2005, Coombs et al., 2006, Völgyi et al., 2009, Sümbül et al., 2014]. The arbor diameter and density of alpha GCs are known to depend strongly on retinal location [Bleckert et al., 2014]. Asymmetry [Kim et al., 2008] and stratification depth [Sümbül et al., 2014] of the J cell also depend on retinal location. Because the cells of our dense sample come from a single location in a single retina, variation within a GC type is expected to be relatively small, facilitating accurate classification of GC types.

Some sparse sampling techniques used in the light microscopic surveys might miss GC types due to selection bias. Our dense GC sample is expected to include examples of almost all GC types. Possible exceptions would be types that are very rare and happen by chance not to occur in our finite sample, or types that are systematically absent in the region of the retina that contains the sample. Regarding the latter possibility, we are not aware of GC types that are known to exist in one retinal region but not another.

Our GC sample also differs from previous light microscopic surveys because it comes with visual responses from calcium imaging. In the physiological dataset (Fig. 1c), regions of interest (ROIs) were manually drawn around cell bodies in the GCL. For each ROI, the time series of the calcium signal was computed as the sum of the pixels in the ROI for each time point (Fig. 1d). We were able to extract calcium signals for 326 cells, 82% of the 396 GCs that were reconstructed from the anatomical dataset. The remaining cells lacked meaningful signals, either because they lacked sufficient calcium indicator, lacked sufficiently strong visual responses, or were too far from the focal plane used for two-photon imaging. Briggman et al. [2011] only reported calcium signals for 25 On-Off direction selective GCs, a small fraction of the current dataset.

The visual stimulus was a light bar on a dark background, with width 0.2 mm and length 1 mm (Fig. 1d). In each stimulus trial, the bar moved at 1 mm/sec along its long axis. Eight directions of movement were used to evaluate direction selectivity (DS). Two-photon imaging was sequentially applied to each tile in a 3×3 array that covered the retinal patch with slight overlap. The vertical lines in Fig. 1d indicate when the leading and trailing edges of the bar crossed the center of the imaged tile. These nominal stimulus onset and offset times differ slightly from the true times, because the receptive fields of cells in the imaged tile vary in their location and size.

Arbor segregation principle

In his description of ganglion cell diversity, Cajal [1893] divided the IPL into five sublaminae (S1-S5). It is unclear whether the five-way division is a subjective convention or has objective meaning. Wässle [2004] proposed a molecular basis (calbindin or calretinin staining) for the three borders between S1, S2, S3, and S4/5. Again, it is unclear whether the stained bands are more than arbitrary conventions. Famiglietti and Kolb [1976] divided the IPL into sublamina a, specialized for processing of Off (dark) stimuli, and sublamina b, specialized for processing of On (light) stimuli. This two-way division has objective meaning, but is rather coarse.

Here we show that the IPL can be divided into sublaminae based on the purely anatomical principle that BC arbors should segregate in the sublaminae. We normalize IPL depth so that it ranges from 0 to 1, where 0 denotes the border with the INL and 1 the border with the GCL (Methods). Suppose that we divide the IPL into inner and outer sublaminae. The inner-outer boundary is located at 0.47 IPL depth for now; the optimality of this value will be demonstrated later. A histogram of the difference between the amounts of inner arbor and outer arbor reveals that BCs separate into two clusters (Fig. 2a). Cells of one cluster have a mostly inner arbor, while cells of the other cluster have a mostly outer arbor (Fig. 2b). The gap in the center of the histogram indicates rarity of cells that evenly straddle the inner-outer boundary (Fig. 2a).

If we now vary the location of the inner-outer boundary, we obtain a family of histograms (not shown). Each histogram can be summarized by a numerical index of segregation, defined as the separation between cluster centers divided by their average width (Methods, Fig. 2d). The segregation index is graphed versus the location of the inner-outer boundary in Fig. 2c. It turns out that arbor segregation is maximized at 0.47 IPL depth; this was the value used for the histogram in Fig. 2a. The global maximum at 0.47 IPL depth is the On-Off boundary, in the sense that On and Off BC types lie on either side, as will be explained more fully later on. Therefore our arbor segregation principle reproduces the Famiglietti and Kolb [1976] two-way division of the IPL into sublamina *a* and *b*. The novelty is that this division emerges from a purely anatomical principle, without use of physiological information.

Flanking the global maximum at 0.47 IPL depth, there are also local maxima at 0.28 and 0.65 IPL depth (Fig. 2c). The local maximum at 0.28 coincides with the Off SAC depth. The local maximum at 0.65 is very close to the On SAC depth of 0.62 (see Supplementary Methods about the small deviation). It should be noted that SAC-specific staining has become the most popular landmark for IPL depth [Manookin et al., 2008, Siegert et al., 2009]. The arbor segregation principle supports the idea that SACs are more than merely convenient landmarks; they are objective divisions in the IPL.

To summarize, BC arbor segregation supports the division of the IPL into four sublaminae, which we will call outer marginal, outer central, inner central, and inner marginal. The inner-outer boundary corresponds to the On-Off boundary and the marginal-central boundaries are at the SACs. Further subdivision of the inner marginal sublamina to create a total of five sublaminae is also defensible but less convincing (Supplementary Methods).

We now examine whether GC arbors segregate across the four sublaminae. It turns out that GC arbors do not segregate well across the inner-outer boundary (Extended Data Fig. 1b). This is not surprising, given that so many On-Off GC types are known. However, good segregation is obtained if we consider the difference between the amount of arbor in the two marginal sublaminae minus the amount of arbor in the two central sublaminae (Fig. 2d, e). The optimal locations for the marginal-central boundaries are exactly at the On and Off SAC depths (Fig. 2f). For simplicity, the preceding analysis omitted the On-Off and On DS cells [Barlow and Levick, 1965, Sabbah et al., 2017]. The full analysis including the DS cells is slightly more complex while preserving the main findings (Extended Data Fig. 1d-f and Supplementary Methods).

Hierarchical clustering

By dividing the IPL into four sublaminae, the preceding section revealed a difference in BC and GC organization. Marginal-central segregation is stronger for GCs, while inner-outer segregation is stronger for BCs. These splits can be used as starting points for divisive hierarchical clustering of cells. For BCs, the inner-outer split of Fig. 2a is followed by marginal-central splits yielding four high-level clusters (Fig. 2h, Supplementary Fig. 2). Each of these four clusters can be further subdivided into BC types using anatomical criteria described previously [Helmstaedter et al., 2013, Kim et al., 2014, Greene et al., 2016]. We define the "stratification profile" as the linear density of arbor as a function of IPL depth (Methods). The profile is treated like a probability distribution, with its area normalized to unity. The three boundaries between the four sublaminae are visible as "notches" in a graph of the average stratification profiles of the BC types (Fig. 2g). As mentioned earlier, the inner-outer boundary divides the On BC types (BC5-9) from the Off BC types (BC1-4). According to recent data, the On-Off distinction is essentially binary for mouse BCs [Franke et al., 2017].

For GCs, a high-level clustering also follows from the four sublaminae (Fig. 2j, Supplementary Fig. 3). We first separate the DS cells based on their strong co-stratification with SACs (Extended Data Fig. 2b, Supplementary Data 1, split a-1). The remaining GCs are split into marginal and central clusters as in Fig. 2d (see also Supplementary Data 1, split a-2). The marginal cluster separates into inner and outer clusters (Extended Data Fig. 2c, Supplementary Data 1, split a-4). The central cluster separates into inner, outer, and inner-outer clusters (Extended Data Fig. 2d, Supplementary Data 1, split a-3). This procedure yields a total of six high-level clusters (Fig. 2j): DS, inner marginal, outer marginal, inner central, outer central, and inner-outer central. The average stratification profiles of the high-level clusters are shown in Fig. 2i. The inner-outer central cells straddle the inner-outer boundary, and are the main cause of poor inner-outer segregation of GCs, which was noted above.

Each high-level cluster was subdivided using a decision tree (Methods), which yielded a total of 47 low-level clusters. The dendrogram of Extended Data Fig. 2a shows how the low-level clusters emerge from splits of the high-level clusters (see also Supplementary Figs. 3 and 4). All decisions in the dendrogram are documented in Supplementary Data 1. We have compiled a one-page gallery (Extended Data Fig. 3) that illustrates each cluster with a single example cell, and a multi-page gallery (Supplementary Data 2) that shows all cells sorted by cluster. All cells and clusters can be viewed interactively at the EyeWire Museum.

Each decision in the tree was made by thresholding an anatomical quantity, which was typically some percentile of the stratification profile (Extended Data Fig. 4c, d). In some cases we restricted the stratification profile to a particular range of IPL depths and renormalized its area to unity (Extended Data Fig. 4e, f). This was usually to exclude the dendritic trunks, which appeared to contribute noise to the classification. Finally, some decisions relied on soma size (Extended Data Fig. 5), SAC contact analysis (Extended Data Fig. 6), and arbor density and complexity (Extended Data Fig. 7).

Figure 3a summarizes the anatomy of each cluster by its stratification profile averaged over cells within the cluster. The stratification profiles are grouped by their membership in the six high-level clusters of Fig. 2i. The physiology of each cluster is summarized

3

by its "temporal response function," defined as visual response versus time, averaged over trials, directions of the moving bar (Fig. 1e), and cells within the cluster. All temporal response functions are normalized in the graphs so that their minimum and maximum values are the same. In reality, response amplitudes varied greatly across clusters, and some clusters responded only very weakly. For example, the 1ws temporal response function is very noisy, because its light-evoked response was so weak, and was averaged over only two cells. More detailed information about visual responses can be found in Supplementary Data 2 and 3. It should be emphasized that the visual responses were not used at all to define the anatomical clusters. The average visual responses in Fig. 3a were computed after the clustering was complete.

Our nomenclature for the clusters uses numbers and letters to indicate anatomical properties (Fig. 3b). The name of each cluster begins with one or more integers in the range 1 to 9 that roughly describe the stratification profile. The first number encodes the location of the global maximum of the average stratification profile, when computed over ten bins corresponding to 10 sublayers of the IPL with equal thickness. Added numbers encode the location(s) of a local maxima if they exist. (No maxima were found in the tenth bin.) Letters are added as suffixes to indicate features that distinguish clusters with similar stratification profiles. The division into 10 sublayers is merely a convention, unlike our earlier objective division into four sublaminae.

For any given cluster, one might want a measure of confidence that the cluster is actually a pure GC type. There are several ways to obtain this. First, our clustering procedure is "transparent" in the sense that every decision, while expressed as a computational algorithm, can also be understood and examined by a human. By examining Supplementary Data 1, one can trace the hierarchical sequence of decisions that lead to any given cluster, which yields some qualitative impression of confidence in the cluster.

Second, many of our clusters can be externally validated because they correspond well with "securely known" types, defined as those that have been extensively characterized by a combination of molecular, physiological, and anatomical techniques [Sanes and Masland, 2015, Rousso et al., 2016]. Correspondences are provided in Fig. 3a (see Supplementary Methods for detailed justifications). Our 1wt, 4ow, and 8w correspond to the classical alpha types [Pang et al., 2012], and our 6sw corresponds to the nonclassical transient On alpha type [Krieger et al., 2017]. Our 37c, 37d, 37r, and 37v correspond to the On-Off DS types, and 7ir, 7id, and 7iv correspond to the classical On DS types. Our 1ws and 9w correspond to the M1 and M2 melanopsin types. Our 51 corresponds to the W3B type [Zhang et al., 2012]. Our 2an, 63, and 6t correspond to F-mini[Off], F-mini[On], and F-midi[On] of Rousso et al. [2016].

Density conservation principle

For a third measure of confidence, we would like some quantitative validation procedure that is "internal," meaning that it depends only on information within the dataset. This could conceivably come from the "mosaic principle," according to which the cell bodies of a GC type are arranged as if they repel each other [Wassle and Riemann, 1978]. Mosaic analysis utilizes the locations of cell bodies, which today are typically obtained using molecular labeling of a GC type [Kim et al., 2008, Huberman et al., 2008, Zhang et al., 2012, Rousso et al., 2016].

Because our reconstructions have both cell bodies and dendritic arbors, it would be more powerful to use the "tiling principle," according to which the dendritic arbors of a GC type should "tile" the retina with little overlap, almost like the tiles of a floor. Tiling is traditionally quantified by the coverage factor, which is defined as the average number of arbors that cover a retinal location. Perfect tiling would yield a coverage factor of 1. However, there have been many reports of GC types with coverage factors markedly greater than one [Wassle et al., 1981, Stein et al., 1996]. More recently, genetic techniques have been used to verify that GC types can exhibit high coverage factors while still satisfying the mosaic principle [Zhang et al., 2012, Rousso et al., 2016]. For our GC clusters, the median coverage factor is between 2 and 3 (Fig. 3d), indicating substantial overlap between neighboring arbors.

To illustrate violation of the tiling principle, Fig. 4a shows how the arbors of an example cluster cover the retina. Each arbor is represented by its convex hull, and from the overlap between hulls we can see that the coverage of a retinal location can be as high as 5 (Fig. 4b). However, it is not the case that a region covered by 5 hulls contains $5\times$ more arbor than a region covered by 1 hull. On the contrary, the aggregate arbor density is almost independent of coverage (Fig. 4c). This example suggests that coverage as quantified by convex hull overlap can be misleading.

Motivated by this example, we propose that the arbors of a type add up to roughly uniform density across the retina. We call this the "density conservation principle," and it reduces to the traditional tiling principle for the special case of arbors with uniform density within their convex hulls. For arbors that vary in density across their convex hulls, our new principle is compatible with arbor overlap. We have found a prior qualitative report of density conservation in the literature [Dacey, 1989], and related arguments have been made about overlap between GC receptive fields [Borghuis et al., 2008]. Here we present the first quantitative analysis of density conservation, and investigate its universality by applying it to all our GC clusters.

We first defined a central "crop region" in e2198 (Fig. 4d). Cropping excluded the parts of e2198 near the borders, which are expected to have lower aggregate arbor density because we did not reconstruct neurites of cells with their somas outside e2198. The crop region was divided into a grid of boxes (Fig. 4d). In each grid box, we computed the aggregate arbor density. Then we computed the coefficient of variation (standard deviation divided by mean) of the aggregate arbor density across the grid boxes (Fig. 4e). We expected the coefficient of variation to be small, and indeed it was for many cells (Fig. 4f).

To assess statistical significance of a small coefficient of variation, we created an ensemble of randomized configurations from the original cluster. The soma positions and arbor orientations were randomized in a way that left the aggregate arbor density in the crop region roughly constant (Fig. 4g and Methods). The coefficient of variation of the aggregate arbor density was calculated for each randomized configuration (Fig. 4h). We deemed a small coefficient of variation to be statistically significant if it had less than 1% probability of emerging from the randomized ensemble.

We found that 24 of the 47 clusters exhibit statistically significant density conservation (Fig. 4f). We also examined securely known types that failed the density conservation test, and found that these failures were clusters containing relatively few cells (Fig. 3c). For example, 37d is the least numerous of the four On-Off DS types in our sample. It contains few cells (6) and just barely fails

the test. 4ow is the second least numerous of the four alpha types in our sample. It also contains few cells (4) and just barely fails the test. Both types might have passed the test had our sample been larger. So there is no strong violation of the conjecture that density conservation is universal for all GC types.

Even if a necessary criterion, density conservation cannot by itself be a sufficient criterion for a pure GC type. To see why, consider a thought experiment in which two pure types exactly satisfy density conservation. A mixture of the two types will also exactly satisfy density conservation. However, we should be able to detect a mixture by its abnormally large coverage factor, which will be the sum of the coverage factors of the pure types.

Therefore, we propose that a cluster is internally validated as a type if (1) it satisfies density conservation and (2) its coverage factor is in the normal range. This range is roughly 1.5 to 3.5 for our clusters (Fig. 3d), which is consistent with many reports of GC coverage factors in the literature. (The coverage factors lower than 1.5 in Fig. 3d are likely underestimates, because they are mainly for clusters containing few large cells, for which the coverage factor computation is corrupted by edge effects.)

If we take the union of our internally validated types with the securely known types, we end up with a total of 33 clusters that are validated as types, plus one cluster (2aw) that appears to contain two securely known types (F-midi[Off] and J). These contain 84% of all cells in the sample of 381 (Fig. 3c). Further discussion of the 2aw and 63 clusters can be found in the Supplementary Methods.

Novel ganglion cell types

While the present work focuses on general principles of IPL organization, as a bonus it also yields six particular GC types that appear to be novel. For six of our internally validated clusters (1ni, 1no, 2o, 85, 27, 5to in Fig. 5), we have been unable to find unambiguous correspondences with previously published types. Outer marginal types 1ni and 1no look remarkably similar in Figs. 5a and 5b, and are novel types that co-stratify with 1ws (M1 melanopsin). Our clustering procedure separates 1ni and 1no based on a small but systematic difference in their stratification profiles (Fig. 5c and Supplementary Data 1, split e-7). Their average temporal response functions also differ slightly (Fig. 5c).

Outer central type 5to appears monostratified in the tangential view of Fig. 5d. However, its average stratification profile is relatively broad and contains more than one peak (Fig. 3a). Multiple peaks can also be seen in the stratification profiles of individual 5to cells (Supplementary Data 2 and Museum). Inner marginal type 85 arborizes throughout the inner IPL, but also extends sparse branches towards the INL. Its average stratification profile (Fig. 3a) exhibits three peaks in the inner marginal, inner central, and outer marginal sublamina. Both 5to and 85 show that the stratification profile can be very rich in information.

Outer marginal types 2o and 27 are depicted in Figs. 5f and 5g. The 2o soma is almost as large as that of the classical alpha types, which have the largest somas of all GCs (Fig. 5h). The 27 soma is of more typical size (Fig. 5h). Of the outer marginal cells, 2o exhibits the most transient Off response (Fig. 5h).

These and all other clusters can be interactively viewed in the EyeWire Museum, which provides search bar access to reconstructions (Fig. 5i), stratification profiles, and visual responses (Fig. 5j).

Nine of our internally validated types (5ti, 5so, 5si, 4i, 4on, 6sn,

7o, 73, 82wi) correspond with GC classes that were recently defined by physiological and anatomical techniques [Baden et al., 2016, Jacoby and Schwartz, 2017, Mani and Schwartz, 2017, Sabbah et al., 2017] but have not yet been confirmed as pure types by molecular techniques and mosaic analysis. Our density conservation test (Fig. 4f) provides supporting evidence that these classes are indeed pure types. Our 5ti, 5so, and 5si co-stratify with 51 (W3b), and correspond to the HD family of Jacoby and Schwartz [2017]. Our 4i and 4on co-stratify with 4ow (transient Off alpha), and our 6sn co-stratifies with 6sw (transient On alpha). These may correspond with "mini alpha" types identified by Baden et al. [2016]. Our 7o corresponds to the nonclassical transient On DS cell found by Baden et al. [2016] and Sabbah et al. [2017]. Our 73 corresponds to the On delayed cell as defined by Mani and Schwartz [2017]. Our 82wi corresponds to the vertically orientation selective cell studied by Nath and Schwartz [2016]. Some of the confirmed types can also be found in the survey of Helmstaedter et al. [2013]. One of our internally validated types (25) confirms a type in Helmstaedter et al. [2013] that has not yet been identified by physiologists. 25 is the fourth most numerous cluster in our sample (Fig. 3c). More detailed evidence for the above correspondences is given in the Supplementary Discussion.

Sustained vs. transient

Our distinction between marginal and central GC clusters occurs at the top of the GC hierarchy (Fig. 2j and Extended Data Fig. 2a), so it seems fundamental. Could this anatomical distinction have functional significance? Physiologists classify retinal neurons as sustained or transient, mainly based on duration of response to a sudden change in illumination [Cleland et al., 1971]. A previous study combined electrophysiology and light microscopic anatomy to provide evidence that central GCs are transient, while marginal GCs are sustained [Roska and Werblin, 2001]. Our dense sample of GCs provides an opportunity to systematically characterize differences in sustainedness.

For each high-level cluster we averaged the temporal response functions of all cells in the cluster. The average response of the inner marginal cluster is markedly more sustained than that of the other high-level clusters (Fig. 6a). We then quantified the sustainedness of each cell by the value of its temporal response function at a fixed time after nominal stimulus onset, relative to the peak value (Fig. 6b and Methods). Inner marginal cells are significantly more sustained on average than the cells in the other high-level clusters (Fig. 6c). Differences between other clusters are not statistically significant.

One might worry that the temporal resolution of calcium imaging is inadequate for revealing sustained-transient differences. The four alpha types are a counterexample to this idea, because they exhibit sustained-transient differences (Fig. 6d) consistent with those that have been observed using electrophysiology [Pang et al., 2012, Krieger et al., 2017].

We refine the analysis by plotting the average sustainedness for the cells of each low-level cluster versus its peak stratification (Fig. 6e). For comparison, we also plot the average temporal response function for each low-level cluster (Extended Data Fig. 8). The numbers for any individual cluster should be taken with caution, especially as some clusters have not been validated as types. The sustainedness of the outer marginal cells spans a wide range; 1wt is

5

most sustained, while 2o is most transient. The sustainedness of the inner marginal cells also spans a wide range. Given this diversity, the idea that marginal cells are more sustained than central cells may be simplistic.

One concern about the analysis is that our stimulus is a moving bar, while a flashed stimulus is customarily used for the sustained-transient distinction. For a moving stimulus, the trailing edge leaves the receptive field later for an arbor with larger diameter. This might cause large cells to spuriously seem more sustained than they actually are. To control for this possibility, we recalculated the sustainedness indices after translating the response of each cell backwards in time by an amount equal to its cluster-specific arbor radius divided by the stimulus speed (Methods). The result is shown in Extended Data Fig. 9, which looks very similar to Fig. 6e.

Discussion

In light of our findings for GCs, it is helpful to reexamine the analogous claim that marginal BCs are sustained while central BCs are transient [Borghuis et al., 2013, Baden et al., 2013]. For a full field flashed stimulus, inner marginal BCs are the most sustained cluster; the differences between the other three clusters are more minor (Extended Data Fig. 6m of Franke et al., 2017). This finding is strikingly analogous to our own finding for GCs, and would be predicted from the hypothesis that GCs inherit their sustainedness from their BC inputs [Awatramani and Slaughter, 2000].

The conventional wisdom that marginal BCs are more sustained than central BCs still appears valid, if the comparison is restricted to inner cells only or outer cells only [Franke et al., 2017]. But it would be misleading to say that marginal is unconditionally more sustained than central, as inner central BCs can be more sustained than outer marginal BCs [Franke et al., 2017], and the same may be true for our GCs (Fig. 6e). One caveat is that sustained versus transient could depend on the specifics of the stimulus. For example, Franke et al. [2017] find that inner central cells are more transient with a full field than a local stimulus. A second caveat is that there may be heterogeneity across the types within a high-level cluster. For example, BC9 is markedly more transient than other inner marginal BC types for a full field stimulus [Franke et al., 2017]. Heterogeneity is likely even greater for GCs, which come in many more types than BCs. A third caveat is that sustainedness may depend on whether conditions are photopic or scotopic [Grimes et al., 2014]. All of these caveats support the overall conclusion that sustained versus transient is more complex and subtle than a simple dichotomy.

Our purely anatomical subdivision of the IPL into marginal and central sublamina supports the idea that using SACs as landmarks is fundamental, and not merely a convenience made popular by the ease of ChAT staining. It has been proposed that SAC arbors could serve as a scaffold for development of some types of GC arbors [Stacy and Wong, 2003], so the marginal-central division is potentially relevant for neural development.

We have proposed a density conservation principle to replace the tiling principle for GC types. Density conservation makes functional sense as it would serve to make uniform the density of synapses from each BC type to each GC type. The tiling principle can be fulfilled if dendrites of different cells repel each other during development [Grueber and Sagasti, 2010]. We speculate that density conservation could be established by a developmental process in which homotypic dendrites repel each other without regard to whether they belong to the same cell or different cells.

Of the 47 clusters that we identified, 24 were internally validated by the density conservation test (Fig. 4f). Most of the remaining clusters contain too few cells to satisfy the criteria with statistical significance, but some of them can be externally validated because their properties match "securely known" types that have been published previously. If we take the union of internally validated types and securely known types, we end up with a total of 35 types. This lower bound on the number of GC types is consistent with the recent identification of 35 to 50 clusters in the large-scale physiological survey of mouse retinal GCs mentioned previously [Baden et al., 2016].

References

David Arthur and Sergei Vassilvitskii. k-means++: The advantages of careful seeding. In *Proceedings of the 18th annual ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035. Society for Industrial and Applied Mathematics, 2007.

G B Awatramani and M M Slaughter. Origin of transient and sustained responses in ganglion cells of the retina. *J Neurosci*, 20 (18):7087–95, Sep 2000.

Tudor Constantin Badea and Jeremy Nathans. Quantitative analysis of neuronal morphologies in the mouse retina visualized by using a genetically directed reporter. *J Comp Neurol*, 480(4): 331–51, Dec 2004. doi: 10.1002/cne.20304.

Tom Baden, Philipp Berens, Matthias Bethge, and Thomas Euler. Spikes in mammalian bipolar cells support temporal layering of the inner retina. *Curr Biol*, 23(1):48–52, 2013.

Tom Baden, Philipp Berens, Katrin Franke, Miroslav Román Rosón, Matthias Bethge, and Thomas Euler. The functional diversity of retinal ganglion cells in the mouse. *Nature*, 529(7586):345–50, Jan 2016. doi: 10.1038/nature16468.

H B Barlow and W R Levick. The mechanism of directionally selective units in rabbit's retina. *J Physiol*, 178(3):477–504, Jun 1965.

Adam Bleckert, Gregory W Schwartz, Maxwell H Turner, Fred Rieke, and Rachel O L Wong. Visual space is represented by nonmatching topographies of distinct mouse retinal ganglion cell types. *Curr Biol*, 24(3):310–5, Feb 2014. doi: 10.1016/j.cub. 2013.12.020.

Bart G Borghuis, Charles P Ratliff, Robert G Smith, Peter Sterling, and Vijay Balasubramanian. Design of a neuronal array. *Journal of Neuroscience*, 28(12):3178–3189, 2008.

Bart G Borghuis, Jonathan S Marvin, Loren L Looger, and Jonathan B Demb. Two-photon imaging of nonlinear glutamate release dynamics at bipolar cell synapses in the mouse retina. *J Neurosci*, 33(27):10972–85, 2013.

KTS Brazier. Confidence intervals from the rayleigh test. *Monthly Notices of the Royal Astronomical Society*, 268(3):709–712, 1994.

Kevin L Briggman, Moritz Helmstaedter, and Winfried Denk. Wiring specificity in the direction-selectivity circuit of the retina. *Nature*, 471(7337):183–8, Mar 2011. doi: 10.1038/nature09818.

Santiago Ramón y Cajal. La rétine des vertébrés. *Cellule*, 9:17–257, 1893.

BG Cleland, MW Dubin, and WR Levick. Sustained and transient neurones in the cat's retina and lateral geniculate nucleus. *The Journal of Physiology*, 217(2):473–496, 1971.

J Coombs, D Van Der List, G-Y Wang, and LM Chalupa. Morphological properties of mouse retinal ganglion cells. *Neuroscience*, 140(1):123–136, 2006.

Dennis M Dacey. Monoamine-accumulating ganglion cell type of the cat's retina. *Journal of Comparative Neurology*, 288(1):59–80, 1989.

Winfried Denk and Heinz Horstmann. Serial block-face scanning electron microscopy to reconstruct three-dimensional tissue nanostructure. *PLoS Biol*, 2(11):e329, 2004.

E V Famiglietti, Jr and H Kolb. Structural basis for on-and off-center responses in retinal ganglion cells. *Science*, 194(4261):193–5, Oct 1976.

Katrin Franke, Philipp Berens, Timm Schubert, Matthias Bethge, Thomas Euler, and Tom Baden. Inhibition decorrelates visual feature representations in the inner retina. *Nature*, 542(7642):439–444, 2017.

Matthew J Greene, Jinseop S Kim, H Sebastian Seung, and Eye-Wirers. Analogous convergence of sustained and transient inputs in parallel on and off pathways for retinal motion computation. *Cell Rep*, 14(8):1892–900, Mar 2016. doi: 10.1016/j.celrep.2016.02.001.

William N Grimes, Gregory W Schwartz, and Fred Rieke. The synaptic and circuit mechanisms underlying a change in spatial encoding in the retina. *Neuron*, 82(2):460–473, 2014.

Wesley B Grueber and Alvaro Sagasti. Self-avoidance and tiling: mechanisms of dendrite and axon spacing. *Cold Spring Harbor perspectives in biology*, 2(9):a001750, 2010.

Moritz Helmstaedter, Kevin L Briggman, Srinivas C Turaga, Viren Jain, H Sebastian Seung, and Winfried Denk. Connectomic reconstruction of the inner plexiform layer in the mouse retina. *Nature*, 500(7461):168–74, Aug 2013. doi: 10.1038/nature12346.

A.D. Huberman, M. Manu, S.M. Koch, M.W. Susman, A.B. Lutz, E.M. Ullian, S.A. Baccus, and B.A. Barres. Architecture and activity-mediated refinement of axonal projections from a mosaic of genetically identified retinal ganglion cells. *Neuron*, 59(3):425–438, 2008.

Jason Jacoby and Gregory W Schwartz. Three small-receptive-field ganglion cells in the mouse retina are distinctly tuned to size, speed, and object motion. *Journal of Neuroscience*, 37(3):610–625, 2017.

In-Jung Kim, Yifeng Zhang, Masahito Yamagata, Markus Meister, and Joshua R Sanes. Molecular identification of a retinal cell type that responds to upward motion. *Nature*, 452(7186):478–82, Mar 2008. doi: 10.1038/nature06739.

Jinseop S Kim, Matthew J Greene, Aleksandar Zlateski, Kisuk Lee, Mark Richardson, Srinivas C Turaga, Michael Purcaro, Matthew Balkam, Amy Robinson, Bardia F Behabadi, Michael Campos, Winfried Denk, H Sebastian Seung, and the Eye-Wirers. Space-time wiring specificity supports direction selectivity in the retina. *Nature*, 509(7500):331–6, May 2014. doi: 10.1038/nature13240.

Jee-Hyun Kong, Daniel R Fish, Rebecca L Rockhill, and Richard H Masland. Diversity of ganglion cells in the mouse retina: unsupervised morphological classification and its limits. *Journal of Comparative Neurology*, 489(3):293–310, 2005.

Brenna Krieger, Mu Qiao, David L Rousso, Joshua R Sanes, and Markus Meister. Four alpha ganglion cell types in mouse retina: Function, structure, and molecular signatures. *PloS one*, 12(7):e0180091, 2017.

Adam Mani and Gregory W Schwartz. Circuit mechanisms of a retinal ganglion cell with stimulus-dependent response latency and activation beyond its dendrites. *Current Biology*, 27(4):471–482, 2017.

Michael B Manookin, Deborah Langrill Beaudoin, Zachary Raymond Ernst, Leigh J Flagel, and Jonathan B Demb. Disinhibition combines with excitation to extend the operating range of the off visual pathway in daylight. *J Neurosci*, 28(16):4136–50, Apr 2008. doi: 10.1523/JNEUROSCI.4274-07.2008.

Calvin R Maurer, Rensheng Qi, and Vijay Raghavan. A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(2):265–270, 2003.

Amurta Nath and Gregory W Schwartz. Cardinal orientation selectivity is represented by two distinct ganglion cell types in mouse retina. *J Neurosci*, 36(11):3208–21, Mar 2016. doi: 10.1523/JNEUROSCI.4554-15.2016.

C W Oyster and H B Barlow. Direction-selective units in rabbit retina: distribution of preferred directions. *Science*, 155(3764):841–2, Feb 1967.

Ji-Jie Pang, Fan Gao, David L Paul, and Samuel M Wu. Rod, m-cone and m/s-cone inputs to hyperpolarizing bipolar cells in the mouse retina. *J Physiol*, 590(Pt 4):845–54, Feb 2012. doi: 10.1113/jphysiol.2011.224113.

B. Roska and F. Werblin. Vertical interactions across ten parallel, stacked representations in the mammalian retina. *Nature*, 410(6828):583–587, 2001.

David L Rousso, Mu Qiao, Ruth D Kagan, Masahito Yamagata, Richard D Palmiter, and Joshua R Sanes. Two pairs of on and off retinal ganglion cells are defined by intersectional patterns of transcription factor expression. *Cell reports*, 15(9):1930–1944, 2016.

Shai Sabbah, John A Gemmer, Ananya Bhatia-Lin, Gabrielle Manoff, Gabriel Castro, Jesse K Siegel, Nathan Jeffery, and David M Berson. A retinal code for motion along the gravitational and body axes. *Nature*, 2017.

730 Joshua R Sanes and Richard H Masland. The types of retinal ganglion cells: Current status and implications for neuronal classification. *Annu Rev Neurosci*, Apr 2015. doi: 10.1146/annurev-neuro-071714-034120.

735 Mie Sato, Ingmar Bitter, Michael A Bender, Arie E Kaufman, and Masayuki Nakajima. Teasar: Tree-structure extraction algorithm for accurate and robust skeletons. In *Computer Graphics and Applications, 2000. Proceedings. The Eighth Pacific Conference on*, pages 281–449. IEEE, 2000.

740 Sandra Siegert, Brigitte Gross Scherf, Karina Del Punta, Nick Didkovsky, Nathaniel Heintz, and Botond Roska. Genetic address book for retinal cell types. *Nat Neurosci*, 12(9):1197–204, Sep 2009. doi: 10.1038/nn.2370.

Rebecca Colleen Stacy and Rachel Wong. Developmental relationship between cholinergic amacrine cell processes and ganglion cell dendrites of the mouse retina. *Journal of Comparative Neurology*, 456(2):154–166, 2003.
745

JJ Stein, SA Johnson, and DM Berson. Distribution and coverage of beta cells in the cat retina. *Journal of Comparative Neurology*, 372(4):597–617, 1996.

750 Uygar Sümbül, Sen Song, Kyle McCulloch, Michael Becker, Bin Lin, Joshua R Sanes, Richard H Masland, and H Sebastian Seung. A genetic and computational approach to structurally classify neuronal types. *Nat Commun*, 5:3512, 2014. doi: 10.1038/ncomms4512.

755 Béla Völgyi, Samir Chheda, and Stewart A Bloomfield. Tracer coupling patterns of the ganglion cell subtypes in the mouse retina. *Journal of Comparative Neurology*, 512(5):664–687, 2009.

H Wassle and HJ Riemann. The mosaic of nerve cells in the mammalian retina. *Proceedings of the Royal Society of London B: Biological Sciences*, 200(1141):441–461, 1978.
760

H Wassle, BB Boycott, and R-B Illing. Morphology and mosaic of on-and off-beta cells in the cat retina and some functional considerations. *Proceedings of the Royal Society of London B: Biological Sciences*, 212(1187):177–195, 1981.
765

H Wässle, L Peichl, and B B Boycott. Dendritic territories of cat retinal ganglion cells. *Nature*, 292(5821):344–5, Jul 1981.

Heinz Wässle. Parallel processing in the mammalian retina. *Nature reviews. Neuroscience*, 5(10):747, 2004.

770 Yifeng Zhang, In-Jung Kim, Joshua R Sanes, and Markus Meister. The most numerous ganglion cell type of the mouse retina is a selective feature detector. *Proc Natl Acad Sci U S A*, 109(36):E2391–8, Sep 2012. doi: 10.1073/pnas.1211547109.

Aleksandar Zlateski and H. Sebastian Seung. Image segmentation by size-dependent single linkage clustering of a watershed basin graph. *CoRR*, abs/1505.00249, 2015. URL http://arxiv.org/abs/1505.00249.
775

Author Contributions

KLB acquired the e2198 dataset, and performed preliminary analysis of the calcium imaging data (including manual ROI detection). JAB analyzed arbor segregation and density conservation with input from SM, JSK, and NLT. SM analyzed sustained vs. transient responses. EyeWirers reconstructed neurons with supervision from RP, MS, CD, DLJ, and DB. RP devised strategies for expert correction of crowd wisdom ("reaping"). ALRS managed EyeWire operations and analyzed efficacy of crowdsourcing techniques. SM created code for detecting reconstruction errors by finding collisions between neurons. JSK curated neuron reconstructions, and performed the computational flattening. JSK and SM found correspondences between calcium ROIs and EM cell bodies. JAB, SM, and JSK performed hierarchical clustering. JAB and NLT transformed 3D reconstructions into skeletons. SM performed SAC contact analysis. NLT segmented cell bodies with help from M. Moore, and quantified arbor asymmetry. JSK computed convex hulls. IT, WMS, and ADN created the online museum. WMS created the system by which scythes mark branches as complete, with help from K. Radul. WMS added support for multiple languages with help from M. Balkam and K. Radul. NK devised a new point system for gameplay that better incentivized accuracy, and implemented a chatbot that helped with community management. CSJ added EyeWire features used for collaboration with KT. JMP worked on KT promotions with help from H.-J. Park, S.-H. Seo, J. Hong, E. Bae, and S.-B. Yang. HSS wrote the paper with help from SM, JSK, NLT, JAB, NK, WMS, ALRS, JMP, and KLB.

Acknowledgments

Methods

Countdown to Neuropia

On August 12, 2014, KT Corporation and EyeWire signed a memorandum of understanding in which KT pledged to "fulfill its corporate social responsibility by mobilizing Korean people to participate in EyeWire, thereby using telecommunications technology to advance neuroscience research for the benefit of all humanity." The signing ceremony was held at KT Olleh Square in central Seoul and was covered by over 40 mass media outlets including television, newspapers, and websites. The memorandum laid out the KT-EyeWire plan for Countdown to Neuropia, a campaign to reconstruct ganglion cells in the e2198 dataset.

To prepare for the Countdown, the EyeWire site was translated into Korean and a separate Korean chat channel was created and moderated by Korean-speaking lab members. The Countdown officially launched on October 10, 2014 with four months of nationwide television advertising. Banner ads were posted on the main pages of various portal websites. KT created a microsite to promote EyeWire with prizes. More than 2,600 players participated in a six-week long competition which awarded monetary prizes totaling $50,000. A few top players won the opportunity to visit EyeWire headquarters and several U.S. universities. In addition, KT mobilized an existing group of 280 students from 40 colleges who serve as brand ambassadors of KT. From October to December 2014, these students (known as "Mobile Futurists") both played EyeWire to win prizes, and publicized EyeWire at their colleges. From March to July 2015, competitions were organized at five high schools. A total of 511 students competed to win weekly and overall best rank in their respective schools. KT promotions closed in July 2015 with the completion of Phase 3 of the Countdown.

During the KT promotions from October 2014 to July 2015, there were an estimated 4,271 Korean and 9,532 non-Korean participants. (Participant is defined as a player who submitted a non-tutorial cube. Korean vs. non-Korean was inferred based on IP address, language setting of web browser, and participant lists of KT promotions.) 13,878 Korean players registered, 38 times the number from October 2013 to July 2014. Korean players completed 879,713 cubes, 33% of the total cubes played during the period.

The e2198 GCL patch was divided into four zones with borders defined by three concentric squares with side lengths of 0.05, 0.1, and 0.2 mm. There were 27, 79, 242, and 456 cell bodies in Zones 1 through 4, enumerated from inside to outside. In the first three zones, 7, 16, and 62 had already been reconstructed for previous studies [Kim et al., 2014, Greene et al., 2016] and other preliminary studies in lab. The remaining 20, 63, and 180 cells were reconstructed in three successive Countdown Phases.

The Countdown started in October 2014, and Phases 1 through 3 concluded on November 2014, Feburary 2015, and July 2015. The end of Phase 3 was celebrated by the release of a video showing all cells in Zones 1 through 3 (Supplementary Video). After cells were reconstructed, GCs were distinguished from amacrine cells (ACs) by the presence of axons. The total numbers of GCs in the first three zones turned out to be 13, 43, and 112.

Since Zone 4 contained so many cell bodies, we decided to restrict reconstructions to GCs only. We inspected 456 candidate cell bodies in Zone 4, and identified 228 as GCs by detecting axons. Of these 228 cells, 60 had already been reconstructed previously. The remaining 168 "bonus cells" were reconstructed in the fourth and last "Bonus Phase" of the Countdown, which concluded in November 2015.

GC axons can be challenging to detect when they branch from dendrites rather than directly from the soma. Searching for axons in Zones 1 to 3 was done after all cells were fully reconstructed, and was therefore likely more reliable than identification of GCs in Zone 4, which was done prior to reconstruction. Therefore the false negative rate for Zone 4 GCs may be higher than in the first three zones.

Overall, 396 GCs were reconstructed, with a total path length of roughly 1.52 m. We excluded 11 cells that were severely cut off by the borders of the EM volume, and four more "weirdos" that may have resulted from developmental abnormalities, leaving 381 GCs for subsequent analysis.

Computational flattening and downsampling

The IPL in our volume had some curvature and variations in thickness. We computationally flattened the IPL by deforming it so that the Off and On SAC layers became parallel planes [Sümbül et al., 2014, Greene et al., 2016]. Such flattening has previously been shown to increase the reproducibility of stratification profiles [Manookin et al., 2008, Sümbül et al., 2014, Greene et al., 2016]. We adapted the code developed by Sümbül et al. [2014]. Each SAC layer was quasi-conformally mapped to a plane located at its median IPL depth (Extended Data Fig. 4a). The mappings were extended from the SAC planes to the rest of the IPL by using local polynomial approximations. The transformation was applied to the entire e2198 volume, along with $4\times$ downsampling in each direction.

Stratification profiles

For BCs, we defined the stratification profile as the linear density of arbor volume as a function of IPL depth [Kim et al., 2014, Greene et al., 2016]. For GCs, we defined the stratification profile as the linear density of arbor length as a function of IPL depth. The different definitions (volume vs. length) were chosen to lessen the contribution of the arbor trunk to the stratification profile. For BCs, the caliber of the trunk is often less than the caliber of the branches, so using volume tends to weight the trunk less. For GCs, the caliber of the trunk is generally greater than the caliber of the branches, so using volume tends to weight the trunk more.

The 3D reconstruction of each GC arbor was automatically transformed into a 1D skeleton (Extended Data Fig. 4b), and for each skeleton we computed the density of voxels as a function of IPL depth. This "stratification profile" was treated like a probability distribution, with its area normalized to unity.

Automated skeletonization

GC skeletons were computed from 3D reconstructions as follows. We define two graphs on the voxels of the cell, with edges determined by 26-connectivity. In the undirected graph, the weight of the edge between voxels v and v' is given by the Euclidean distance $d(v, v')$, taking on values 1, $\sqrt{2}$, or $\sqrt{3}$. In the directed graph, the weight of the edge from voxel v to v' is

$$w(v \rightarrow v') = d(v, v') \left(1 - \frac{d(v)}{1.01 d_{max}} \right)^{16} \quad (1)$$

9

Here $d(v)$ is the Euclidean distance from voxel v to the boundary of the cell, which is known as the "distance boundary field" and computed using a Euclidean Distance Transform algorithm from Maurer et al. [2003]. This procedure ignores voxel anisotropy . The maximum of $d(v)$ over all voxels in the cell is d_{max}. Equation (1) differs from Sato et al. [2000], who used a sum rather than a product.

A root is selected as the first voxel in the dictionary ordering of the *xyz* voxel locations, which is typically at the end of a dendrite. A destination is chosen as the farthest voxel in the undirected graph. Then the shortest path between root and destination is computed in the directed graph to yield part of the skeleton. Then the undirected graph is modified by removing all points in cubes centered on the skeleton voxels, where the cube at voxel v has length $6d(v) + 6$. The dependence on $d(v)$ means that the cube is bigger where the dendrite is thicker. The numerical values set the size of "spines" that will be ignored. Based on this modified undirected graph, a new destination point is selected as the farthest point from the root. Then we compute the shortest path from the root to the destination in the directed graph. This process is iterated until no points remain in the undirected graph.

A cell may consist of multiple connected components due to small inaccuracies of the reconstruction process. Each connected component is skeletonized separately, and the stratification profile is computed from the set of skeletons.

k-means clustering in 1D

In our hierarchical clustering, every split was made by applying *k*-means clustering in 1D. We used $k = 2$ for all splits except one, for which we used $k = 3$. The centroids of the clusters were randomly initialized using the method of Arthur and Vassilvitskii [2007].

For some of the high-level splits, we defined the segregation index

$$I_{segregation} = \frac{|\mu_1 - \mu_2|}{\sqrt{(\sigma_1^2 + \sigma_2^2)/2}} \qquad (2)$$

where μ_1 and μ_2 are the centroid locations and σ_1^2 and σ_2^2 are the cluster variances. The segregation index was averaged over 1000 random initializations of the *k*-means algorithm.

Hierarchical clustering of GCs

The hierarchical clustering of GCs is depicted by the dendrogram of Extended Data Fig. 2a. The top levels of the hierarchy are the same as in the smaller dendrogram of Fig. 2j. In the first division, DS cells are separated by cosine similarity with On and Off SACs (Extended Data Fig. 2b). The remaining cells are separated into marginal and central clusters (Fig. 2d). The marginal cluster separates into inner and outer clusters (Extended Data Fig. 2c). The central cluster separates into inner, inner-outer, and outer clusters (Extended Data Fig. 2d). The preceding yields six high-level clusters, which are further divided into 47 clusters based mainly on features computed from the stratification profile (Extended Data Figs. 4c-f). Soma size, SAC contact, and arbor density and complexity are used for a few divisions.

From the outer central cluster, 5to separates by 10th percentile of IPL depth (split b-1). Then, 4ow separates by its large soma size (split b-2). The remaining cells separate into 4on and 4i based

on the difference between the 70th and 15th percentile IPL depths (split b-3).

From the inner-outer central cluster, 63 separates via cosine similarity with the Off SAC stratification profile (split c-1). From the remaining cells, 5si and 5so separate via 5th percentile of IPL depth (split c-2), and are distinguished from each other based on the 80th percentile IPL depth (split c-4). The remaining cells divide into 51 and 5ti by arbor complexity (split c-3).

From the inner central cluster, 6t separates via 95th percentile IPL depth (split d-1). The remaining cells divide into 6sw and 6sn based on soma size (split d-2).

From the outer marginal cluster, 27 and 28 separate via 85th percentile IPL depth of the stratification profile restricted to the marginal IPL (split e-1), and are distinguished from each other via 95th percentile IPL depth (split e-12). From the remaining cells, 1wt and 2o separate based on soma size (split e-2), and are distinguished from each other via arbor complexity (split e-11). 1ws, 1no, and 1ni can be separated via 50th percentile IPL depth (split e-3) because their dendrites are adjacent to the INL. 1ws can be distinguished from 1no and 1ni by arbor density (split e-5), and 1no and 1ni can be separated via 85th percentile IPL depth (split e-7). From the remaining cells, 2an and 25 separate based on arbor complexity (split e-4), and are distinguished from each other by the difference of the 90th and 45th percentile IPL depths (split e-10). From the remaining cells, 3i and 3o separate via the difference between the 80th and 10th percentile IPL depths of the stratification profile restricted to the outer IPL (split e-6), and are distinguished from each other by the difference between the 90th and 10th percentile IPL depths of the same restricted stratification profile (split e-9). The remaining cells divide into 2aw and 2i based on arbor density (split e-8).

From the inner marginal cluster, 8w separates via soma size (split f-1). From the remaining cells, 72, 73, 81o, and 82wo separate by 95th percentile IPL depth (split f-2), and 81o and 82wo separate from 72 and 73 by 80th percentile IPL depth (split f-6). 72 and 73 are distinguished from each other via straticiation profile area of central region (split f-9), while 81o and 82wo are distinguished from each other via 25th percentile IPL depth (split f-10). From the remaining cells, 85 separates via the inner central stratification profile area (split f-3). From the remaining cells, 8n,9n, and 9w separate by 5th percentile IPL depth (split f-4), and 8n is distinguished from 9n and 9w by 50th percentile IPL depth in detailed resolution (split f-8). Then, 9n and 9w are distinguished from each other by dendritic field area (split f-13). Among the remaining cells, 915 and 82n have extra arbors in the central region so they are separated via the central stratification profile area (split f-5), and distinguished from each other by the difference between marginal and central stratification profile areas (split f-12). 91 is separated from the rest of the cells by the 50th percentile IPL depth of the stratification profile restricted to the inner region (split f-7). The remaining cells are clustered into 81i and 82wi via 25th percentile IPL depth (split f-11).

Soma size

Semiautomated reconstruction in EyeWire is based on a convolutional network described previously [Kim et al., 2014]. This convolutional network worked well in the IPL, but was inaccurate in the GCL because of large gaps in boundaries between cell bodies caused by incomplete staining. To segment the GCL,

10

we started with a $2\times$ downsampled image. Then we applied minimum filtering with a $9 \times 9 \times 1$ sliding window and a stride of 5 in all dimensions. This yielded a new image with $10\times$ downsampling in each dimension relative to the original image. The convolutional network was applied, and its output was segmented using a modified watershed transform (http://github.com/seung-lab/watershed and Zlateski and Seung, 2015). A human annotator examined all ganglion cell bodies and removed all but small parts of the primary dendrites. Often no action was required, as almost all of the dendritic arbor of a cell was generally split off from the cell body due to the downsampling.

SAC contact analysis

On DS and On-Off DS types separate from each other via SAC contact analysis. The flattened and downsampled EM volume between IPL depths 0.1 and 0.8 was divided into a regular grid of rectangular cuboids. Each cuboid was 15×11 voxels (roughly 1 μm) in the tangential plane, and 181 voxels (12 μm) along the light axis. The grid was exactly two cuboids deep along the light axis. One cuboid was outer IPL (depth 0.1 to 0.45) and the other cuboid was inner IPL (depth 0.45 to 0.8).

For each GC, we examined all reconstructed SACs and recorded two sets of SAC voxels. Firstly, we found all *contacting* voxels, defined as SAC surface voxels contacting the GC. Secondly, we found all *collocating* voxels, defined as SAC surface voxels in the grid cuboids occupied by the GC. The contacting voxels are a fraction of the collocating voxels. This "SAC contact fraction" has a numerical value with the following interpretation. If a portion of a SAC dendrite intermingles with the arbor of a GC arbor that has SAC contact fraction f, on average the dendrite portion will devote a fraction f of its surface area to contact with the GC.

For each contacting or collocating voxel, we recorded the direction from the corresponding SAC soma centroid to the voxel in the plane perpendicular to the light axis (Extended Data Fig. 6a). Based on these directions, the voxels were divided into 8 bins equally spaced on the circle. For each bin, we computed the ratio of contacting voxels to collocating voxels. This yielded SAC contact fraction versus direction, shown in the polar plots of Extended Data Fig. 6. The overall preferred direction of SAC contact for a GC is computed by taking the vector mean of the polar plot for that GC.

The SAC contact fraction for a GC is a normalized rather than absolute measure of SAC contact. The normalization is intended to make the SAC contact fraction robust to incomplete sampling of SAC dendrites. Reducing the number of SACs in the analysis tends to reduce both contacting and collocating voxels by the same factor, leaving the ratio unchanged. Our sample of SAC dendrites is biased, because it contains no SAC dendrites that come from a SAC soma outside the patch. The bias is least for GCs near the center of the retinal patch, and greatest for GCs near the borders. Our normalization procedure does a good job of correcting for the biased sampling, as shown by the reproducibility of the polar plots within each type in Extended Data Fig. 6.

Arbor density, complexity, and asymmetry

We projected the voxels of each cell's volumetric reconstruction onto the 2D plane orthogonal to the light axis, and found the convex hull of this projection (Extended Data Fig. 7a). The arbor density was then computed as the length of arbors divided by the area of the convex hull. The length of arbors is calculated from the number of skeleton voxels scaled by the correction factor, which is the average ratio of arbor length and the number of skeleton nodes of 381 cells. For each skeleton, a branch point was defined as a voxel with degree greater than or equal to 3. Path length was defined as the sum over all edge lengths in the skeleton, where edge length is the Euclidean distance (corrected for voxel anisotropy) between the pair of nodes joined by an edge. Arbor complexity was defined as the number of branch points in a skeleton divided by its path length.

The arbor vector was defined as pointing from the soma centroid to the skeleton centroid, after projection onto the 2D plane orthogonal to the light axis and scaling the y and z coordinates to match in physical length (Extended Data Fig. 7d). For each type, we computed the mean of arbor vectors over the cells in the type. To quantify asymmetry for an individual cell, we constructed a plane through the soma centroid that was perpendicular to the mean arbor vector for its type. An index of asymmetry was defined as

$$I_{\text{asymmetry}} = \frac{P - N}{P + N} \tag{3}$$

where P is the amount of skeleton on the same side of the plane as the mean arbor vector, and N is the amount of skeleton on the other side. Statistical significance of arbor asymmetry for a type was assessed by applying the Rayleigh z-test [Brazier, 1994] to the arbor vectors. The resulting p values were evaluated under a Bonferroni correction. Note that arbor asymmetry was only computed after the clustering was complete; it was not used to create the clusters.

Coverage factor

Each cell's dendritic field was approximated as the 2D convex hull of its volumetric reconstruction. The coverage factor for a cluster was computed as the sum of cells' hull areas divided by the area of the union of hulls. Border effects will lead to underestimation of the coverage factor, similar to the above-mentioned underestimation of the arbor density. To minimize border effects, we considered only the intersection of each convex hull with the crop region.

Density conservation test

Before quantifying density conservation, dendritic trunks were removed by eliminating skeleton nodes with IPL depth greater than the peak of the stratification profile plus 0.1.

Our sample is missing those GC dendrites inside the e2198 volume that come from cell bodies lying outside the volume. The missing dendrites lead to underestimation of aggregate arbor density. To minimize this border effect, we cropped away 1000 pixels (65 μm) on all four sides of the e2198 patch (Fig. 4d). The remaining "crop region" was divided into unit grid squares of 600×600 pixels, with physical dimensions of (40 μm)^2. For each cluster, the aggregate arbor density was computed for each grid square (Fig. 4e, h). To quantify density conservation for a cluster, we computed the coefficient of variation (CV), defined as the standard deviation divided by the mean, for the densities of the grid squares. A small CV meant that the aggregate arbor density was approximately uniform.

To quantify statistical significance for each cluster, we needed to compare with randomized configurations of the arbors in the clus-

ter. Naively we would have done this by randomly relocating every cell's soma to a different location within the patch. However, this naive randomization would have typically decreased the amount of arbor within the patch, by causing some full branches to extend outside the patch, and some cutoff branches to terminate in the interior. Therefore, we constrained the soma relocation to approximately preserve the total amount of dendrites within the patch. For each soma, we drew a rectangular "orbit" that preserved the distance of the soma from the border of the patch (Fig. 4g). Each cell's soma was randomly relocated to another position on its orbit. The arbor was also rotated by some multiple of 90° so that the cutoff branches remained oriented towards the border of the patch (Fig. 4g).

The true CV was compared with the CVs of 10,000 randomized configurations. A cluster was judged to have passed the significance test if its CV was less than 99% of randomized configurations.

Analysis of calcium imaging data

Imaging was performed in a single focal plane for each tile in a 3×3 array of tiles. For each tile, the location of the moving bar was adjusted so that it was centered on the tile. Then the responses of all cells in the tile were recorded for this common stimulus. The bar's speed, width, and location in the receptive field could not be tuned on a cell-by-cell basis to optimize visual responses, as is commonly done with intracellular recordings.

ROI extraction and detrending

Ellipsoidal ROIs were manually drawn around fluorescent cells. All pixels in an ROI were summed to yield fluorescence versus time for that cell. We converted to fractional fluorescent change ($\Delta F / F$) relative to the mean computed over the entire recording. We then detrended the fractional fluorescence change by subtracting a Gaussian smoothed version of itself with $\sigma = 20$s, to reduce slow drifts in the baseline.

Correspondence between ROIs in calcium imaging and cells reconstructed in EM were established as follows: First, starting from the somata of On-Off DS cells from our classification, two experts visually examined their surroundings in the ganglion cell layer in the EM volume. Using the vasculature in the images as the guiding landmark and the shapes and arrangements of neighboring cells as the deciding features, 25 corresponding ROIs were identified in the fluorescent image. EM coordinates were recorded for roughly the center point of each of these somata, and fluorescent image coordinates were recorded for the center of the corresponding ROIs. A linear transformation between these two sets of coordinates is then computed in the least squares sense from their respective 3D or 2D homogeneous coordinates. Using this transformation, we computed the "nominal" EM coordinates for the center point of each of the ROIs. Finally, a human expert examined each of these EM locations to identify the correct cell that corresponds to the original ROI, similar as in the initial step.

Response quality index

To quantify reproducibility of visual responses, we used the quality index of Baden et al. [2016], defined as

$$QI = \frac{\mathrm{Var}\left[\langle C \rangle_r\right]_t}{\langle \mathrm{Var}\left[C\right]_t \rangle_r} \qquad (4)$$

where C is a $T \times R$ "response matrix" and $\langle \rangle_x$ and $\mathrm{Var}[]_x$ denote the mean and variance across the indicated dimension, respectively. The number of time steps per trial is T and the number of trial repetitions is R. Cells were included in the sustainedness index calculations of Figs. 6c, 6e and Extended Data Fig. 9 only if the quality index exceeded a threshold of 0.5.

Our recording consisted of 5 trials. Each trial contained 8 stimuli corresponding to 8 directions of motion presented in the same order (Figs. 1d, e). The recording was continuous; there was no gap between consecutive trials or stimuli. For all analyses we discard the first trial, which exhibited adaptation to the stimulus, and use only the remaining $R = 4$ trials for the response matrix C.

Each of the eight stimuli in a trial lasted for 31 time steps. For computing the quality index of On responses, the response matrix omitted all but the 4th to 15th time steps of each stimulus, so that each trial had $T = 96$ time steps (12 time steps by 8 directions). For computing the quality index of Off responses, the response matrix omitted all but the 15th to 31st time steps, so that each trial had $T = 136$ time steps (17 time steps by 8 directions). For computing the quality index of general responses not specific to On or Off (used by Fig. 5c, h, with the same threshold of 0.5), the response matrix contained all time steps, so that each trial had $T = 248$ time steps (31 time steps by 8 directions).

Sustainedness index

To compute the sustainedness index of a cell in Fig. 6, we define a reference time at 0.8 sec after the nominal stimulus onset (just before stimulus offset). For the On response, we normalize the temporal response function such that the minimum is 0 and the maximum at or before the reference time is 1, and then we define the On sustainedness index as the value at the reference time. For the Off response, we normalize the temporal response function such that the minimum is 0 and the maximum after the reference time is 1. We similarly define the Off sustainedness index as the value at 0.8 sec after the nominal stimulus offset time. For inner-outer central and inner cell clusters, we computed the On sustainedness index; for outer cell clusters, we computed the Off sustainedness index (Fig. 6e).

A cell with a larger dendritic arbor might spuriously appear more sustained if that means the stimulus exits the receptive field later. To compensate for this effect, we estimate the additional time taken by the leading or trailing edge of the moving bar to exit the dendritic field. For a given low-level cluster in our hierarchy, we find the cell with the maximum diameter in the cluster, and compute the time for the leading or trailing edge of the moving bar to exit the dendritic field as half of the arbor diameter divided by the speed of the moving bar. To the "exit time" for that cluster we add 0.6 seconds, which yields a cluster-specific reference time. The cluster-specific reference times vary between 0.6 to 0.8 seconds after the nominal stimulus onset or offset time. For each cluster, we computed the sustainedness index as before, but using the cluster-specific reference time (Extended Data Fig. 9). Linear in-

terpolation was performed for adjusted time points falling between frames.

Comparison of sustainedness index for high-level clusters

One-way ANOVA (Matlab `anova1`) confirmed that the high-level clusters have different mean sustainedness indices ($p < 0.01$, $df_{error} = 172$, $df_{total} = 176$) (also see Fig. 5c). Post hoc pairwise comparison between groups were performed with Tukey's honest significant difference criterion (Matlab `multcompare`).

Cluster-average temporal response functions

Temporal response functions (as defined in Fig. 1e) of individual cells were first normalized in the same way as when computing their sustainedness index, and then averaged to produce cluster averages (Figs. 5c, h, 6a, d, Extended Data Fig. 8). When standard deviation or standard error is computed and shown in a figure (Figs. 5c, h, 6a, d), the cluster average is rescaled again to have a minimum of 0 and maximum of 1 (which could have already been the case if cells in a cluster are very consitent with each other in their responses), and the standard deviation and standard error shown are scaled using the same scaling factor.

Inference of the e2198 retinal patch orientation

The orientation of the compass rosette in Fig. 1a was inferred as follows. We used the average arbor vector for 2an as an estimate of the ventral direction, because 2an corresponds to a genetically defined type (F-mini[Off]) that is known to have ventrally directed dendrites [Rousso et al., 2016]. We used the preferred direction (PD) for motion of 7o to define the rostral direction on the retina (posterior direction in the world), based on the assumption that the three 7i types correspond with the classical On DS cells, which have PDs in the superior, inferior, and anterior directions in the visual world [Oyster and Barlow, 1967]. Equivalently, the PD of SAC contact of 7o defines the caudal direction on the retina, as it is anti-parallel to the PD for motion. The end result is that 2an dendrites point roughly toward 240°, and 7o cells have SAC contact PDs of ~170° (Extended Data Fig. 6d) and physiological motion PDs of ~340°.

Code availability

Code for the analyses of the 3D cell reconstructions and calcium imaging data will be made available on Github at https://github.com/seung-lab/e2198-gc-analysis upon publication.

Data availability

Cell reconstructions and other data in summary forms can be viewed at http://museum.eyewire.org/. Raw calcium imaging data will be made available together with the analysing code at https://github.com/seung-lab/e2198-gc-analysis. EM volume and reconstruction data will be made available at http://seunglab.org/data/.
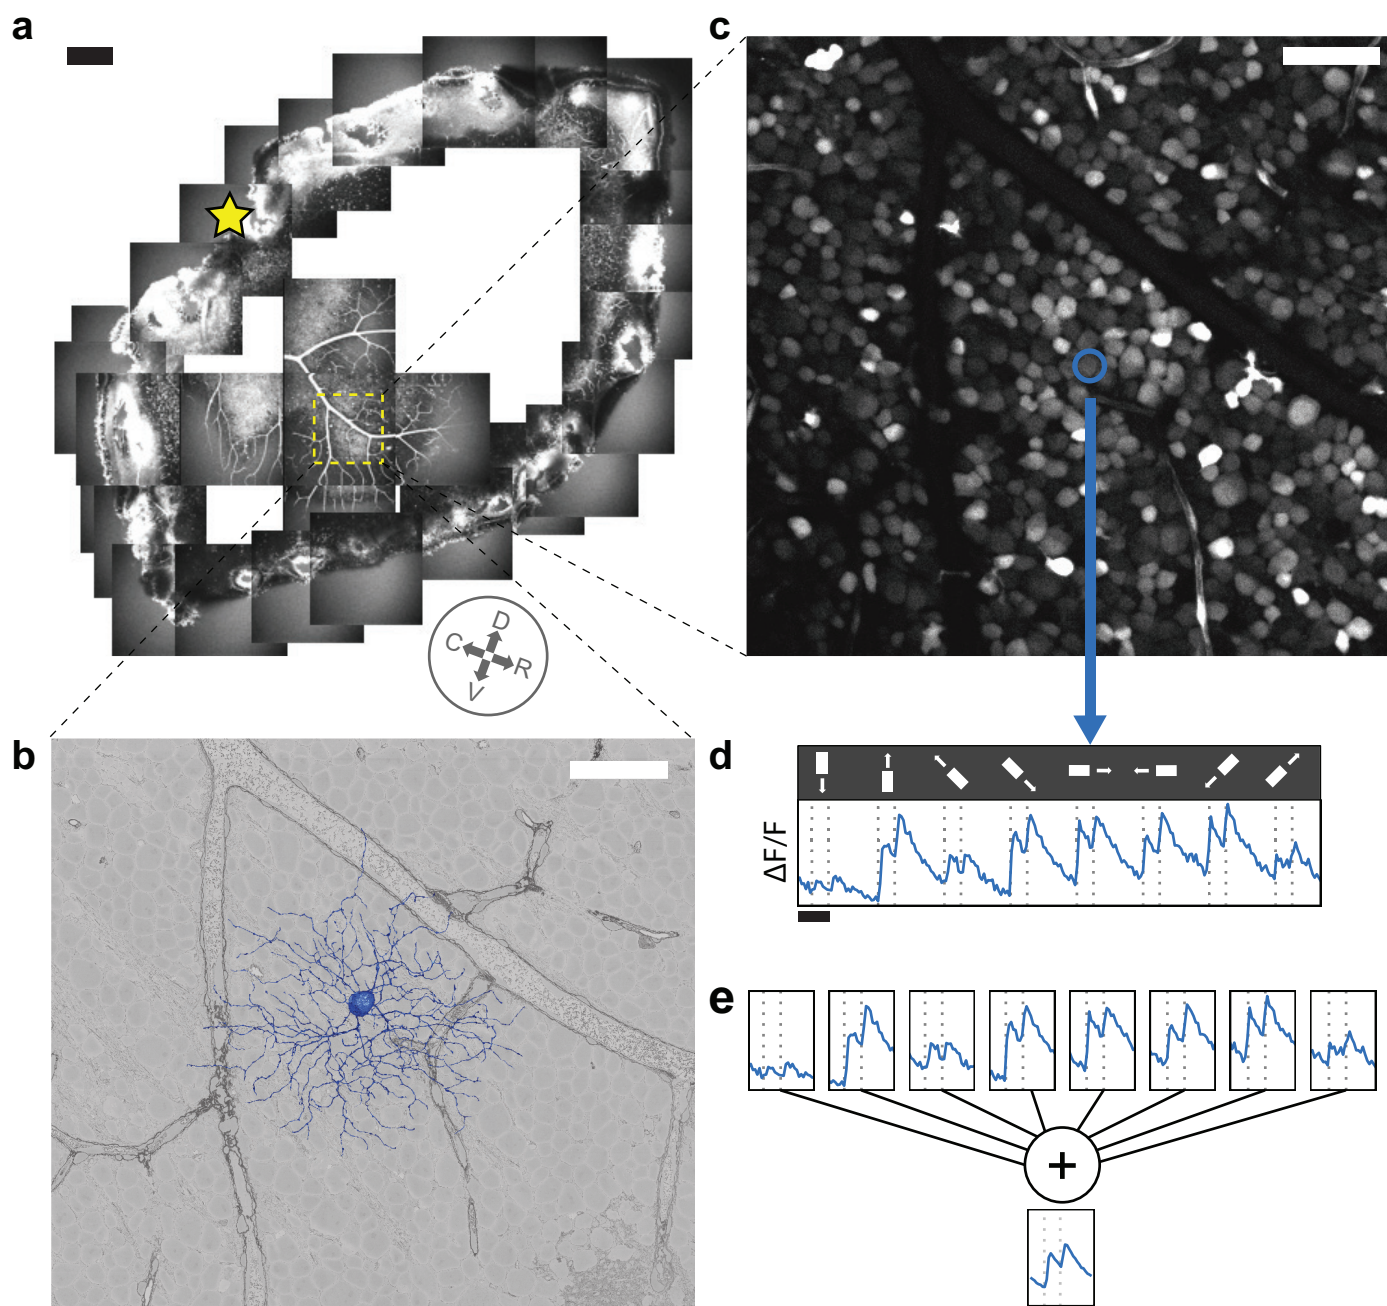
Figure 1: Anatomy and physiology of retinal ganglion cells via electron and light microscopy. **a,** Hemiretina containing imaged 0.3×0.35 mm^2 patch (yellow square). Star, optic disk. Compass rosette, inferred cardinal directions (dorsal, ventral, rostral, caudal, see Methods). **b,** 3D reconstruction of GC dendritic arbor (blue) and 2D cross section through GCL in serial EM image (grayscale). **c,** Soma of same GC (blue) in image of GCL obtained via two-photon microscopy. **d,** Fluorescence versus time for same GC along with stimulus sequence of light bar moving in eight directions on dark background (see main text for true aspect ratio of bar). **e,** Averaging over stimulus directions (shown) and trials (not shown) yields temporal response function for GC. Scale bars, 200 μm (**a**), 50 μm (**b**, **c**) and 2 sec (**d**).
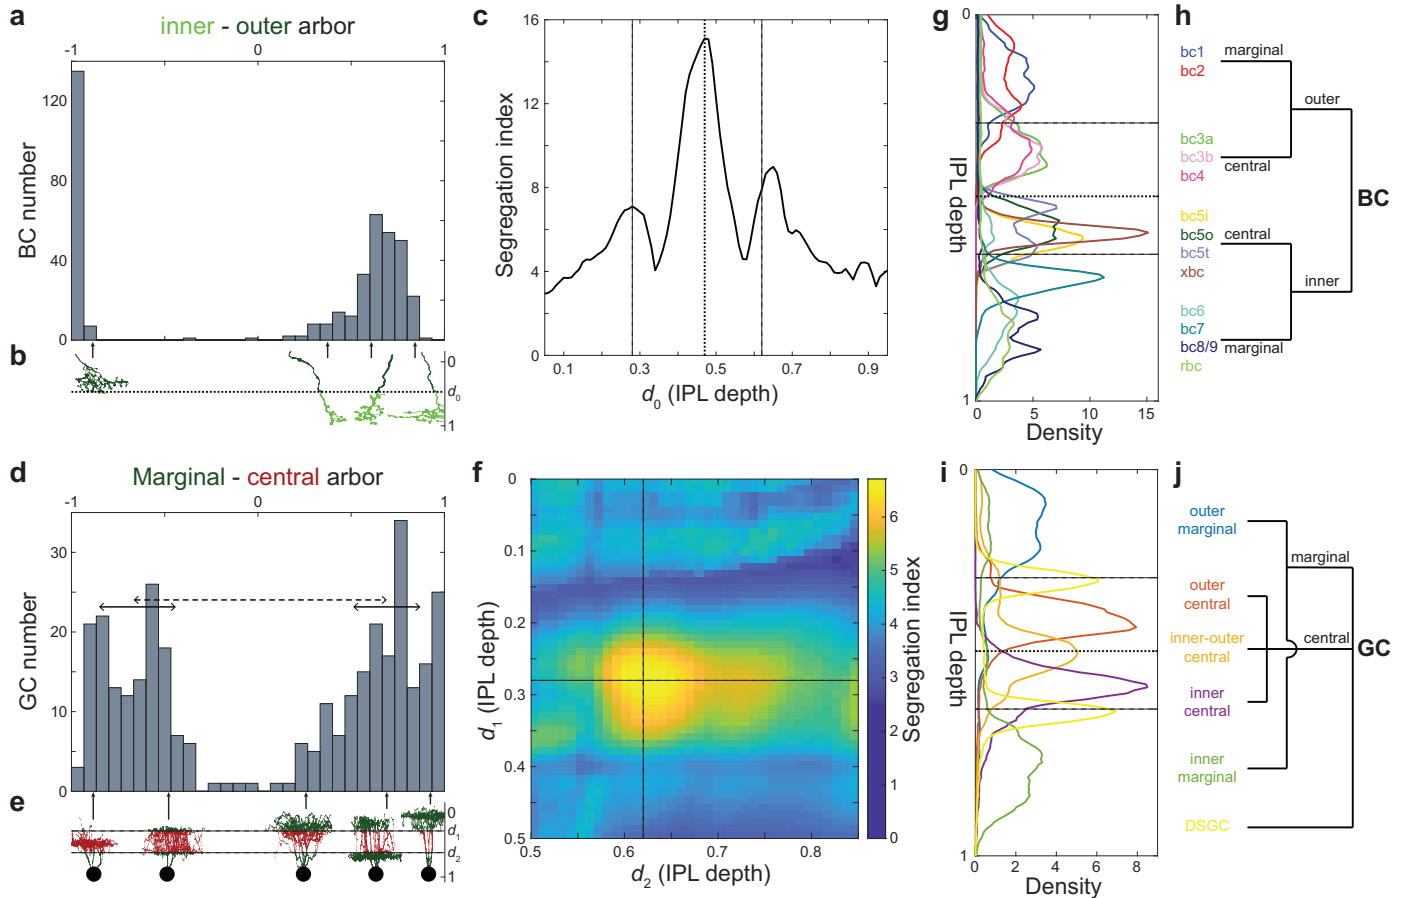
Figure 2: Maximizing laminar segregation of arbors yields marginal-central and inner-outer divisions of the IPL. **a**, Histogram of the difference between inner and outer arbor volume for BCs (total volume normalized to one). BC axonal arbors are either mostly inner (right cluster) or mostly outer (left cluster); intermediate cases are rare. **b**, Inner arbor (light green) and outer arbor (dark green) of example BCs. The depth of the inner-outer boundary is denoted by d_0. **c**, BC inner-outer segregation is maximized for $d_0 = 0.47$ (dotted line, same value used in **a**), with two flanking local maxima at or near the SAC depths (dashed lines). **d**, Histogram of the difference between marginal and central arbor length for GCs (total length normalized to one). GC dendritic arbors are either mostly marginal (right bump) or mostly central (left bump). The segregation index is defined as the separation between the clusters (dashed line), divided by the square root of the average of the half widths of the clusters (full widths are solid lines). **e**, Marginal arbor (green) and central arbor (red) of example GCs (aspect ratio of the cells are distorted for visualization). The IPL depths of the marginal-central boundaries are denoted by d_1 and d_2. **f**, GC marginal-central segregation index is maximized for d_1 and d_2 at the SAC depths (dashed lines, same values used in **d**). **g**, Average stratification profiles (linear density of arbor volume vs. IPL depth) of BC types. **h**, BC types belong to four high-level BC clusters created by inner-outer and marginal-central splits. **i**, Average stratification profiles (linear density of arbor length vs. IPL depth) of six high-level GC clusters that emerge from the first splits (**j**) in the hierarchical clustering of GCs.
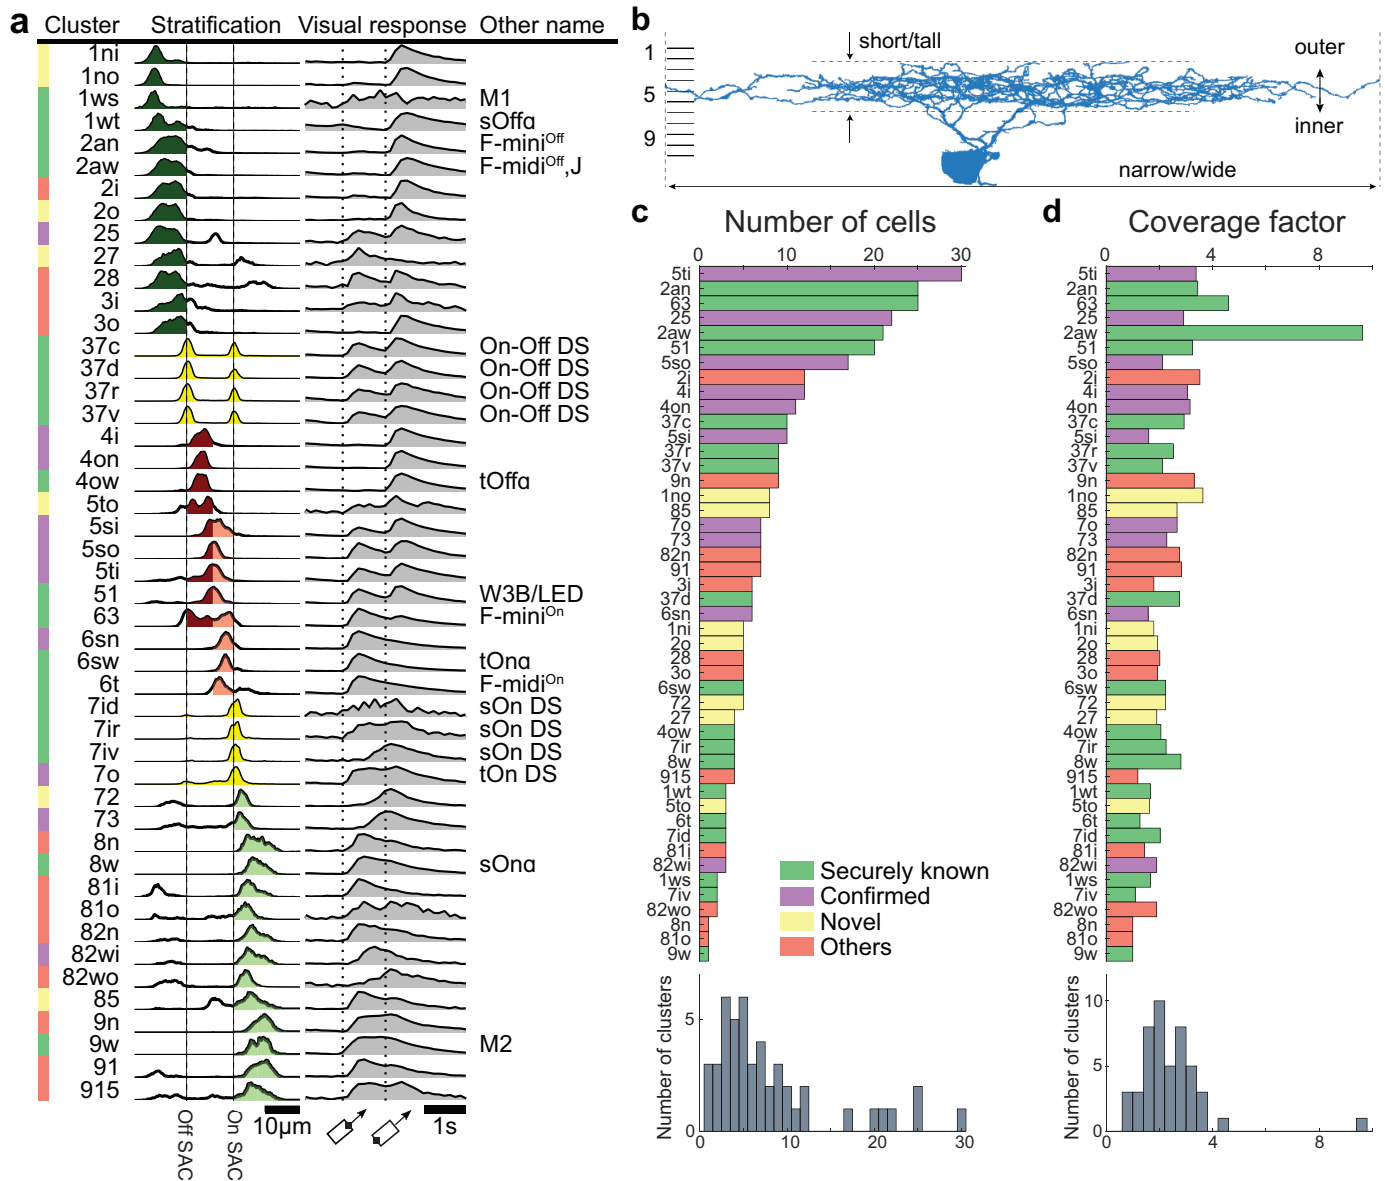
Figure 3: Classification of ganglion cells. **a**, Summary of clusters with anatomical name, stratification profile, and temporal response function defined in Fig. 1e. Alternative names in black are "securely known" types (see main text for definition). **b**, Each cluster name begins with a number in the range 1-9 indicating which tenth of the IPL depth contains the most stratification profile area. More numbers are appended for multistratified clusters. Letters (s, t, n, w, o, i, a) are added to distinguish between clusters with similar stratification, where "a" denotes asymmetric arbor. **c**, Number of cells in each cluster. **d**, Coverage factors.
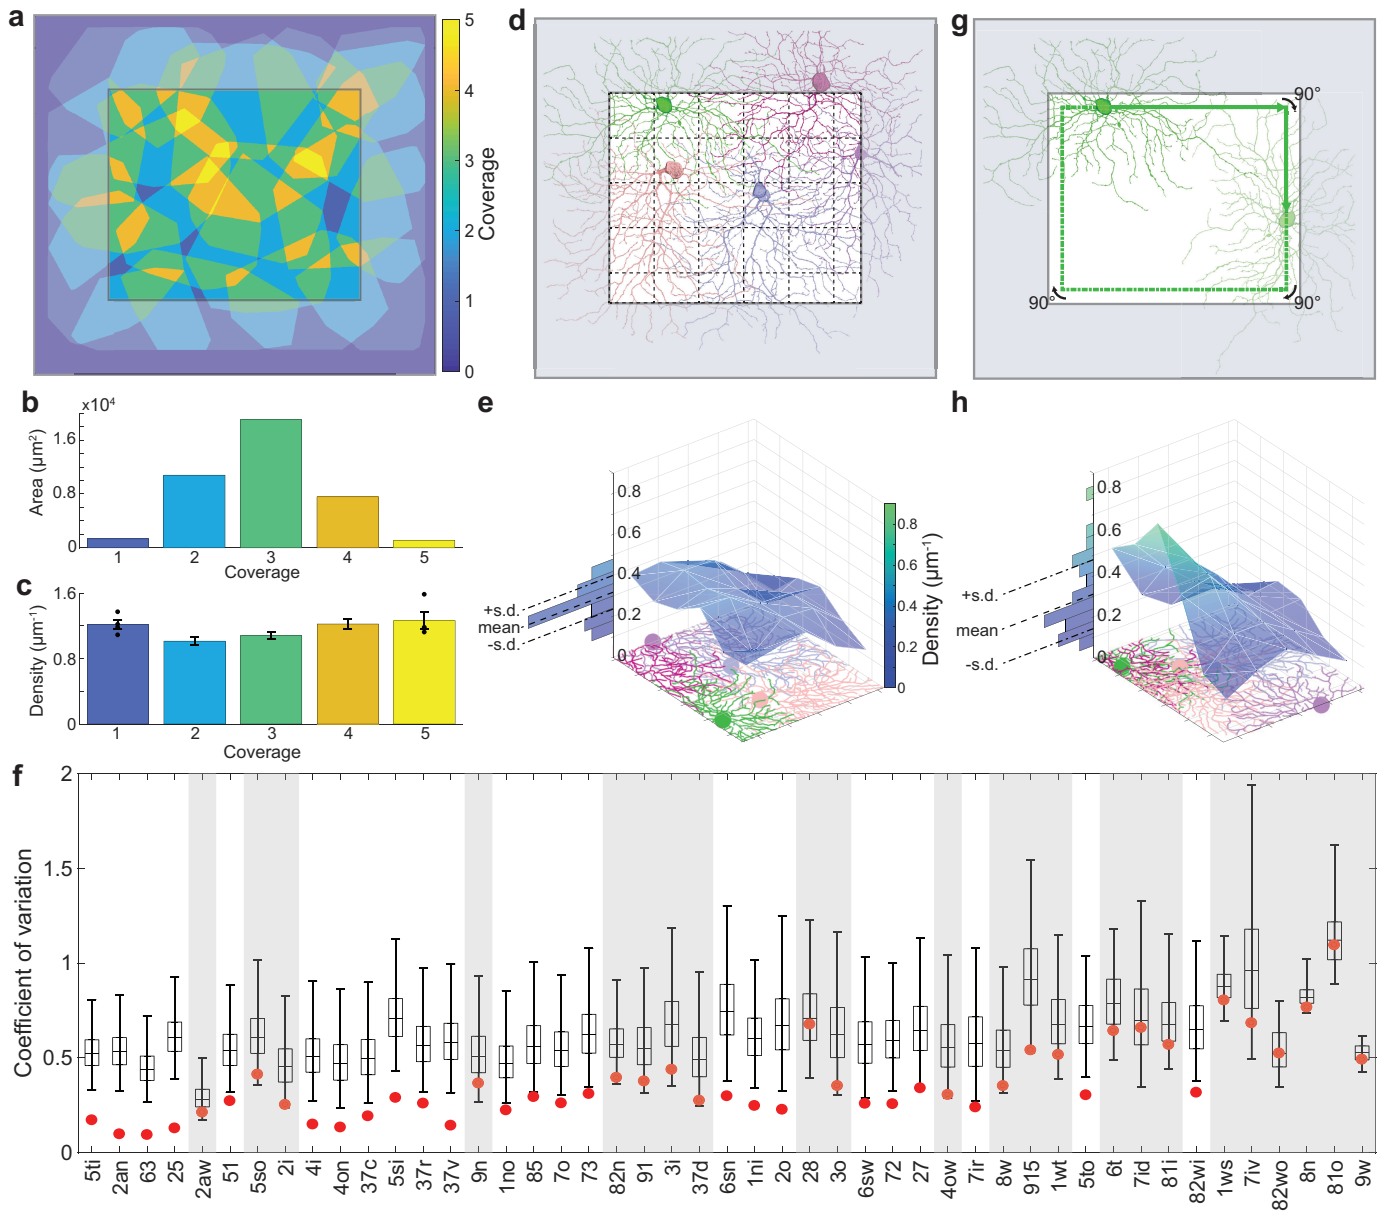
Figure 4: According to our density conservation principle, the arbors of a GC type should have an aggregate density that is approximately uniform. **a**, Arbor convex hulls of an example cluster (25) overlap substantially. Colors indicate how many hulls cover each retinal location inside the crop region. **b**, Retinal area versus coverage inside the crop region. Each bar represents the area devoted to the corresponding color/coverage in the crop region. **c**, The aggregate arbor density of the cluster varies relatively little with coverage. Each bar represents the density within the area devoted to the corresponding color/coverage in the crop region (standard error, $n = 4, 19, 33, 20, 4$). **d**, The crop region is divided into grid boxes, and the aggregate arbor density is computed for each box, as illustrated for an example cluster (6sw). **e**, The aggregate arbor density is close to uniform across the crop region, as quantified by the coefficient of variation (standard deviation divided by mean). **f**, The density conservation test is satisfied by a cluster (non-shaded) when the coefficient of variation is significantly smaller for the real configuration (red dot) than for 99% of all randomized configurations (99/1 percentiles, black bar; quartiles and median, box; $n = 10,000$). **g**, To test statistical significance, the arbors of a cluster are randomized by relocating the soma somewhere on its "orbit" (green line) and rotating the arbor to have the same orientation relative to the nearest side of the retinal patch. **h**, The aggregate arbor density typically varies more after randomization. Example cluster is 25 in **a-c** and 6sw in **d**, **e**, **g**, **h**.
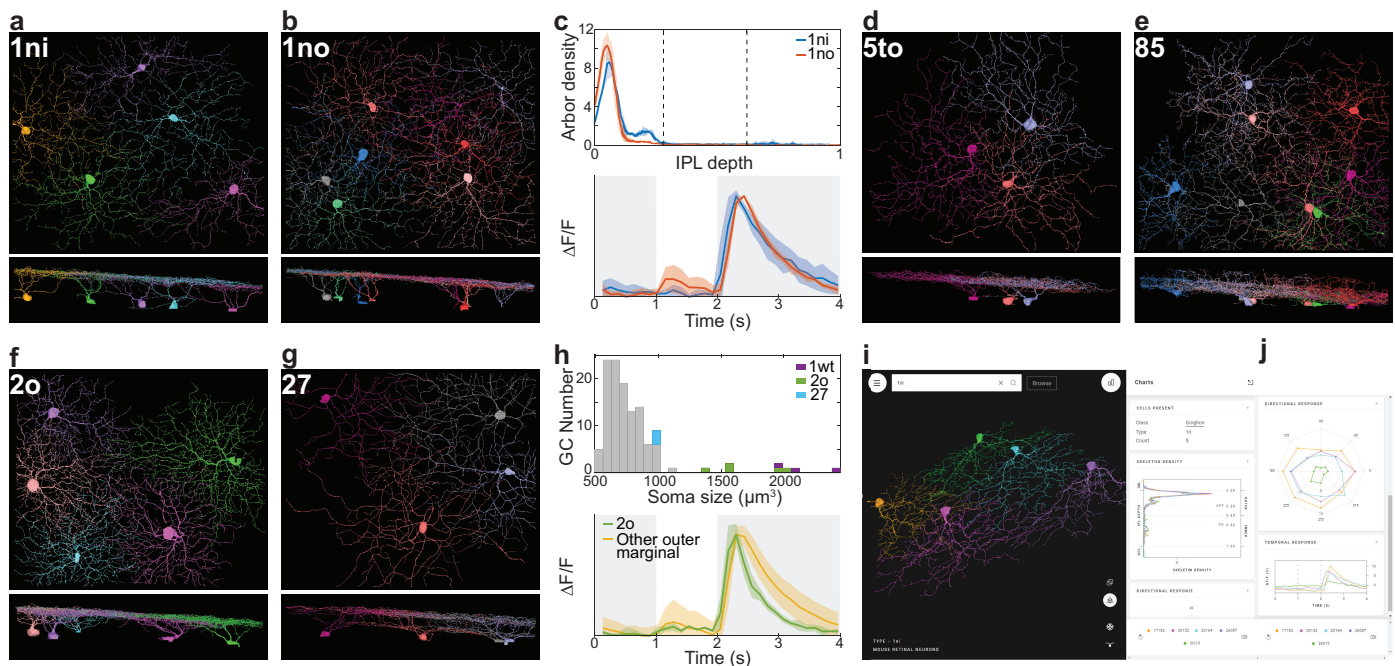
17

Figure 5: Novel ganglion cell types, views along the light axis and a tangential axis. **a-c**, 1ni and 1no are types with very similar stratification profiles (**c**, top) and temporal response functions (**c**, bottom). **d**, 5to looks monostratified in the tangential view but its stratification profile (Fig. 3a) is surprisingly complex. **e**, 85 stratifies throughout the inner IPL but also extends sparse branches towards the INL. **f-h**, 2o and 27 are outer marginal types. Histogram of soma size for outer marginal cells shows that 2o somas are much larger than those of 27 and other typical cells, and smaller than 1wt (transient Off alpha) somas (**h**, top). The Off response of 2o decays more rapidly than that of other outer marginal cells (**h**, bottom). **i,j**, Screenshot of EyeWire Museum (`museum.eyewire.org`), where 3D renderings of all cells in our sample (**i**) can be viewed along with their stratification profiles, directional tuning curves, and temporal response functions (**j**). Shaded regions around curves in **c** and **h** represent standard deviations.
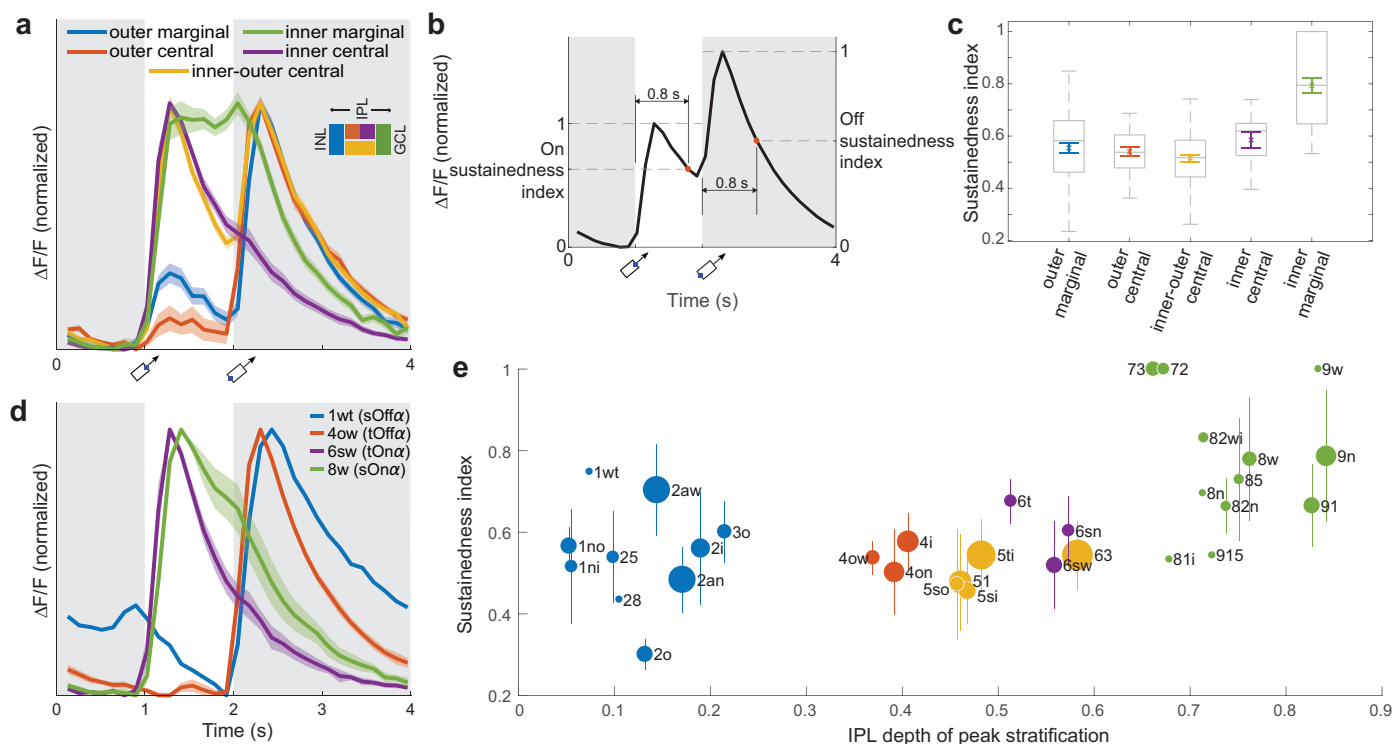
Figure 6: Sustainedness of visual responses and dendritic stratification. **a**, Average temporal response function (Fig. 1e) for high-level GC clusters (Fig. 2i). Each response function is averaged over cells in a cluster, and normalized to have the same maximum and minimum. The inner marginal cluster is markedly more sustained than the others. Shading indicates standard error ($n = 102, 26, 78, 12, 55$ for outer marginal, outer central, inner-outer central, inner central, inner marginal). **b**, The sustainedness index is defined as the response at 0.8 sec after nominal stimulus onset, divided by peak response in the 0.8 sec interval. **c**, The cells in the inner marginal cluster are significantly more sustained than cells in the other clusters (ANOVA with post hoc $p < 0.01$). The differences between other clusters are not statistically significant. Colored, mean and standard error; grey box and bars, quartiles, median, and extrema. **d**, Marginal is more sustained than central for the four alpha types. Shading indicates standard error of the mean ($n = 4, 5, 4$ for 4ow, 6sw, 8w). There is no standard error for 1wt, because only a single 1wt cell had calcium signals. Applying t-tests to sustainedness indices yield $p = 0.02$ for 1wt against 4ow, and $p = 0.03$ for 8w against 6sw. **e**, Sustainedness index for cells with high response quality (Methods, Eq. 4), sorted by clusters. Bars indicate standard deviations for the clusters, except for clusters containing only a single cell with high response quality (1wt, 28, 81i, 8n, 915, 9w). Dot area indicates the number of cells in the cluster; the largest dot (63) represents 18 cells. Size-adjusted sustainedness index for individual cells is shown in Extended Data Fig. 9.

19
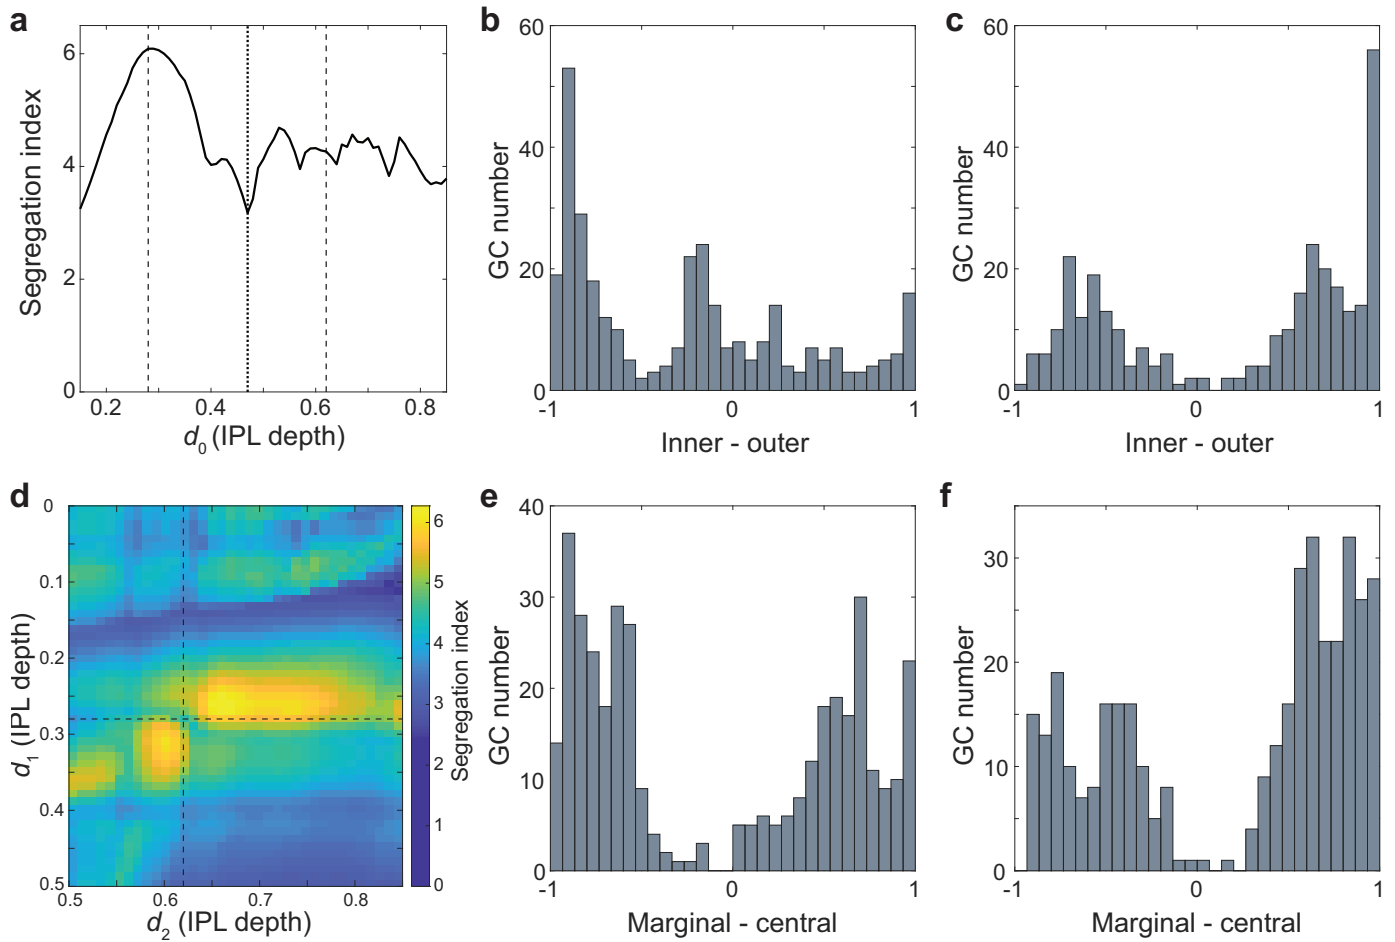
Extended Data Figure 1: GC arbor segregation. **a**, GC segregation index versus inner-outer boundary location. **b**, GCs differ markedly from BCs, as inner-outer segregation is very poor when the boundary is at 0.47 IPL depth. **c**, Maximal inner-outer segregation is achieved by placing the boundary at 0.28 IPL depth, but the result is not as good as the best marginal-central segregation (Fig. 2d, f). **d**, If DSGCs are not excluded, the GC segregation index exhibits two local maxima. The marginal-central boundaries are just outside the SACs in one maximum (**e**), and just inside the SACs in the other maximum (**f**). The only difference is that DSGCs switch from the marginal cluster to the central cluster, supporting the idea that their assignment is arbitrary.
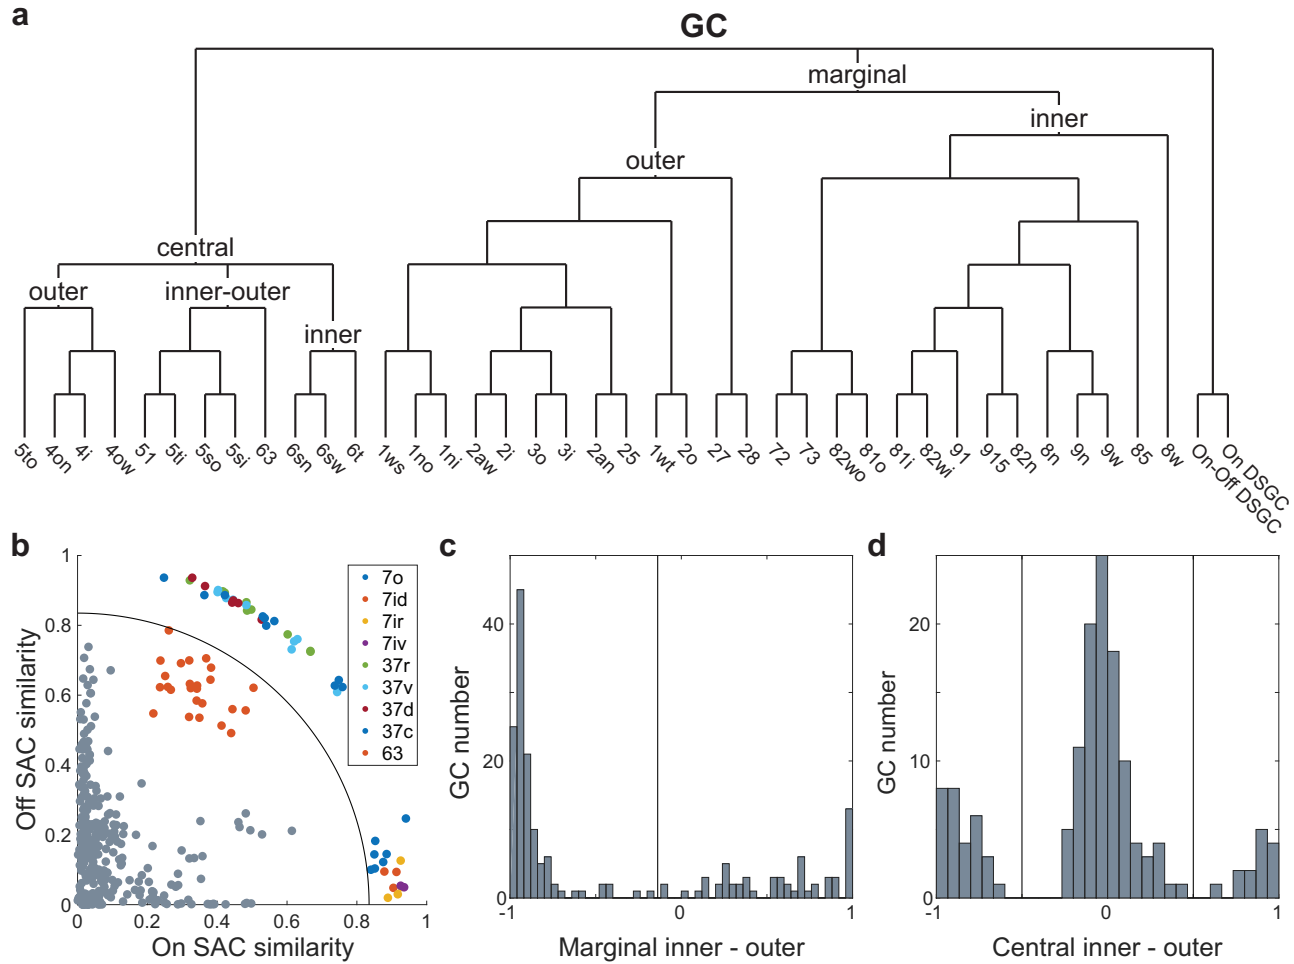
Extended Data Figure 2: Hierarchical clustering of GCs. **a**, The GC hierarchy starts with the high-level clusters of Fig. 2j, which divide into 47 low-level clusters. **b,** On-Off and On DS cells (37 and 7 types, outside the quarter-circle) are the most similar in stratification to SACs, as shown by a scatter plot of cosine similarity with Off SAC vs. On SAC stratification profiles. After removing DS cells, the remaining GCs fall into two clusters that stratify predominantly either in the marginal or central sublamina (Fig. 2d). **c,** The marginal cluster divides into two clusters, based on the difference between arbor length in the outer marginal and inner marginal sublamina. **d,** The central cluster divides into three clusters, based on the difference between arbor length in the inner central and outer central sublaminae.
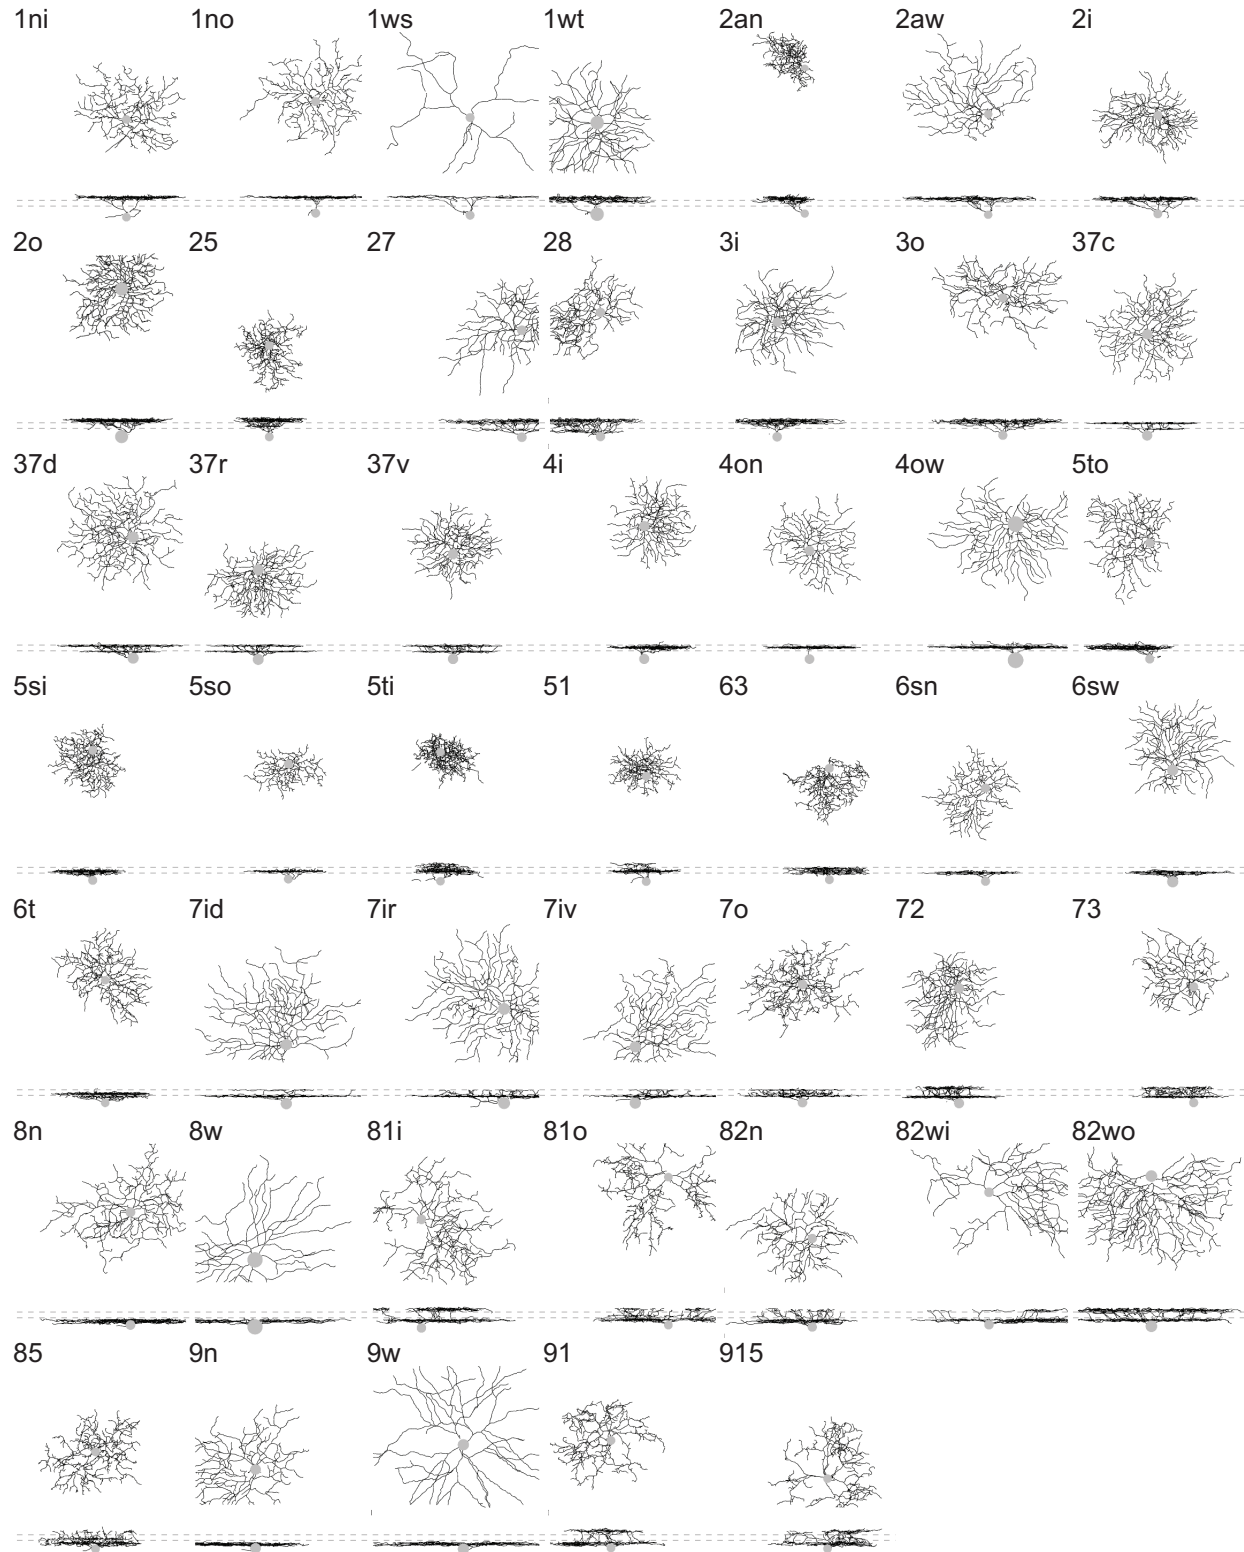
Extended Data Figure 3: Gallery of ganglion cell clusters, light axis (top) and tangential (bottom) views. Renderings of skeletons after computational flattening. Cell body locations are schematically indicated with shaded circles. Dashed lines delineate the marginal-central boundaries or SAC depths. Only one example of each cluster is shown here. All cells can be viewed in Supplementary Data 2 and the EyeWire Museum.
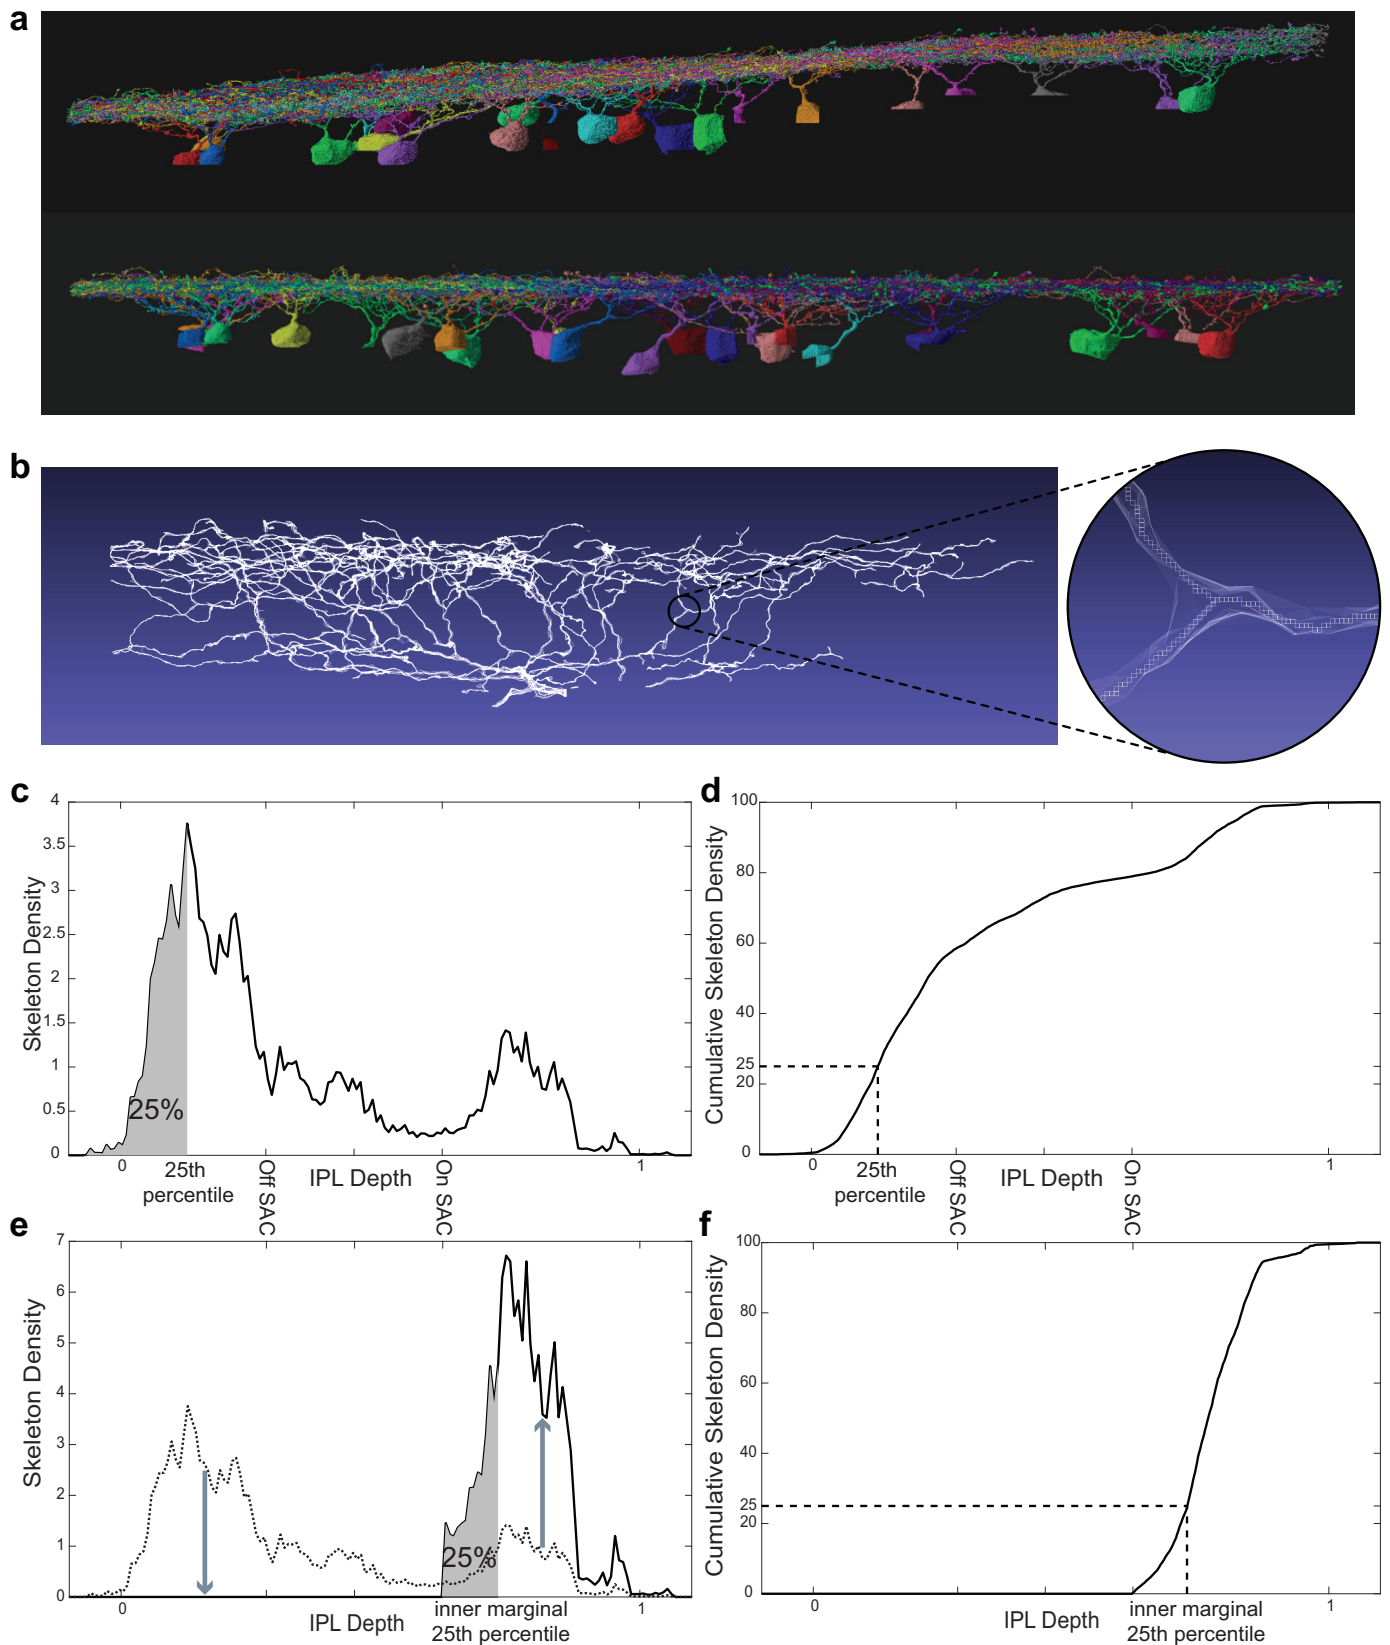
Extended Data Figure 4: Computation of stratification profiles. **a,** Computational flattening of the On SAC layer. **b,** Transformation of 3D reconstruction into skeleton. **c,** The hierarchical clustering procedure relies on various percentiles of the stratification profile. For example, the 25th percentile is defined as that IPL depth below which the stratification profile contains 25% of its area (shaded). **d,** Equivalently, the cumulative distribution obtained from the stratification profile takes on the value 25% at 25th percentile IPL depth. **e-f,** Density and cumulative distribution for the same stratification profile restricted to the inner marginal sublamina (Methods)
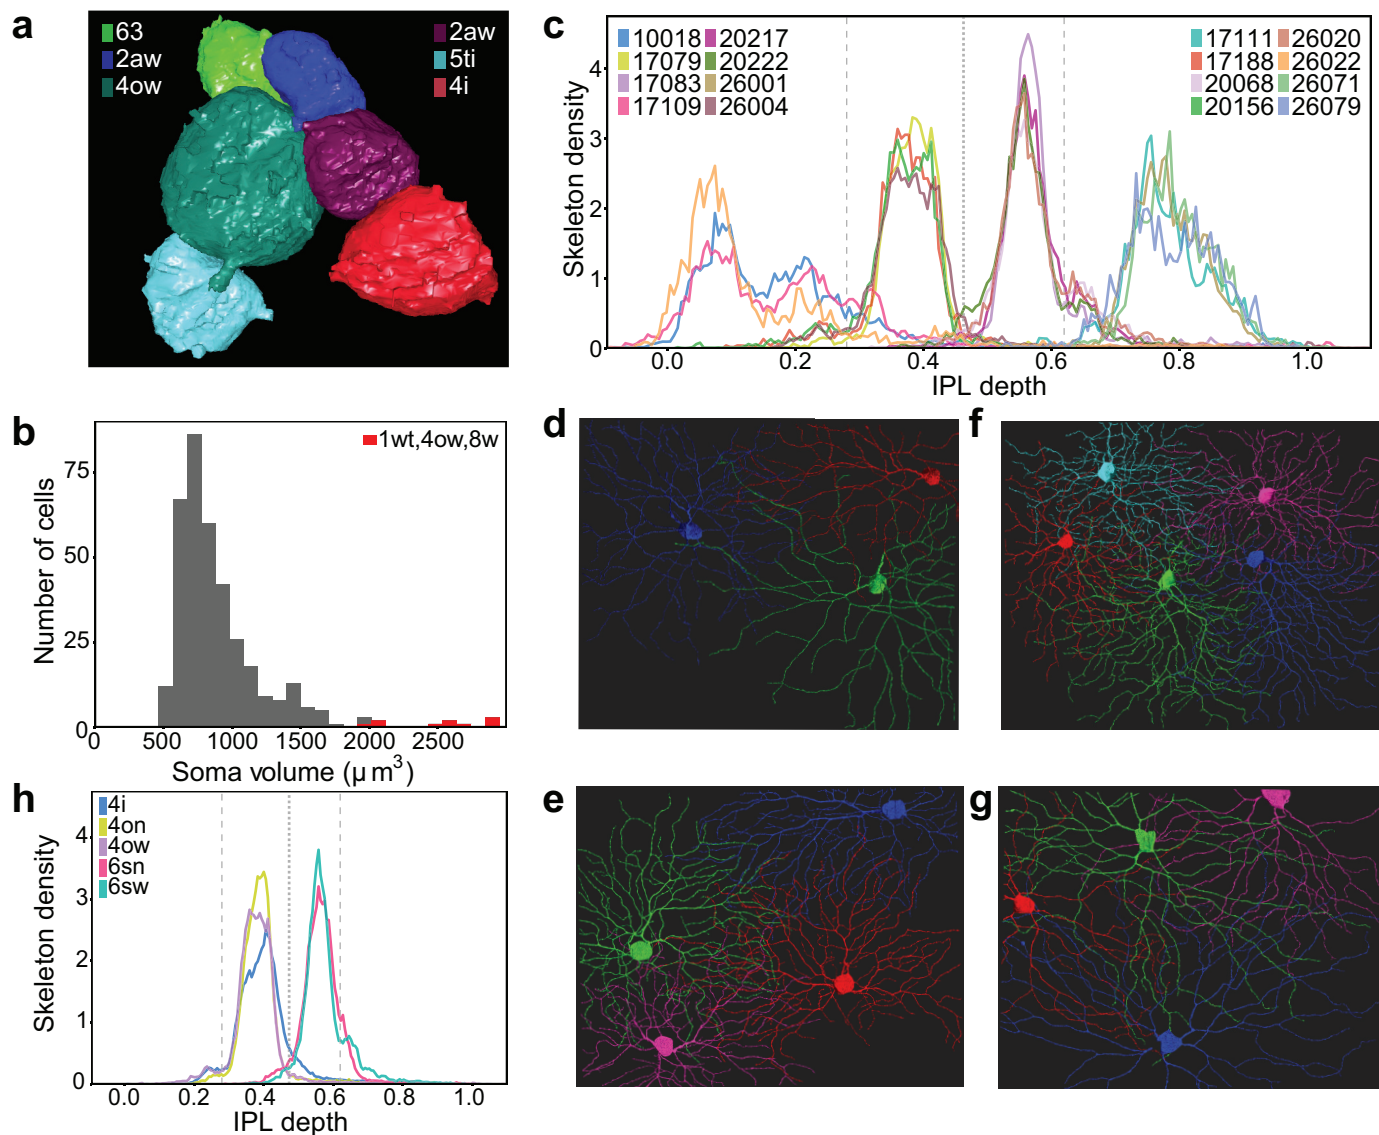
Extended Data Figure 5: Soma size distinguishes certain cell types. **a,** Reconstructed soma of example alpha cell (dark green) is much larger than surrounding somas. **b,** Classical alpha cells (1wt, 4ow, 8w) are outliers in histogram of ganglion cell soma sizes. **c,** Stratification profiles for individual alpha cells (classical and 6sw) show high reproducibility within each type. Five digit numbers are identifiers that can be used to locate cells in the EyeWire Museum. **d-g,** Cells of each alpha type tile the retinal patch with little overlap (**d**, 1wt; **e**, 4ow; **f**, 6sw; **g**, 8w). Some gaps in coverage are visible, and presumably due to cells with somata outside the e2198 volume. **h,** Candidates for "mini" alpha types (4i, 4on, 6sn) are almost identical to classical alpha types (4ow, 6sw) in average stratification profiles.
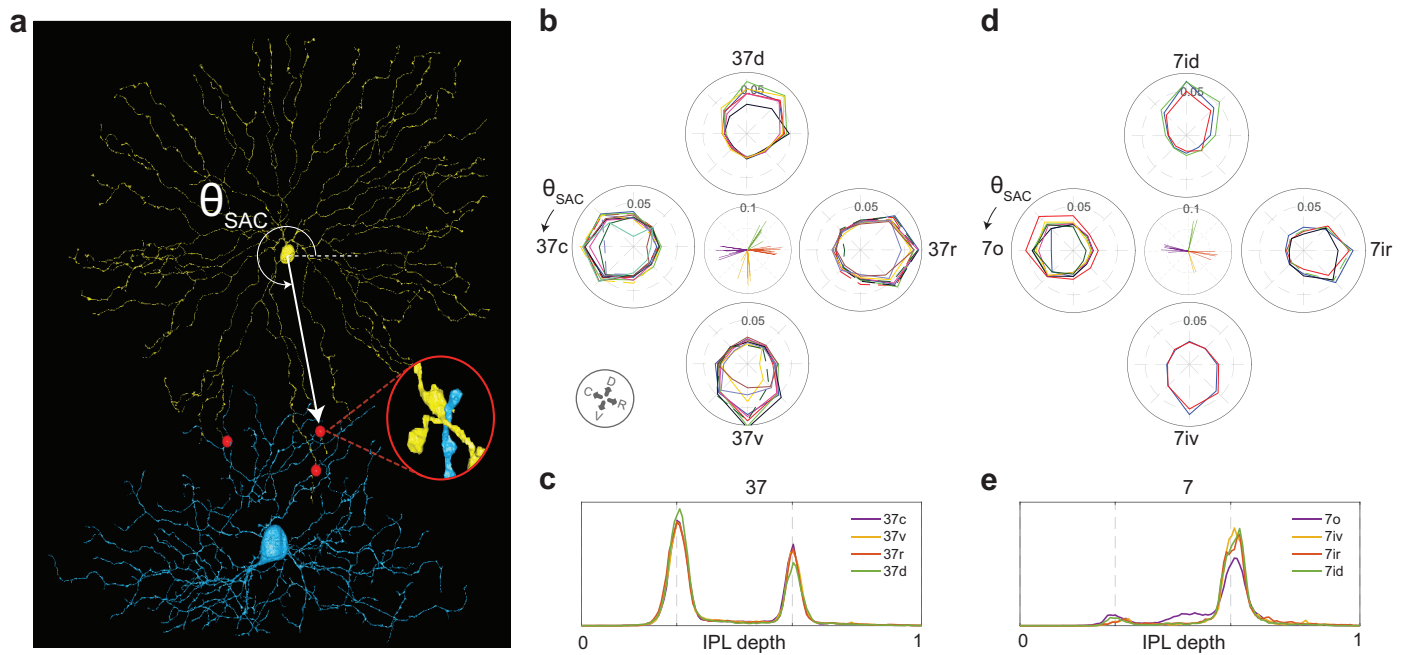
Extended Data Figure 6: On-Off and On DS cells separate into types by preferred directions of SAC contact. **a,** For each SAC-GC contact (red dots, inset), SAC dendrite direction θ_{SAC} is defined by a vector from SAC soma to the contact. **b,** For each On-Off DS cell in our sample, the fraction of intermingling SAC dendrite in contact with the cell is graphed versus θ_{SAC}. The polar plots are highly reproducible within each type, as are the vector sums of the polar plots (center). **c,** The four On-Off DS types (37v, 37r, 37c, 37d) differ little in their average stratification profiles. **d,** For each On DS cell, the fraction of intermingling SAC dendrite in contact with the cell is graphed versus θ_{SAC}. The polar plots are highly reproducible within each type, as are the vector sums of the polar plots (center). **e,** The three classical On DS types (7id, 7ir, 7iv) have similar average stratification profiles, which differ slightly from that of the transient On DS type (7o). The coordinate systems in **b** and **d** are the same as in Fig. 1a, as indicated by the diagram showing our estimates of the four cardinal directions.
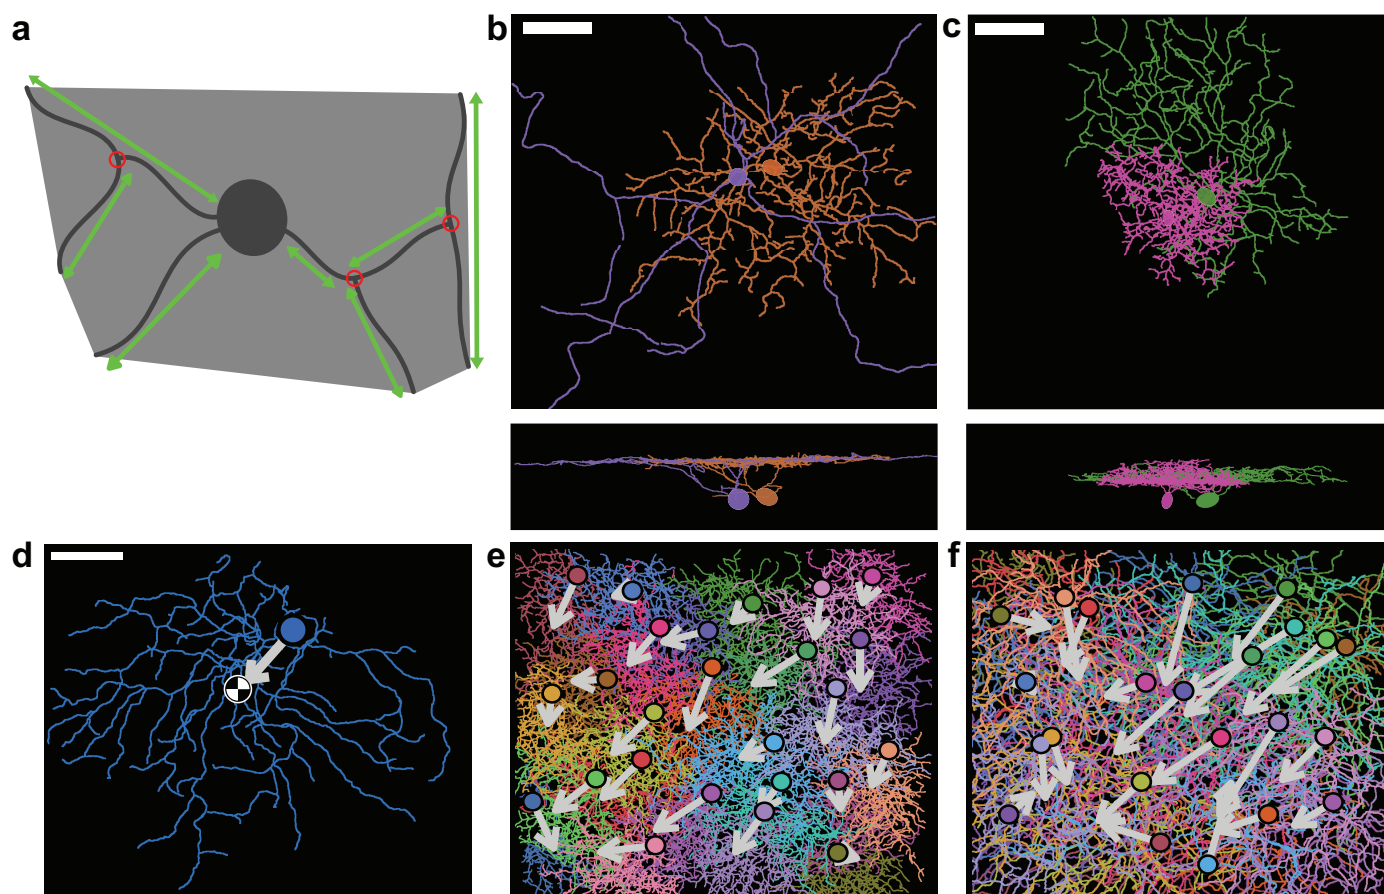
Extended Data Figure 7: Arbor density, complexity, and asymmetry. **a,** From the skeletonized arbor, we extract total path length (sum of green lengths), branch points (red), and convex hull area (shaded). **b,** Of all types, 1ws (purple) has the lowest arbor density, defined as ratio of total path length to convex hull area. 1ni (red) is shown for comparison. **c,** Of all types, 5ti (purple) has the highest arbor complexity, defined as ratio of branch point number to total path length. 5to (green) is shown for comparison. **d,** The arbor vector of a cell is drawn from the soma centroid to the skeleton centroid. **e-f,** For an asymmetric type, the directions of arbor vectors are nonuniformly distributed. **e**, 2an; **f**, 2aw.
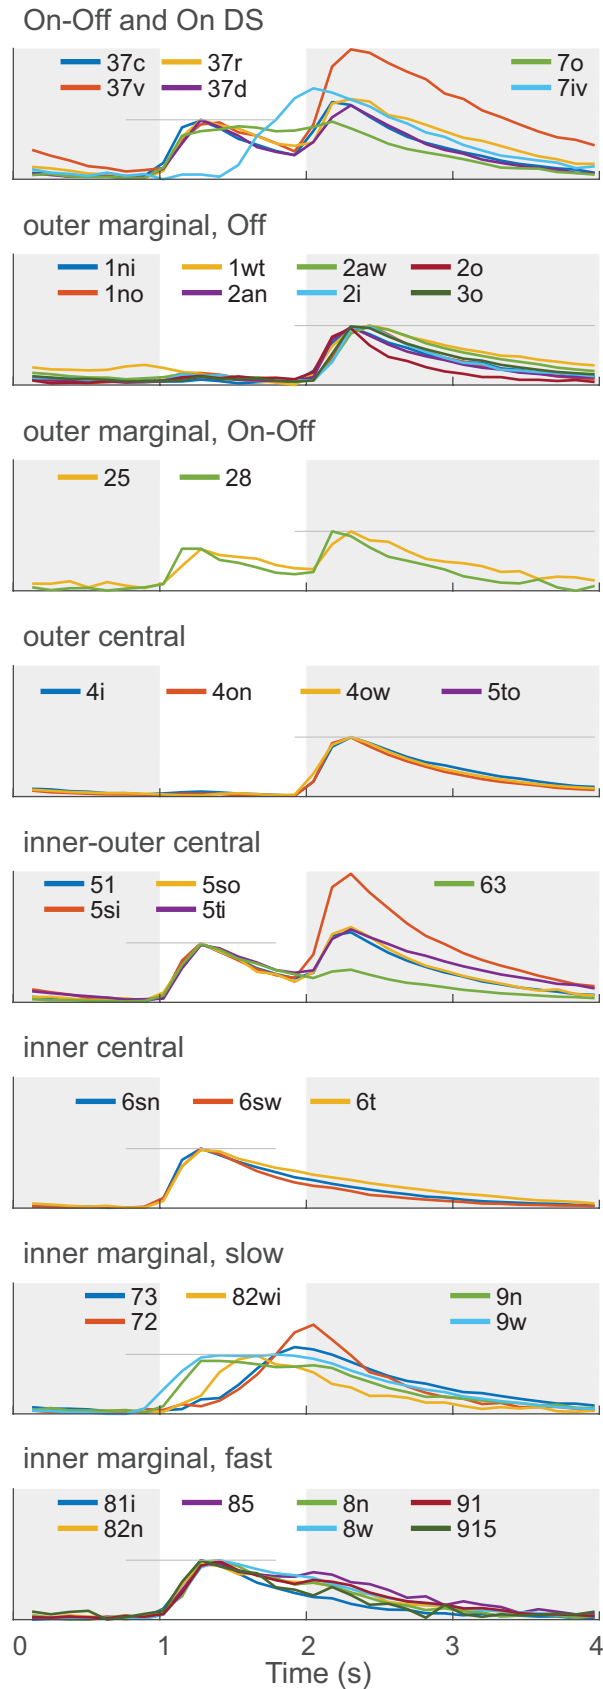
Extended Data Figure 8: Average temporal response functions for all clusters containing cells with high response quality (Methods, Eq. 4). Each graph shows the temporal response functions for the low-level clusters contained in a high-level cluster. For easier visualization, outer marginal clusters have been separated into Off and On-Off groups, and inner marginal clusters have been separated into "slow" and "fast" groups.
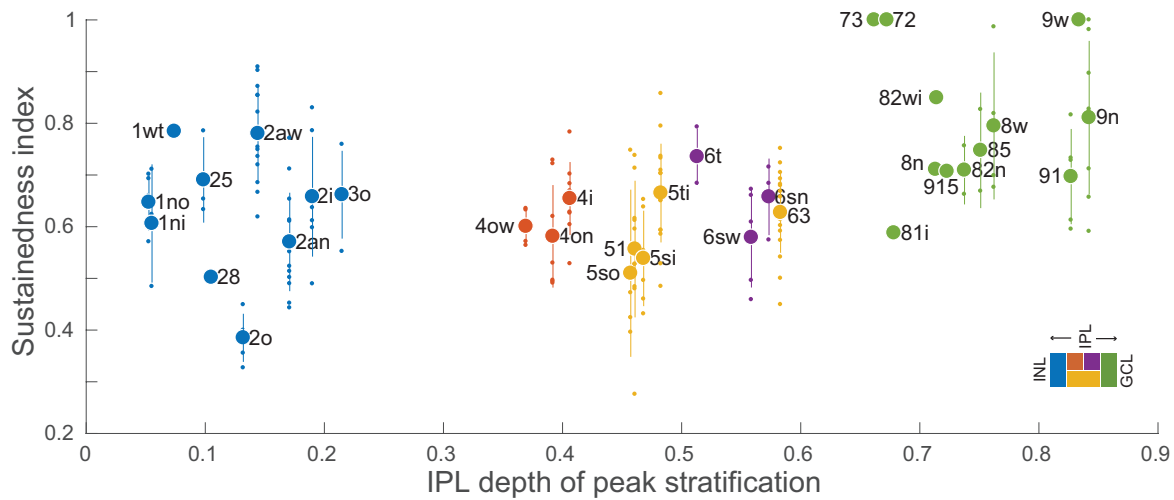
Extended Data Figure 9: Size-adjusted sustainedness index (Methods) for cells with high response quality, similar to Fig. 6e. Larger dots are means of each cluster, and individual cells in each cluster are plotted as small dots with the same horizontal coordate. Bars indicate standard deviations for the clusters. Clusters with no or invisible bars either contain only a single cell with high response quality (1wt, 28, 81i, 8n, 915, 9w), or has two (82wi), three (72), or four (73) high response quality cells with nearly identical sustainedness index.