The optimal spatial arrangement of ON and OFF receptive fields

Na Young Jun[†], Greg Field[†], and John Pearson[†‡]
[†]*Department of Neurobiology* [‡]*Department of Biostatistics & Bioinformatics*
Duke University

Many sensory systems utilize parallel ON and OFF pathways that signal stimulus increments and decrements, respectively. These pathways consist of ensembles or grids of ON and OFF detectors spanning sensory space. Yet encoding by opponent pathways raises a question: How should grids of ON and OFF detectors be arranged to optimally encode natural stimuli? We investigated this question using a model of the retina guided by efficient coding theory. Specifically, we optimized spatial receptive fields and contrast response functions to encode natural images given noise and constrained firing rates. We find that the optimal arrangement of ON and OFF receptive fields exhibits a second-order phase transition between aligned and anti-aligned grids. The preferred phase depends on detector noise and the statistical structure of the natural stimuli. These results reveal that noise and stimulus statistics produce qualitative shifts in neural coding strategies and provide novel theoretical predictions for the configuration of opponent pathways in the nervous system.

INTRODUCTION

Across many sensory systems, neurons encode information about either increments or decrements of stimuli in the environment: so-called 'ON' and 'OFF' signals. This division between ON and OFF signaling has been observed in visual (Joesch et al. 2010, Kuffler 1953), thermosensory (Gallio et al. 2011), auditory (Smith and Lewicki 2006), olfactory (Chalasani et al. 2007), and electro-sensory (Bell 1990) systems. This organization has the advantage that neurons can be tasked with signaling changes in the environment rather than the steady-state or mean stimulus condition, thereby resulting in more efficient neural codes (Barlow 1961, Gjorgjieva et al. 2014). Moreover, when the number of potential stimuli is large, neurons often specialize: for example, they only respond to a small region of visual space or a narrow auditory frequency band. The combination of these coding strategies raises two questions: First, how should specialized detectors, the ON and OFF cells, arrange themselves most efficiently cover stimulus space? Second, what is the optimal relative arrangement of ON and OFF detector grids? For one system, the retina, the answer to the first question is clear from previous work: detectors tile stimulus space and exhibit overlap near the 1-sigma boundary of a Gaussian receptive field (Borghuis et al. 2008, Devries and Baylor 1997, Doi et al. 2012, Gauthier et al. 2009a,b). The answer to the second question, what might be called the "sensor alignment problem," has received comparatively little attention and is the focus of this study.

Conceptually, there are three general possibilities for how the sensor alignment problem could be solved: One possibility is that the grids of sensors are statistically independent, meaning the locations of receptive fields in one grid provide no information about the receptive field locations in the other grid. A second possibility is that the two grids are aligned, meaning the receptive field centers in one grid are closer than expected by chance. The third possibility is that the two grids are anti-aligned, meaning the receptive field centers in the two grids are further apart than expected by chance. On general information theory grounds, the optimal solution is likely to depend on noise in the encoding process and the statistics of the encoded stimuli (Atick and Redlich 1990, 1992).

To determine the optimal solution to the sensor alignment problem, we rooted our analysis in efficient coding theory (Barlow 1961). This theory argues that sensory systems should aim to reduce the redundancy present in sensory input while minimizing metabolic costs, thereby, reliably encoding natural stimuli with fewer spikes. Efficient coding theory has been successful at explaining many aspects of sensory processing and retinal physiology, including center-surround receptive fields, grids of receptive fields (henceforth called 'mosaics'), and a greater proportion of OFF than ON cells (Atick and Redlich 1992, Barlow 1961, Doi et al. 2012, Karklin and Simoncelli 2011, Ratliff et al. 2009). Thus, we asked whether efficient coding theory might predict the optimal spatial arrangement of ON and OFF receptive field mosaics within the retina. Our approach to this question involved optimizing a model that approximates the processing performed by many retinal ganglion cells (Karklin and Simoncelli 2011). By maximizing the mutual information between an (input) library of natural images and (output) spike rates, we could study the effects of image statistics and encoding noise on the optimal arrangement of ON and OFF mosaics.

In this model, we found that the optimal spatial arrangement was a pair of approximately hexagonal mosaics of ON and OFF receptive fields. However, surprisingly, the relative alignment of these mosaics depended on the input noise, output noise, and the statistics of the natural image set. When output noise was low, the mosaics were aligned, with ON and OFF receptive fields centered at nearby locations more often than expected by chance. When output noise was relatively high, anti-alignment became the

favored arrangement. Surprisingly, the content of the image set also strongly influenced the transition between aligned and anti-aligned mosaics. In particular, when image sets contained more 'outlier' images with particularly large luminance or contrast values, anti-alignment became the favored state for fixed input and output noise. As we demonstrate analytically and confirm computationally, the boundary between these two regimes is characterized by a second-order phase transition: as noise parameters or stimulus statistics vary, mutual information changes smoothly, while the optimal mosaic arrangements undergo a sudden, qualitative shift. Finally, we confirm these predictions by showing that systematic manipulations of the training data set change the phase boundary in a manner predicted by an analytical model. These findings underscore the crucial role played by both noise and the statistics of natural stimuli for understanding specialization and coordination in sensory processing.

METHODS

Linear-nonlinear efficient coding model

The model used to examine efficient coding is based on previous work (Karklin and Simoncelli 2011), designed to find the optimal encoding of visual scenes. The model takes patches of natural scenes (Doi and Lewicki 2007) as the input, each of which has 18x18 (324) pixels. Each input image was multiplied by a circular mask to avoid cases in which corners in the optimization region produce corner artifacts that dominate the final configuration. The circular masking effectively left about 255 pixels in each input image patch. The model consisted of linear-nonlinear (LN) architecture with learned linear filters and nonlinearities. Given an input stimulus $\mathbf{x}$, an input noise $\mathbf{n}_x$, and the linear filter $\mathbf{w}_j$ for each neuron j, the nonlinearities were modeled as softplus functions of the linearly filtered noisy stimulus for each neuron $y_j = \mathbf{w}_j^\top(\mathbf{x} + \mathbf{n}_x)$:

$$h(y_j) = \log(1 + \exp(\beta \cdot y_j))/\beta \tag{1}$$

using a fixed value of $\beta = 2.5$. We empirically chose a fixed value of $\beta = 2.5$ which produced a stable optimization trajectory. The output of the model consisted of 100 units that mimic retinal ganglion cells (RGCs), whose firing rates are modeled using neuron-specific gain and threshold parameters γ_j and θ_j, plus an additive output noise $n_{r,j}$:

$$r_j(y_j) = \gamma_j \cdot h(y_j - \theta_j) + n_{r,j}. \tag{2}$$

Note that scaling r, γ, and n_r by the same factor leaves the model invariant, so we choose units in which the mean firing rate is $\mathbb{E}r = 1$. We used a first-order approximation that treats the conditional distribution of activation as a Gaussian, which allowed us to derive a closed-form expression for the mutual information. A detailed derivation is provided in the supplement. The model was trained to maximize the mutual information between the input images and the firing rates of the output neurons. To model the metabolic cost of spiking, the average firing rate of each neuron across images was constrained to 1, which was enforced via an augmented Lagrangian method with quadratic penalty. Input and output noise were modeled as i.i.d. Gaussian noise (zero means, and standard deviations of σ_{in} and σ_{out} respectively).

The parameters of the model were the receptive field filter weights,the gain of the nonlinearity, and the nonlinearity threshold of each output RGC. To preserve positive gain for each output unit, we optimized the logarithms of the corresponding parameters. Filter weights were initialized by drawing samples from a Gaussian distribution (mean zero, standard deviation 1; see **Fig 1B**, top row). The parameters of the nonlinearities (log gain and threshold) were initialized by sampling from a uniform distribution [0, 1]. The parameters were then optimized using stochastic gradient descent with a learning rate of 0.001. During training, we calculated gradients using mini-batches of 100 image patches at a time, renormalizing the filters to have unit norm after each gradient step. Models were trained for 1,000,000 iterations, which we found sufficient in practice to ensure convergence of the mutual information. As a means of escaping local optima, between 200,000 and 500,000 iterations, two modifications were applied to the optimization procedure, which we call jittering and centering. These were used to speed convergence: The jittering operation was applied every 5,000 iterations and consisted of raising each element of the kernels to the power 1.25; this makes high-amplitude portions of the filters more pronounced while attenuating low-amplitude portions. The centering operation penalizes the spatial spread of the kernel by adding the mean of the spatial variance of each kernel as an additional loss term; this encourages the kernels to be localized around their centers. These two operations allowed kernels that were trying to 'squeeze' into an already formed mosaic the space to do so, effectively speeding the optimization and allowing all kernels to contribute to the encoding. After these operations, the model was optimized without the modifications for another 500,000 iterations. We confirmed that jittering and centering do not alter the final converged shape of the kernels.

Median nearest neighbor analysis of the spatial relationships between the mosaics

To quantify the relative alignment of ON and OFF mosaics, we used a metric based on the median nearest-neighbor distances between heterotypic receptive field centers. For each RF center, we found the nearest RF center with the opposite polarity to define its heterotypic nearest neighbor distance, and calculated the median of these values(See Equation 3). When ON and OFF mosaics are nearly aligned, this value is close to zero as in **Fig 2A-B**.

$$\begin{array}{c}\text{Median}\\ \text{NN distance}\end{array} = \underset{c\in\text{RF centers}}{\text{median}}\left(\min_{\substack{\text{RF centers}\\ c'\in\text{with opposite}\\ \text{polarity to c}}} \|c-c'\|\right) \quad (3)$$

One-shape model for rapid exploration of mosaic arrangements

The full 18x18 system required many hours to optimize, motivating the development of a smaller, scaled-down system for exploring how mosaic arrangements depended on different noise parameters (**Fig 3**) or different stimulus statistics (**Fig 6**). Here, we restricted input images to 7x7, fixed the number of ON and OFF cells at 7 each (14 total), and adopted a fixed parametric form for the receptive field. Thus, the optimization only learned receptive field center locations and a small number of receptive field shape parameters. The receptive field function was parameterized as a difference of two Gaussian curves, with a third parameter determining their relative magnitude.

$$\kappa(r) \propto e^{-ar^2} - ce^{-br^2}, \qquad a,b,c>0 \quad (4)$$

Optimization in this setup involves fewer parameters, as well as allowing us to use more efficient first-order optimization methods like Adam (Kingma and Ba 2014). To examine the dependence of mosaic arrangement on noise a grid search was performed over combinations of input and output noise values, with input noise (σ_{in}) ranging from 0.0 to 0.75 and output noise (σ_{out}) from 0.75 to 3.0. As previously argued (Karklin and Simoncelli 2011), these values like within the physiological range for naturalistic stimuli. The optimization was run three times for every pair of noise values and the model resulting in the highest mutual information was retained **Fig 3**. This model was also used to examine the dependence of mosaic alignment on the statistics of natural scenes **Fig 6**. In these simulations the input noise was held fixed while the output noise and distribution of natural images was varied to examine the interaction between these factors on the resulting mosaic arrangements.

RESULTS

Linear-nonlinear efficient coding model of the retina

We asked how mosaics of ON and OFF cells should be arranged to efficiently encode natural scenes. To answer this question, we trained a model of retinal processing to maximize the mutual information between a library of natural images and a set of output firing rates (**Fig 1A**). The model consists of a single layer of linear-nonlinear units intended to approximate the encoding performed by retinal ganglion cells (Baccus and Meister 2002, Borghuis et al. 2008, Chichilnisky 2001, Doi et al. 2012, Karklin and Simoncelli 2011, Meister and Berry 1999). These model cells linearly filter the input stimulus weighted by a spatial mask or 'receptive field' and sum the resulting values. This number, loosely analogous to a membrane potential, is then fed through a nonlinearity to approximate rectified neural firing rates. The model also included two sources of noise: (1) a pixel-wise input noise, which models stochasticity in phototransduction; and (2) an output noise, which models stochasticity in output firing rates. A key difference between the two is that the input noise is subject to the nonlinearity, and thus its effect becomes stimulus-dependent, while the output noise is not. We trained the model by optimizing mutual information under the constraint of a 1 activation (over all images) for each neuron (see Methods). This constraint is intended to model the metabolic cost of generating action potentials.

Spatial tiling of receptive fields is predicted by efficient coding

Following training, units in the model exhibited either ON-center or OFF-center receptive fields with strongly rectified nonlinearities (**Fig 1B-C**). Each population of ON-center units and OFF-center units (referred henceforth as ON cells and OFF cells, respectively) exhibited a mosaic-like organization: their receptive fields tiled space (**Fig 1D**) (Karklin and Simoncelli 2011). When varying noise levels, we reliably found circular ON and OFF receptive fields with input noise between 0.00 and 0.75 and output noise between 0.75 and 3.00. Input and output noise levels outside these ranges produced irregular receptive field shapes (**Sup Fig 1**). Therefore, we restricted the analysis to this range of noise regimes.

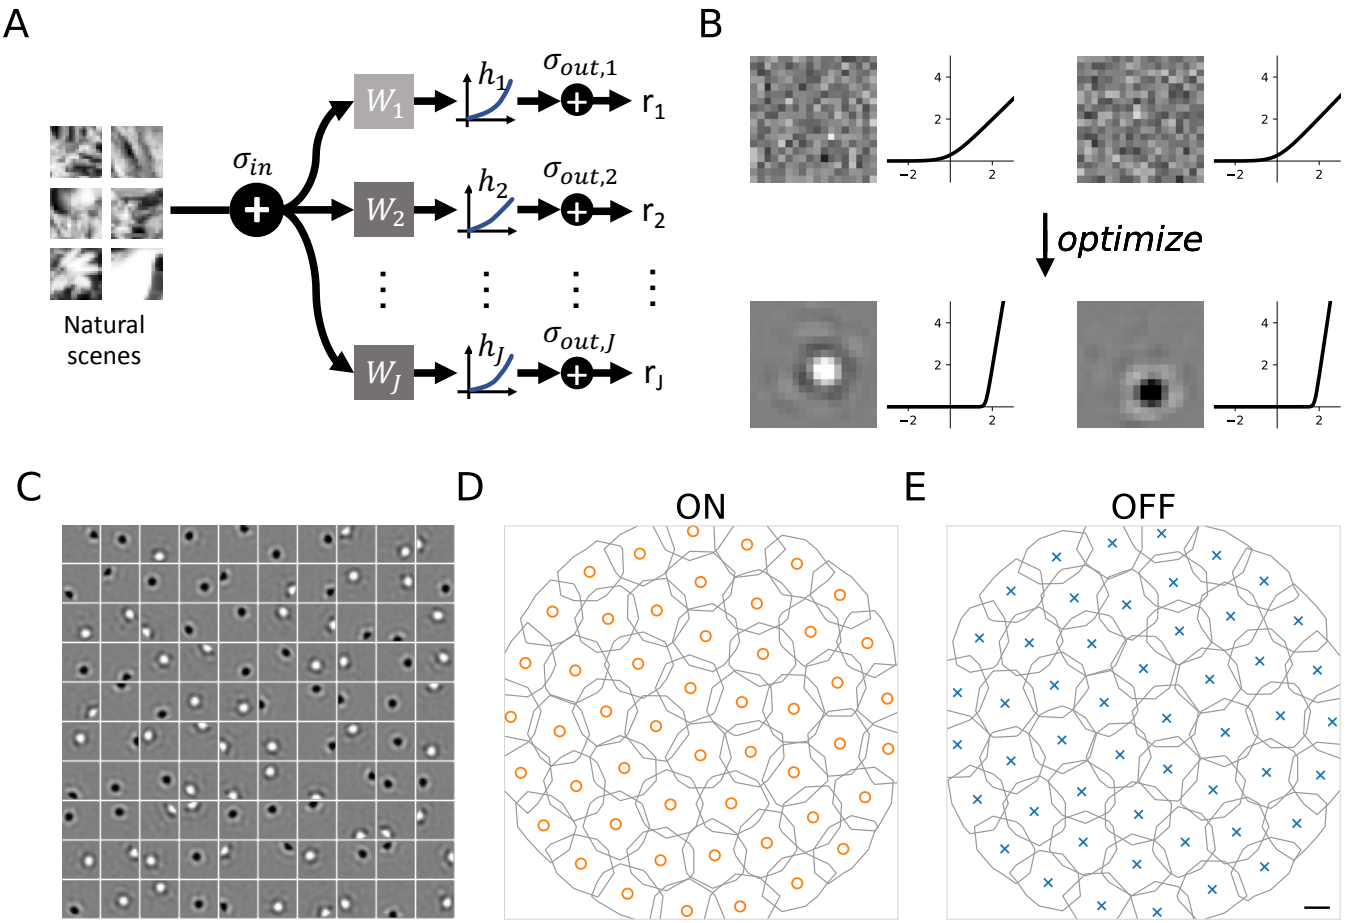

FIG. 1. **Spatial tiling of RFs is predicted by efficient coding.**
A. Efficient coding model architecture. Images of natural scenes (plus input noise σ_{in}) are multiplied by linear filters $\mathbf{w}_j$, passed through a nonlinear function h_j, and perturbed by output noise σ_{out}, resulting in a firing rate r_j for neuron j. **B.** Examples of initial and optimized filters and nonlinearities, where the filters are initialized from a white noise distribution and converge to ON and OFF kernels with a center-surround pattern. The nonlinearities are initialized as unscaled (but slightly perturbed) softplus functions and converge to nonzero threshold values. **C.** A plot of 100 kernels from a trained model with $\sigma_{\text{in}} = 0.4$, $\sigma_{\text{out}} = 3.0$. **D.**A contour plot showing the tiling of the ON and OFF kernels. The contours of the two types of kernels are drawn where the normalized pixel intensity is $\pm$ 0.21. Orange circles and blue x's indicate RF centers of mass for ON and OFF cells, respectively. The scale bar shows a 1-pixel distance.

Spatial arrangements of receptive field mosaics depend on input and output noise level

We next analyzed how these optimized mosaics of ON and OFF receptive fields were spatially arranged with respect to one another. Each optimized mosaic formed a nearly hexagonal grid (**Fig 1D**). Analysis of the spatial relationships between the mosaics in **Fig 1D-E** revealed they were anti-aligned (**Fig 2D**). However, as we explored the effect of changing the amounts of input and output noise on the optimization of this model, we observed other mosaic combinations that were aligned (e.g. **Fig 2D**). Generally, small pairs (input and output) of noise values yielded aligned mosaics, while larger paired noise values produced anti-aligned mosaics (**Fig 2A-D**). Importantly, for a fixed amount of input and output noise, multiple instances of the optimization produced consistent results, indicating that mosaic alignment and anti-alignment depended on the specific amount of input and output noise, and did not depend on different initializations to the optimization. These results motivated a more thorough analysis of how the relative spatial organization between the two mosaics depended on input and output noise.

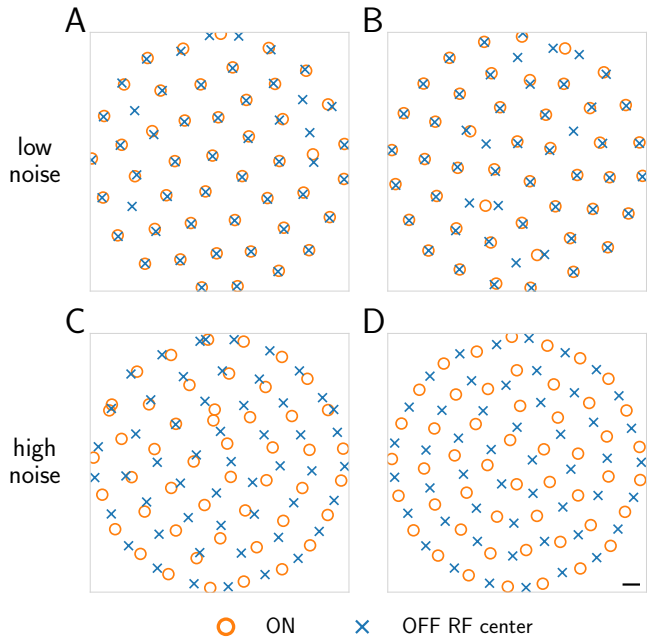

FIG. 2. **Spatial coordination of receptive field mosaics depends on input and output noise level.** Receptive field centers for ON (orange circle) and OFF (blue x) cells under differing sets of noise parameters. **A:** (σ_{in}=0.02, σ_{out}=1.0), **B:** (σ_{in}=0.1, σ_{out}=1.0), **C:** (σ_{in}=0.4, σ_{out}=2.0), **D:** (σ_{in}=0.4, σ_{out}=3.0). The first two parameter sets result in aligned mosaics, while the latter two, at higher levels of output noise, are anti-aligned. The scale bar shows a 1-pixel distance.

A simplified 'one-shape' model captures the noise dependence of mosaic arrangements

To examine the dependence of the mosaic arrangements on noise over a range of densely sampled values, we utilized a simpler 'one-shape' model that trained much more quickly than the full model (see Methods). This model assumed that all receptive fields shared a common shape (a difference of concentric Gaussians), while optimizing over both the parameters of this shape and the locations of the receptive fields (see Methods). The learned kernels of this simplified model also exhibited ON and OFF RFs (**Sup Fig 2A**) and ON and OFF RF mosaics (**Sup Fig 2B**). In addition, the optimal radial kernel exhibited a center-surround organization (**Sup Fig 2C**). To further speed training, we also reduced the number of units in the model to 7 ON and 7 OFF cells (their polarities were fixed during the optimization). These simplifications allowed us to rapidly judge if the optimization preferred alignment or anti-alignment (**Fig 3**). To ensure these simplifications did not strongly bias the results, we compared the mosaic arrangements for particular input and output noise combinations across the one-shape model and the full model. In all cases examined, the two models produced matching aligned or anti-aligned results (**Sup Fig 3**). Thus, the one-shape model is a useful proxy for larger-scale and more general optimizations and could be used to reliably determine if mosaics were aligned or anti-aligned for a given set of noise parameters.

Mosaics exhibit a phase change as a function of input and output noise

To determine the optimal arrangement of ON and OFF mosaics for different amounts of input and output noise, we performed a grid search over input and output noise values, using the one-shape model. We ran the model with input noise ranging from 0.00 to 0.75 in steps of 0.05, and output noise from 0.75 to 3.00 in steps of 0.10. Optimizing the location of seven RF centers in a circular space nearly always resulted in a single kernel in the center surrounded by the remaining six kernels at the vertices of a hexagon. (see **Fig 3**). When output noise level was low, ON and OFF RF mosaics were aligned as in the first row of **Fig 3B** and resulted in near-zero median nearest neighbor distances (dark purple colored area in **Fig 3A**), whereas ON and OFF RF mosaics were anti-aligned under higher output noise levels as in the last row of **Fig 3B** and resulted in larger median nearest neighbor distances (lighter colored area in **Fig 3A**). In particular, this analysis reveals an abrupt transition between mosaic alignment and anti-alignment (**Fig 3A**). Thus, our model optimization indicates the solution to the sensor alignment problem depends on noise present in the encoding process. As argued below, this is a second-order phase transition (Binney et al. 1992). Below, we build an analytical model for understanding this result.

Retinal phase transitions from an analytic model of efficient coding

To understand the factors that give rise to the phase transition between aligned and anti-aligned mosaics, we developed a simplified 1-dimensional analytic model in which ON and OFF grids have identical nonlinearities, fixed kernels, fixed spacing, and are only allowed to shift relative to one another (see supplementary section 2). Moreover, we assume that images are drawn from a Gaussian distribution with a $1/f$ frequency spectrum that is the 1D equivalent of the two-dimensional scale-free distribution of natural images (Atick and Redlich 1990, Ruderman 1994). We verified via simulation that, even with all of these restrictions, the phase transition occurs just as in the less constrained models in both one and two dimensions, suggesting that the remaining free parameters — gain, threshold, and

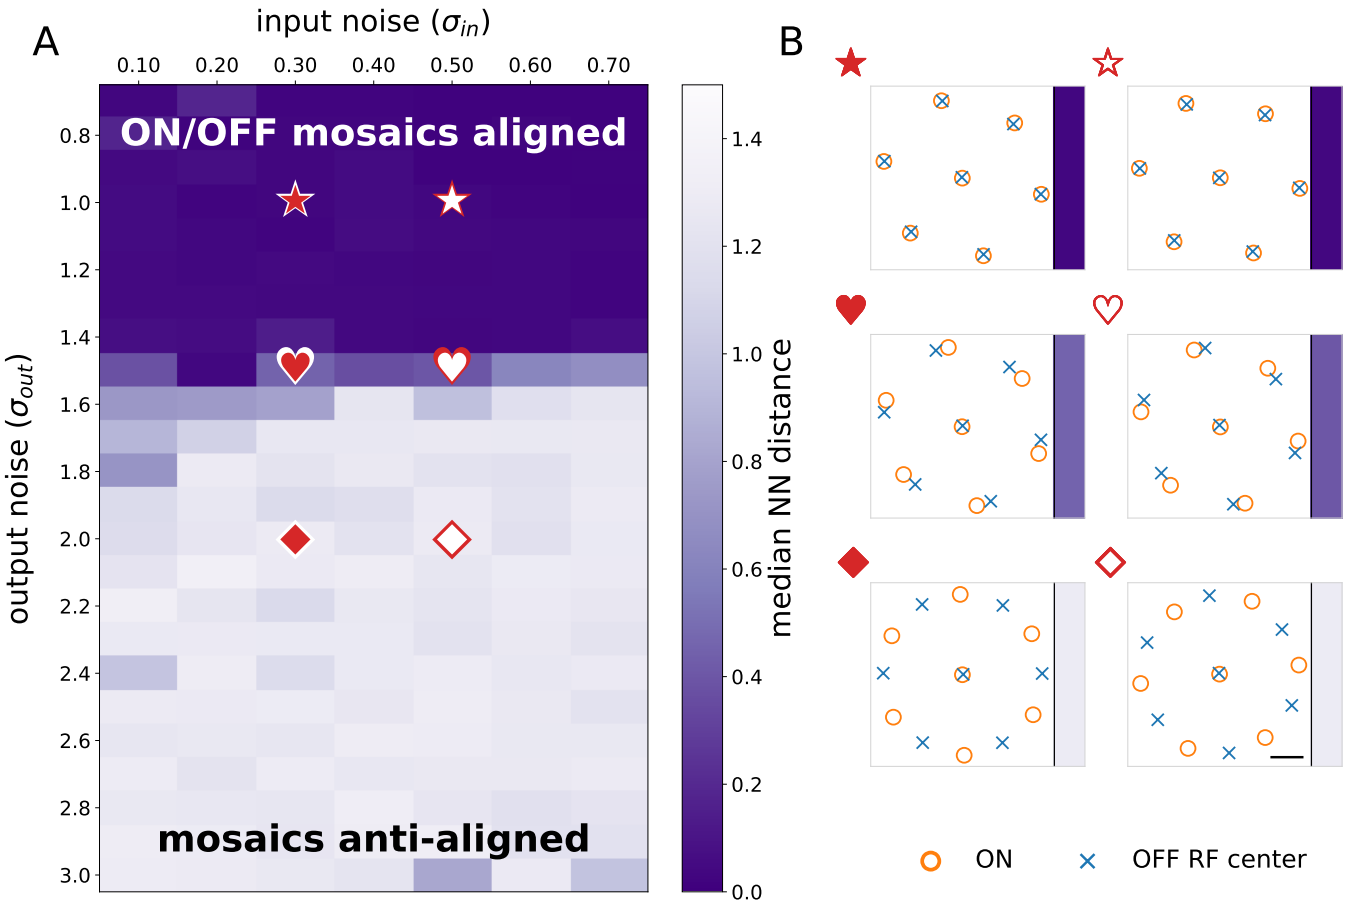

FIG. 3. **Optimal mosaic configurations transition from aligned to anti-aligned as a function of noise.** **A.** As output noise increases, optimal mosaic configurations shift from aligned to anti-aligned, and this holds over a range of input noise values. Color indicates median NN distance, with darker indicating more clearly aligned (see Equation 3 and Methods). The existence of a phase boundary between the two alignments is clear. **B.** Examples of optimal mosaic arrangements for representative input and output noise combinations. Symbols denote the corresponding locations in 3A. The scale bar shows a 1-pixel distance. Note the existence of optimal configurations between aligned and fully anti-aligned (hearts) for some parameter values.

mosaic alignment — are sufficient to account for the phase transition (see **Sup Fig 5** and **Sup Fig 6**)).

To gain additional insight into the origins of the transition, we analyzed the behavior of this model in the more analytically tractable limit of near-independent neurons. The starting point for this analysis is a recognition that the mutual information between stimulus and neural response can be conveniently written as a sum of single-neuron terms and higher-order corrections:

$$\mathcal{I}(\mathbf{X};\mathbf{R}) = H(\mathbf{X}) + 2N\mathcal{I}_1 + \mathcal{I}_2 + \mathcal{I}_{2'} + \cdots \qquad (5)$$

Here, $H(X)$ is the (constant) entropy of the stimulus, $2N$ is the number of ON + OFF cells, $\mathcal{I}_1$ is the contribution of each neuron individually, and $\mathcal{I}_2$ and $\mathcal{I}_{2'}$ are correction terms due to coactivation (and thus redundancy) of ON-ON/OFF-OFF (same polarity) and ON-OFF (opposite polarity) cells, respectively (see supplementary section 2B). Importantly, these corrections are always negative, and vanish for widely separated cells. For this analysis, we do not consider third-order and higher terms, since the terms above are sufficient to model the transition from alignment to anti-alignment. In the following two sections we show (1) that noise and image statistics set the optimal neuron response thresholds, and (2) that these thresholds control the amount of redundancy in the population responses, thereby dictating whether mosaic alignment or anti-alignment is the preferred state.

Output noise and image outliers drive increases in optimal neuron response thresholds

First, we considered optimizing only the $\mathcal{I}_1$ term above, which captures the contribution of each neuron individually to encoding stimulus information. In particular, we focused on the influence of output

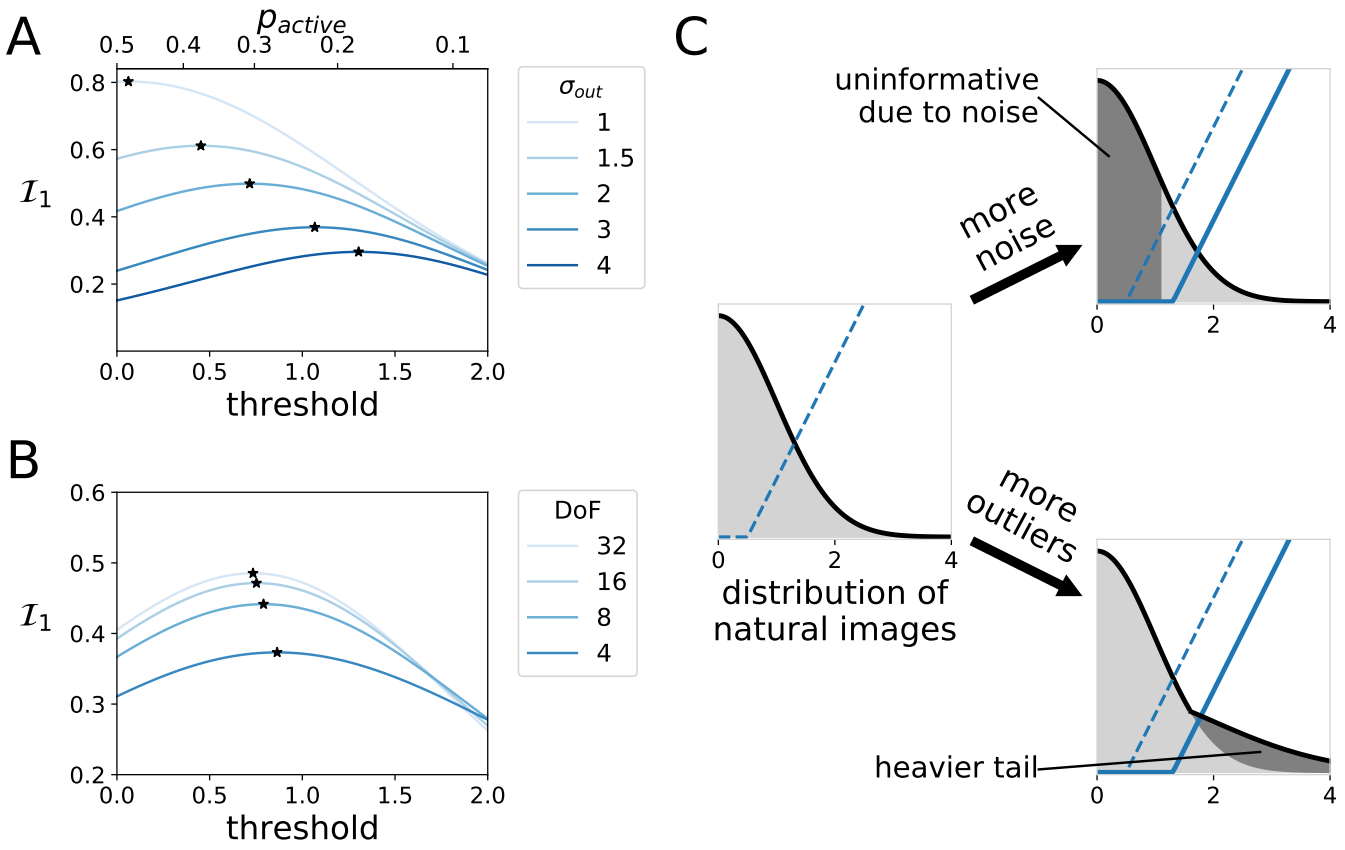

FIG. 4. **Optimal threshold increases with higher noise and heavier tails of image distribution.** **A-B.** A single neuron's mutual information contribution as a function of threshold. Stars mark optimal thresholds for multiple parameter values. **A.** When the input distribution is Gaussian, the optimal threshold increases with higher output noise, even as overall information decreases. **B.** When the input distribution has heavier tails, modeled with Student's t distributions with varying degrees of freedom, the optimal threshold again increases with heavier tails. **C.** An intuitive schematic using the distribution of natural images linearly projected onto activation levels of a neuron, to show the effects of increased noise or outliers. With increased noise (upper right), low-magnitude activations become uninformative, driving the threshold to increase. Similarly, a distribution with heavier tails (lower right) would require a higher threshold to maintain the same activation probability, since more mass is contained in outliers. That is, both output noise and heavy tails lead to higher degrees of optimal rectification.

noise and the distribution of natural images on the output nonlinearities of the neurons (we focused on output noise because the transition depended much more on output than input noise, see **Fig 3A**). As the output noise of neurons increases, maximizing $\mathcal{I}_1$ dictates that neurons should respond by increasing their firing threshold (**Fig 4A**). This allows the neurons to reserve spikes for stimuli with high signal-to-noise ratios (Laughlin 2001). Thus, units respond less often but do so with higher gain when active, thereby mitigating the impact of noise. In addition, the optimal firing threshold depends on the nature of the image distribution. For distributions with more outliers (see Methods), thresholds increase as the tail of the distribution becomes heavier (**Fig 4B**). Intuitively, this is because heavier-tailed distributions require higher thresholds to maintain the same overall response probability. Finally, it can be shown that decreasing the input noise has the effect of increasing the effective threshold (see supplementary section 2B, Equation 18). Thus, the optimal spiking threshold is determined both by the input and output noise and the distribution of natural scenes (schematized in **Fig 4C**) (Gjorgjieva et al. 2014).

Mosaic anti-alignment achieves redundancy reduction at high neuron thresholds

Next, we considered the effect of ON-OFF pair corrections to Equation 5. These terms, represented by $\mathcal{I}_{2'}$ above, capture the effects of redundancy in ON-OFF encoding: they are always negative and are the only contributions to mutual information that change as a function of alignment. Thus, while the overall trend

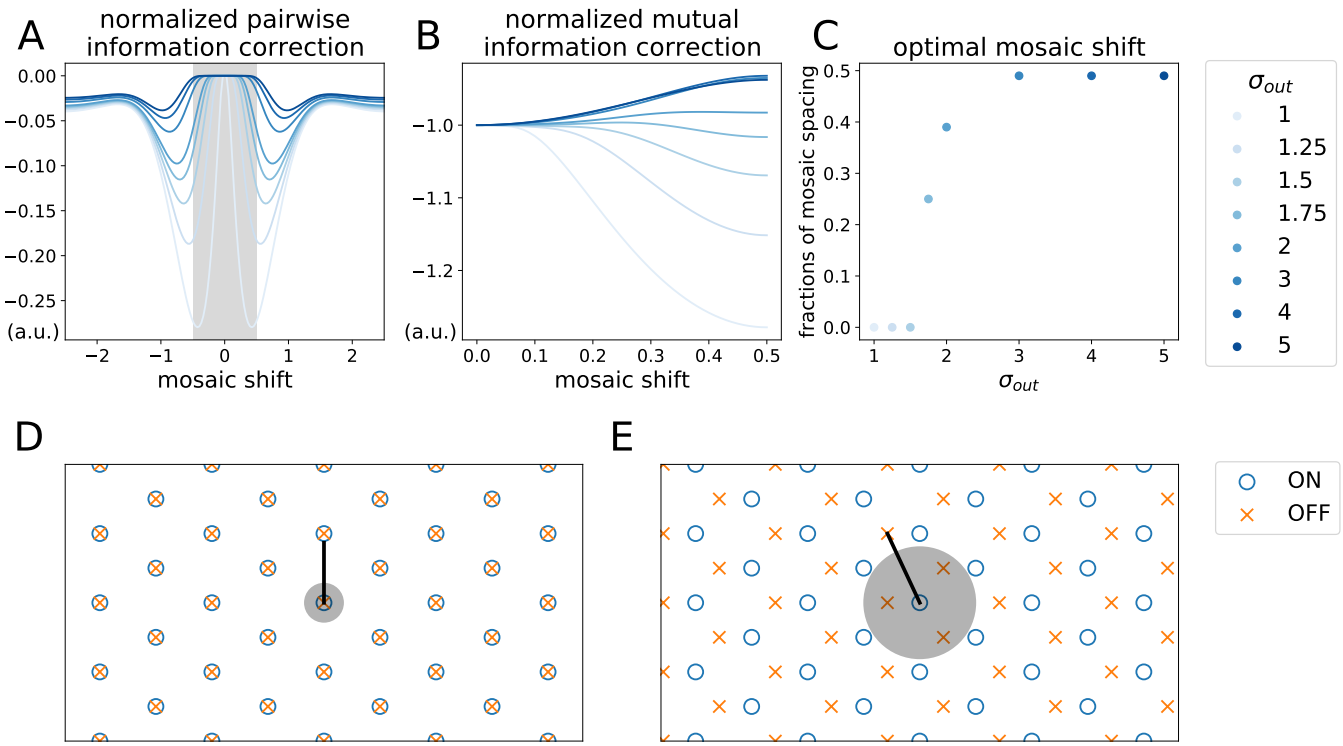

FIG. 5. **Pairwise mutual information corrections explain the phase transition between alignment and anti-alignment.**
A. Normalized pairwise information correction ($\mathcal{I}_{2'}/Np_1^2$) for pairs of ON and OFF neurons as a function of mosaic shift. One unit is the spacing between nearest-neighbor RF centers of the same polarity. With low output noise, an aligned position (i.e., no shift) is the only placement to avoid a large negative information correction, while with higher output noise, an ON-OFF pair can still have a finite shift within the gray area without incurring any penalty. **B.** The total normalized pairwise information correction as a function of mosaic shift, assuming equispaced ON and OFF RF centers filling the 1-D space. With low output noise, any mosaic shift lowers mutual information, while with high output noise, the pairwise mutual information becomes less negative when the mosaic is shifted by 0.5, i.e. when anti-aligned. **C.** The optimal mosaic shift as a function of output noise, showing the phase transition. For $\sigma_{\text{out}} = 0.1.7, \sigma_{\text{out}} = 2$, the optimal mosaic shift lies between 0 (alignment) and 0.5 (anti-alignment), resembling the arrangements we obtained in the middle row of **Fig 3B**. **D.** Optimal 2-D mosaics in a low-noise regime. The gray circle indicates the area within which an OFF cell can be located with zero pairwise correction to mutual information. **E.** Optimal 2-D mosaics in a high-noise regime. Here, the gray area has grown (as in **A**), allowing three neighbors of the opposite polarity to fit inside without loss in mutual information.

is that mutual information decreases with increasing threshold, mosaic configurations may be more or less efficient for information transmission as they limit the impact of these terms. More specifically,

$$\mathcal{I}_{2'} = \sum_{i,j} p_2(d_{ij}) \log R(d_{ij}) \tag{6}$$

Where d_{ij} is the distance between ON cell i and OFF cell j, p_2 is the probability that the pair is coactive, and $0 < R < 1$ is a function that is approximately independent of both output noise and neuron nonlinearity (see supplementary section 2B). This correction term is small and negative for large separations, large and negative around the intra-mosaic spacing, and almost zero in a region around zero inter-cell distance (**Fig 5A**). Importantly, the width of this 'almost zero zone' around zero inter-cell distance widens as output noise grows larger. This is explained by the fact that, as output noise increases, so does the threshold (**Fig 4A**), and at large threshold values, nearby opposite polarity pairs are almost never coactive (see supplementary section 2B).

This analysis gives rise to the following account of the phase transition: at low levels of output noise or stimulus outliers, neuronal thresholds are low, and only aligned RFs avoid the large $\mathcal{I}_{2'}$ penalty. Thus, in one dimension, each ON cell is non-redundant (save one bit of sign information; supplementary section 2C) with its aligned OFF partner, but image correlations make it highly redundant with the two OFF (and two ON) cells on either side. Likewise, in two dimensions, each ON cell is non-redundant (save one bit of sign information; see supplementary section 2C) with its aligned OFF partner, but it suffers large $\mathcal{I}_{2'}$ penalties to encoding efficiency from its six other OFF neighbors.

However, as activity thresholds rise, a finite-sized region develops around each cell within which other RFs of opposite polarity are almost never coactive (**Fig 5A**, gray). This region of non-redundancy grows with threshold, to the point that, when it is sufficiently large, more than one OFF cell can fit inside when the mosaics are anti-aligned (**Fig 5D-E, gray circle**). Thus, anti-alignment results in lower redundancy (via reduced $\mathcal{I}_{2'}$ penalties) and increased mutual information. This picture is born out quantitatively by the model: as output noise (and thus threshold) increases, the optimal shift interpolates smoothly between aligned and anti-aligned configurations (**Fig 5B-C** and cf. **Fig 3B**, second row), suggesting a second-order Landau-Ginzburg phase transition (Binney et al. 1992).

Outliers in natural image statistics modify the transition from aligned to anti-aligned mosaics

The analysis above identifies output noise and image (stimulus) outliers as driving optimal encoding strategies toward higher neural response thresholds, which we claim is the key factor mediating the phase transition from aligned to anti-aligned mosaics. However, this analysis makes numerous simplifying assumptions, including a restriction to one dimension. Thus, we sought to test whether these predictions generalized to a less-restricted two-dimensional model in which image statistics were manipulated.

First, we identified image patches that generated the highest and the lowest firing rates in our fitted model (**Sup Fig 4**). We identified these images using the firing rates under the two noise regimes, one that produced aligned and one that produced anti-aligned mosaics (**Fig 2B** ($\sigma_{\text{in}} = 0.1, \sigma_{\text{out}} = 1.0$) and **Fig 2D** ($\sigma_{\text{in}} = 0.4, \sigma_{\text{out}} = 3.0$)). In both noise regimes, the 100 image patches producing the highest firing rates were extreme luminance values (nearly all black or all white) or extreme contrast values (e.g. contained one or more edges between black and white regions). On the other hand, the 100 images producing the lowest firing rates were nearly homogeneous gray patches (**Sup Fig 4**). A 2-D histogram of the mean and standard deviation across the pixels of all of the available sample image patches confirmed that these images exhibited outlier mean or contrast values (**Sup Fig 4E**). Thus, under both noise regimes, the largest firing rates (across all neurons) were produced by rare, outlier images.

We next tested what impact these outlier images exert on the phase transition between aligned and anti-aligned mosaics. We re-ran the 'one-shape' model (e.g. **Fig 3A**), but with the altered image sets that either reduced or increased the frequency of outlier images. From the distribution of all of the 18,587,848 possible image patches in the dataset (**Fig 6A**), we considered a 2-D space composed of the z scores of the patch mean and SD values, and we drew the unit circle in this 2-D space. Within this space, we considered regions within 1 to 3 standard deviations from the mean of the average pixel intensity (per image) and from the mean of the variance over pixel intensities (per image). We then trained the one-shape model (see **Fig 3**) on only the patches within each boundary (**Fig 6B**). As predicted, we found that the transition between aligned and anti-aligned mosaics occurred at lower noise levels as more outlier images were included in the training set and occurred at higher noise levels when these outlier images were removed. Furthermore, these changes were mediated by an increase in activation threshold as outliers grew more prevalent (**Fig 6C**). Thus, we were able to reproduce the effects of outlier images on the phase transition between aligned and anti-aligned mosaics.

DISCUSSION

Efficient coding theory provides a basis for understanding many features of early sensory coding. This includes the structure of individual receptive fields, how those receptive fields adapt, and the mosaic-like arrangement of receptive fields to uniformly tile space (Atick and Redlich 1992, Barlow 1961, Doi et al. 2012, Karklin and Simoncelli 2011, Ratliff et al. 2009). We have extended this framework to investigate how the receptive fields of distinct cell types should be spatially arranged, given sensory noise and the statistics of natural stimuli. We term this the 'sensor alignment problem.' We discovered a second order phase transition between aligned and anti-aligned mosaics, in which the preferred phase depends on noise and image statistics. We show that this phase transition arises because both sensor noise and heavy-tailed stimulus distributions promote greater thresholds (rectification) in the output nonlinearities of the individual ON and OFF cells. Higher thresholds decorrelate the responses of nearby cells, causing an abrupt transition between aligned and anti-aligned configurations that occurs when the decorrelation region grows large enough to include multiple nearest neighbors.

This relationship between the response threshold and the phase transition can be understood in terms of redundancy reduction, specifically, reduction of co-activation of opposite polarity neurons. As our analysis in supplementary section 2B shows, when an ON-OFF pair of adjacent neurons are coactive, it means they encode an overlapping portion of the image distribution and are thus redundant. This diminishes information efficiency. Therefore, the best strategy to minimize information redundancy is for each neuron to minimize the number of opposite polarity adjacent neurons with which it is coactive. At a low threshold, this optimization only allows one cell at the same spatial location (**Fig 5D**), while at a high threshold,

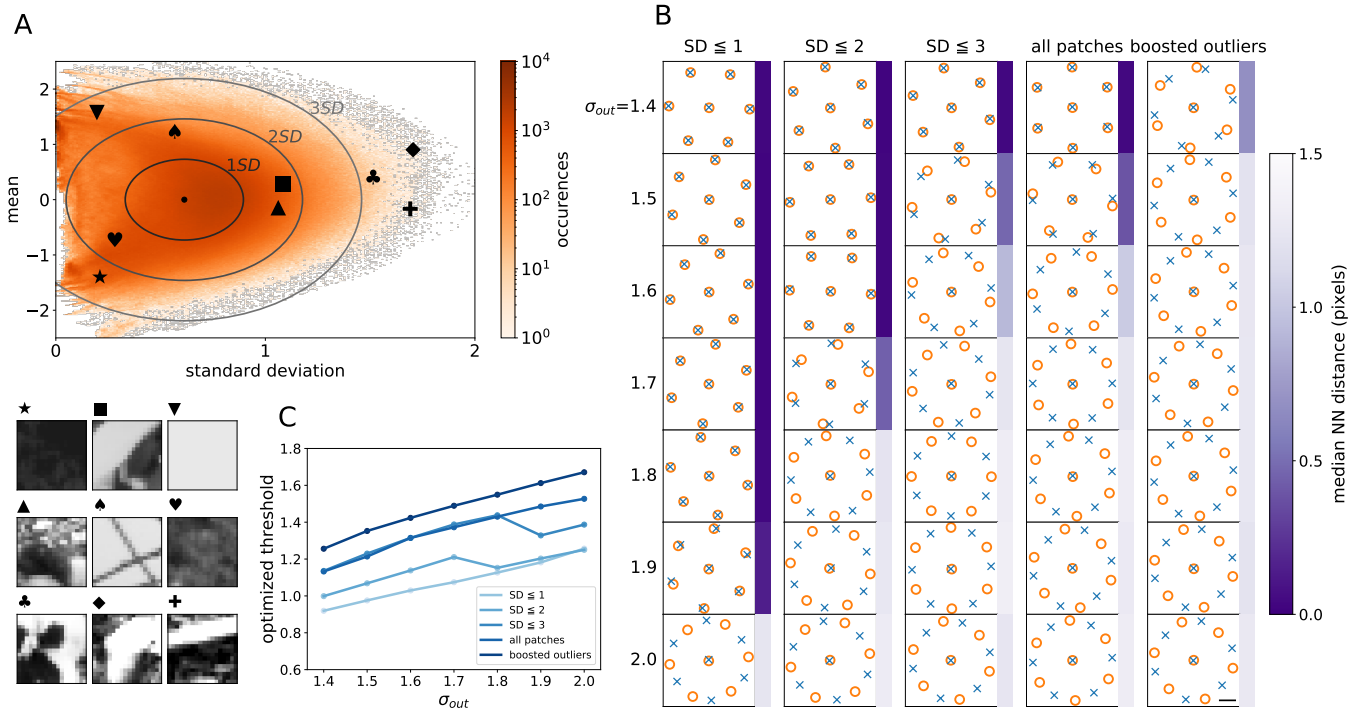

FIG. 6. **Outlier image patches drive an anti-aligned relationship between ON/OFF RF mosaics via higher optimal thresholds.**
A. Distribution of all 18,587,848 image patches from the dataset, plotted as a 2d-histogram with respect to the mean and the standard deviation of each image patch (orange). A few example patches are shown in the lower left, with the corresponding locations in the histogram marked with matching symbols. The means are centered at zero (sd =0.575), and the standard deviation values are centered at 0.623 (sd = 0.187). Ellipses in the histogram represent data within 1-3 standard deviations of the mean, which we use to construct artificial datasets with varying contributions of outliers. **B.** Optimized ON and OFF receptive field mosaics using the image patches within the boundaries denoted in **A** with σ_{in}=0.4, σ_{out}=1.4 to 2.0, plotted with the same scheme as in 3. As predicted, when trained on image sets that systematically excluded outliers (i.e., had lighter distributional tails), phase transitions happened only at larger values of output noise. The fourth column denoted as "all patches" used the full dataset as in 3. The last column, denoted as "boosted outliers," used an augmented dataset to which we added horizontal and vertical mirror flips of patches outside the 1.5 SD circle, thereby creating a distribution with even more outliers. This corresponds to the top 20% of z-scores among all patches in the dataset. Here, the phase transition happened at lower output noise values. The scale bar shows a 1-pixel distance. **C.** Optimal thresholds as a function of output noise, using input distributions with different levels of outliers. As predicted, the optimal thresholds tended to increase as the input distribution contained more outliers, allowing the phase transition to happen at a lower output noise level.

this optimization allows three cells (**Fig 5E**). In the latter case, a high threshold of nonlinearity enhances efficient coding even when their receptive fields have substantial (but not perfect) overlap (Pitkow and Meister 2012). It is worth noting that there is one bit of information shared between an ON and OFF neuron at the same spatial location (the sign of the stimulus) but the contribution from this is negligible compared to a pair of coactive neurons which redundantly encode much richer information about the input signal (supplementary section 2C). The above analysis provides an intuitive yet rigorous reasoning as to why anti-alignment is favored between the two mosaics when the threshold is high.

Our approach is closely related to several strands of work on efficient coding in the retina. As noted above, efficient coding as redundancy reduction subject to signaling costs gave theoretical weight to the idea of center-surround receptive fields as whitening filters for visual stimuli (Atick and Redlich 1990, 1992, Barlow 1961, Laughlin 2001, Srinivasan et al. 1982) and are thus matched to the statistics of natural scenes. However, as later work has made clear, nonlinearities, particularly response thresholds, are perhaps even more important in optimizing retinal codes (Gjorgjieva et al. 2019, 2014, Pitkow and Meister 2012). In addition, noise in both sensory transduction (inputs) and neural outputs affects coding efficiency (Atick and Redlich 1990, 1992, Röth et al. 2020). This work, by placing all of these factors in an optimization-based framework (Karklin and Simoncelli 2011), allows investigating the relative importance of each in the optimal spatial layout of receptive fields. This problem has received comparatively less attention than efficient coding in the temporal domain (Ocko et al. 2018, Palmer et

al. 2015, Sederberg et al. 2018). As in previous work, we found that noise and image statistics set optimal neural response thresholds and that these drive redundancy reduction. However, the crucial link between mosaic structure and response thresholds has not previously been demonstrated. Importantly, this analysis relies only on information maximization arguments and the optimization does not rely on fitting data from electrophysiological recordings, as required by recent deep learning models of retinal coding (Maheswaranathan et al. 2018, Ocko et al. 2018).

By focusing on the factors that govern the transition between aligned and anti-aligned mosaics from an optimal encoding standpoint, we have necessarily ignored several real biological complexities. Most importantly, we have not addressed whether mosaics in the retina are in fact aligned or anti-aligned. Most anatomical studies of retinal mosaics indicate they are statistically independent (Wässle et al., 1981; Rockhill et al., 2000; Eglen, 2012; but see Jang and Paik, 2017). However, this work uses cell body locations to judge the spatial arrangements between the mosaics, and cell body positions do not strictly determine the locations of receptive field centers. Recently, we have shown in rat and primate retinas that pairs of RF mosaics from ON and OFF RGCs that encode similar visual features are anti-aligned (Roy et al. 2021). This indicates that either retinal noise is relatively high and/or that retinal processing is optimized for low signal-to-noise conditions such as detecting dim or low contrast stimuli (Dhingra et al. 2003, Field and Sampath 2017). We have also found that measured RF mosaics from ON and OFF RGCs that encode distinct visual features appear to be statistically independent (Roy et al. 2021). Our current efficient coding model only generates two types of units that encode similar visual features, but with opposite polarity, thus we are not yet able to determine whether RGC types that encode distinct visual features should be independent or coordinated. Developing an efficient coding model that is optimized to encode natural movies may yield a greater diversity of cell types (Ocko et al. 2018). Many RGC types that encode particular space-time features such as direction of motion or 'looming' detectors also form mosaics (Vaney et al. 2012). Whether these detector grids are coordinated or should be coordinated remains an open question. Finally, our results motivate understanding the implications of efficient coding theory to other other sensory systems such as touch receptors on the skin: might they be spatially coordinated for the efficient coding of touch (Kuehn et al. 2019)? More generally, these observations demonstrate that efficient coding theory can make predictions about emergent properties present in the organization of the nervous system, such as how large populations composed of multiple cell types should be arranged to optimally encode the natural environment.

Atick JJ, Redlich AN (1990) Towards a theory of early visual processing. *Neural computation* 2:308–320.

Atick JJ, Redlich AN (1992) What does the retina know about natural scenes? *Neural computation* 4:196–210.

Baccus SA, Meister M (2002) Fast and slow contrast adaptation in retinal circuitry. *Neuron* 36:909–919.

Barlow HB (1961) Possible principles underlying the transformation of sensory messages. *Sensory communication* 1:217–234.

Bell C (1990) Mormyromast electroreceptor organs and their afferent fibers in mormyrid fish. iii. physiological differences between two morphological types of fibers. *Journal of neurophysiology* 63:319–332.

Binney JJ, Dowrick NJ, Fisher AJ, Newman ME (1992) *The theory of critical phenomena: an introduction to the renormalization group* Oxford University Press.

Borghuis BG, Ratliff CP, Smith RG, Sterling P, Balasubramanian V (2008) Design of a neuronal array. *Journal of Neuroscience* 28:3178–3189.

Chalasani SH, Chronis N, Tsunozaki M, Gray JM, Ramot D, Goodman MB, Bargmann CI (2007) Dissecting a circuit for olfactory behaviour in caenorhabditis elegans. *Nature* 450:63–70.

Chichilnisky E (2001) A simple white noise analysis of neuronal light responses. *Network: Computation in Neural Systems* 12:199–213.

Cover TM (1999) *Elements of information theory* John Wiley & Sons.

Devries SH, Baylor DA (1997) Mosaic arrangement of ganglion cell receptive fields in rabbit retina. *Journal of neurophysiology* 78:2048–2060.

Dhingra NK, Kao YH, Sterling P, Smith RG (2003) Contrast threshold of a brisk-transient ganglion cell in vitro. *Journal of neurophysiology* 89:2360–2369.

Doi E, Gauthier JL, Field GD, Shlens J, Sher A, Greschner M, Machado TA, Jepson LH, Mathieson K, Gunning DE et al. (2012) Efficient coding of spatial information in the primate retina. *Journal of Neuroscience* 32:16256–16264.

Doi E, Lewicki MS (2007) A theory of retinal population coding In *Advances in neural information processing systems*, pp. 353–360.

Eglen SJ (2012) Cellular spacing: Analysis and modelling of retinal mosaics In *Computational Systems Neurobiology*, pp. 365–385. Springer.

Field GD, Sampath AP (2017) Behavioural and physiological limits to vision in mammals. *Philosophical Transactions of the Royal Society B: Biological Sciences* 372:20160072.

Gallio M, Ofstad TA, Macpherson LJ, Wang JW, Zuker CS (2011) The coding of temperature in the drosophila brain. *Cell* 144:614–624.

Gauthier JL, Field GD, Sher A, Greschner M, Shlens J, Litke AM, Chichilnisky E (2009a) Receptive fields in primate retina are coordinated to sample visual space more uniformly. *PLoS Biol* 7:e1000063 Publisher: Public Library of Science.

Gauthier JL, Field GD, Sher A, Shlens J, Greschner M, Litke AM, Chichilnisky E (2009b) Uniform signal redundancy of

parasol and midget ganglion cells in primate retina. *Journal of Neuroscience* 29:4675–4680 Publisher: Soc Neuroscience.
Gjorgjieva J, Meister M, Sompolinsky H (2019) Functional diversity among sensory neurons from efficient coding principles. *PLoS computational biology* 15:e1007476.
Gjorgjieva J, Sompolinsky H, Meister M (2014) Benefits of pathway splitting in sensory coding. *Journal of Neuroscience* 34:12127–12144.
Jang J, Paik SB (2017) Interlayer repulsion of retinal ganglion cell mosaics regulates spatial organization of functional maps in the visual cortex. *Journal of Neuroscience* 37:12141–12152.
Joesch M, Schnell B, Raghu SV, Reiff DF, Borst A (2010) On and off pathways in drosophila motion vision. *Nature* 468:300–304.
Karklin Y, Simoncelli EP (2011) Efficient coding of natural images with a population of noisy linear-nonlinear neurons In *Advances in neural information processing systems*, pp. 999–1007.
Kingma DP, Ba J (2014) Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* .
Kuehn ED, Meltzer S, Abraira VE, Ho CY, Ginty DD (2019) Tiling and somatotopic alignment of mammalian low-threshold mechanoreceptors. *Proceedings of the National Academy of Sciences* 116:9168–9177.
Kuffler SW (1953) Discharge patterns and functional organization of mammalian retina. *Journal of neurophysiology* 16:37–68.
Laughlin SB (2001) Energy as a constraint on the coding and processing of sensory information. *Current opinion in neurobiology* 11:475–480.
Maheswaranathan N, Kastner DB, Baccus SA, Ganguli S (2018) Inferring hidden structure in multilayered neural circuits. *PLoS computational biology* 14:e1006291.
Meister M, Berry MJ (1999) The neural code of the retina. *Neuron* 22:435–450.
Ocko S, Lindsey J, Ganguli S, Deny S (2018) The emergence of multiple retinal cell types through efficient coding of natural movies In *Advances in Neural Information Processing Systems*, pp. 9389–9400.
Palmer SE, Marre O, Berry MJ, Bialek W (2015) Predictive information in a sensory population. *Proceedings of the National Academy of Sciences* 112:6908–6913.
Pitkow X, Meister M (2012) Decorrelation and efficient coding by retinal ganglion cells. *Nature neuroscience* 15:628–635.
Ratliff N, Zucker M, Bagnell JA, Srinivasa S (2009) Chomp: Gradient optimization techniques for efficient motion planning In *2009 IEEE International Conference on Robotics and Automation*, pp. 489–494. IEEE.
Rockhill RL, Euler T, Masland RH (2000) Spatial order within but not between types of retinal neurons. *Proceedings of the National Academy of Sciences* 97:2303–2307.
Röth K, Shao S, Gjorgjieva J (2020) Efficient population coding depends on stimulus convergence and source of noise. *bioRxiv* .
Roy S, Jun NY, Davis EL, Pearson J, Field GD (2021) Inter-mosaic coordination of retinal receptive fields .
Ruderman DL (1994) The statistics of natural images. *Network: computation in neural systems* 5:517–548.
Sederberg AJ, MacLean JN, Palmer SE (2018) Learning to make external sensory stimulus predictions using internal correlations in populations of neurons. *Proceedings of the National Academy of Sciences* 115:1105–1110.
Smith EC, Lewicki MS (2006) Efficient auditory coding. *Nature* 439:978–982.
Srinivasan MV, Laughlin SB, Dubs A (1982) Predictive coding: a fresh view of inhibition in the retina. *Proceedings of the Royal Society of London. Series B. Biological Sciences* 216:427–459.
Vaney DI, Sivyer B, Taylor WR (2012) Direction selectivity in the retina: symmetry and asymmetry in structure and function. *Nature Reviews Neuroscience* 13:194–208.
Wässle H, Peichl L, Boycott BB (1981) Morphology and topography of on-and off-alpha cells in the cat retina. *Proceedings of the Royal Society of London. Series B. Biological Sciences* 212:157–175.

I. SUPPLEMENTARY FIGURES

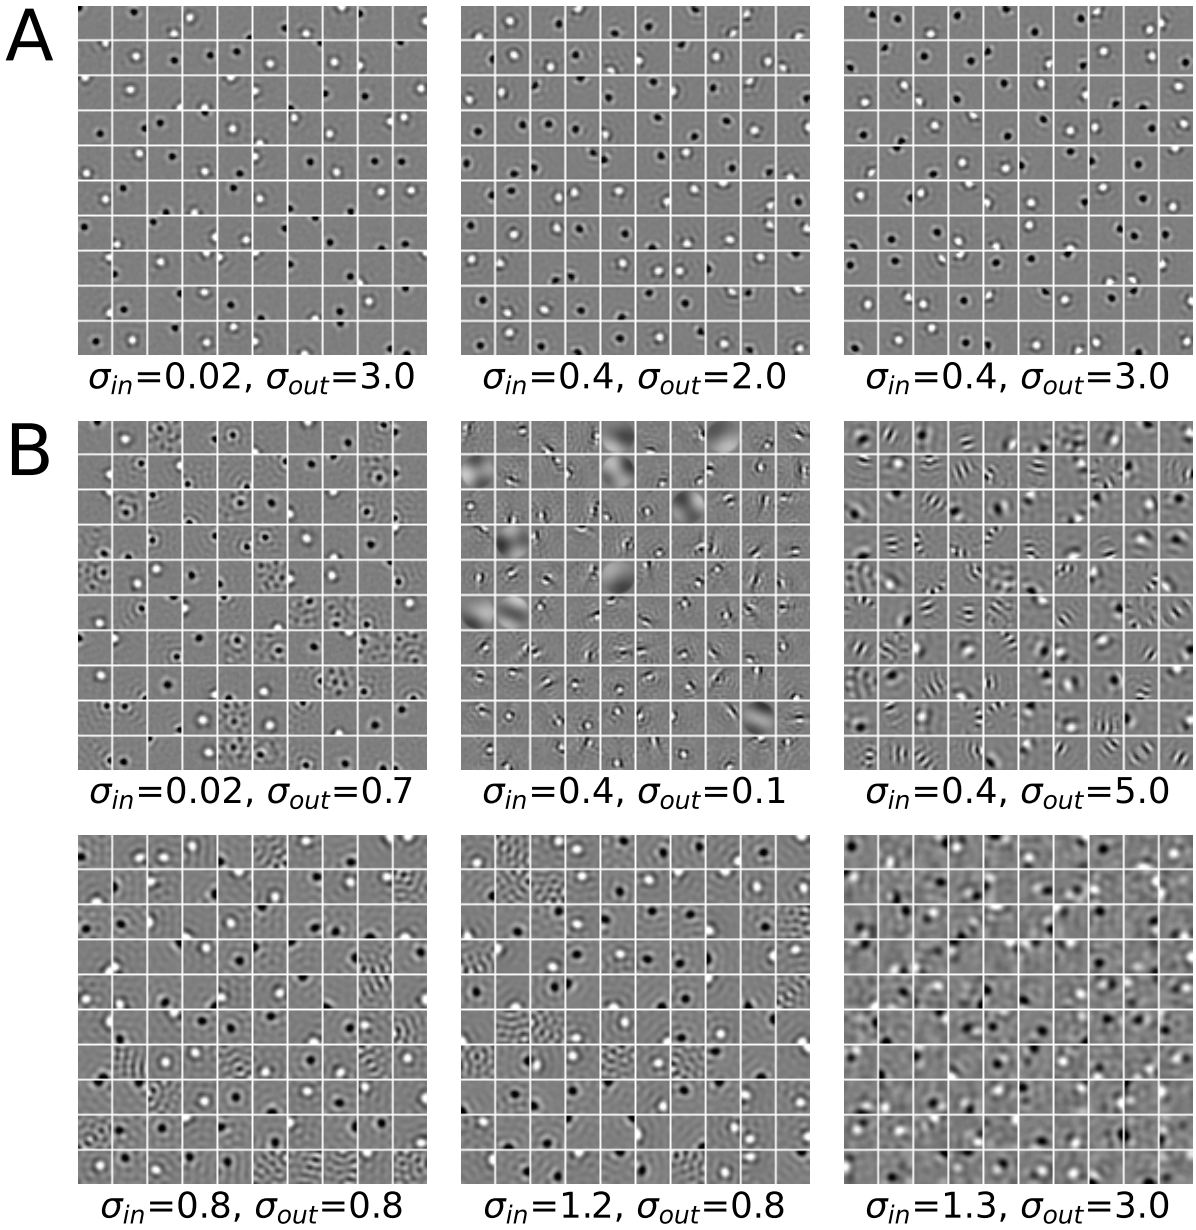

Supplementary Figure 1. **Circular ON and OFF receptive fields are robustly obtained only within specific input and output noise regimes.**
A. With input noise in the range between 0.00 to 0.75 and output noise between 0.75 to 3.00, circular center-surround ON and OFF receptive fields are obtained reliably. **B.** Outside these ranges, i.e., when input noise is higher than 0.75 or output noise is lower than 0.75, either Gabor-shaped or noisy kernels were obtained.

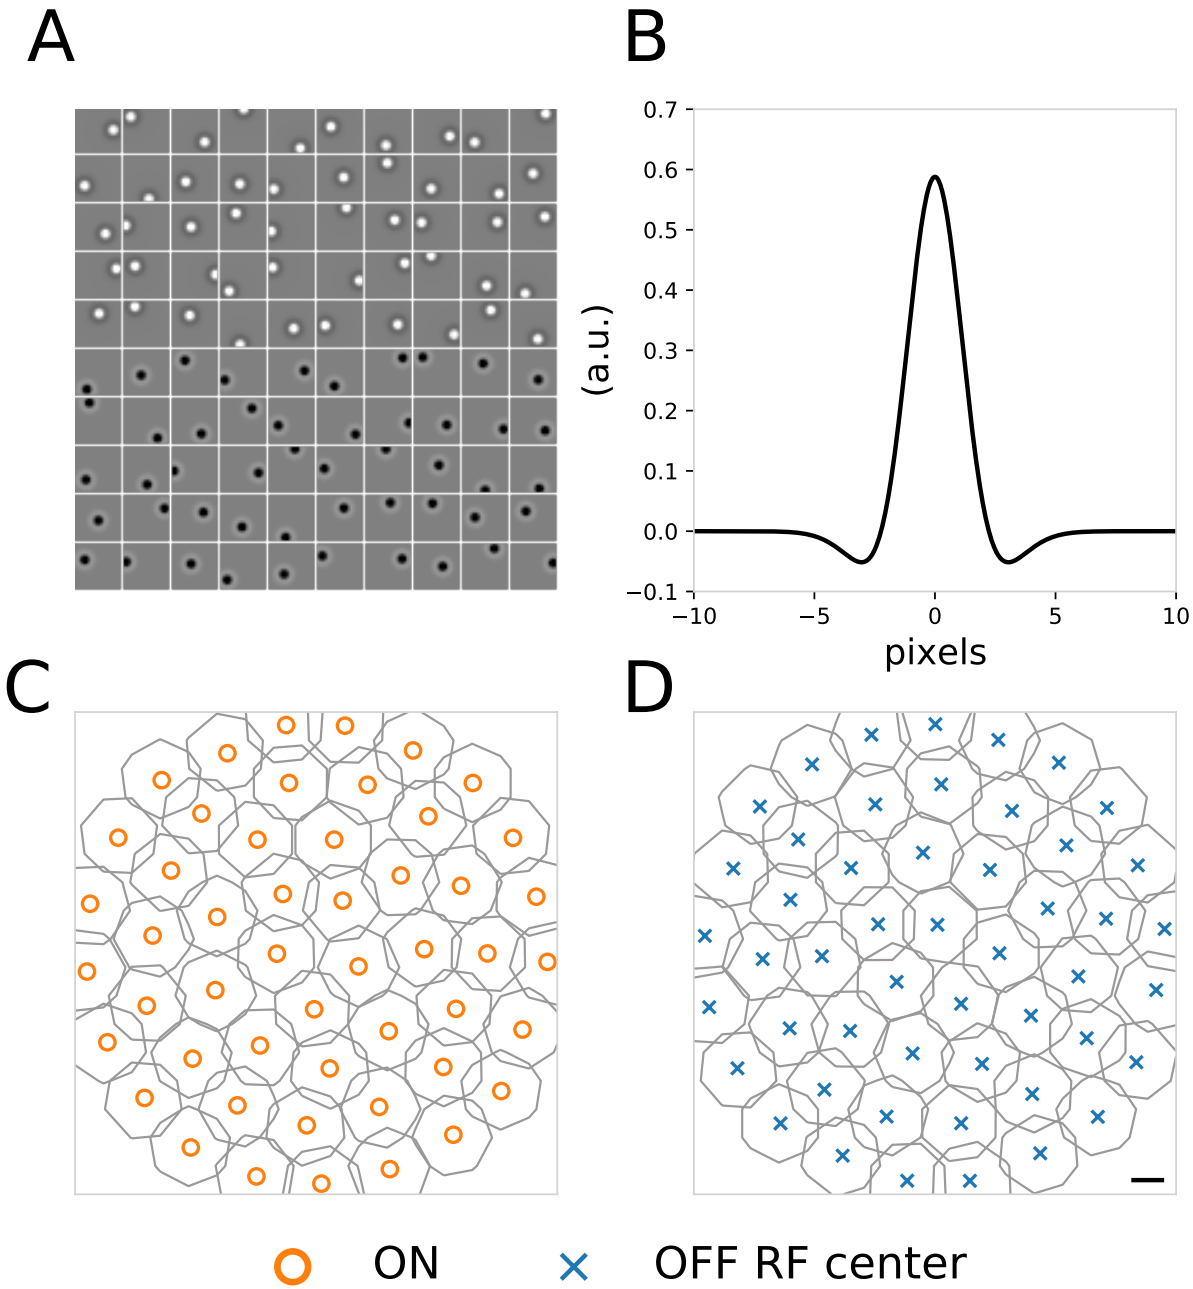

Supplementary Figure 2. **When all filters share the same kernel, filter locations still form mosaics.** **A.** A plot of 100 kernels that maximize information efficiency when all kernels are constrained to take the same shape but are free to move their centers (cf. **Fig 1C B.** The radial shape function is parameterized as a difference of Gaussian curves. Learned radial shape parameters are: a=0.3332, b=0.1559, c=0.4121. **C.** Learned ON kernels sharing the same radial shape fill the visual field (similar to **Fig 1D**). **D.** Learned OFF kernels sharing the same radial shape fill the visual field (similar to **Fig 1E**). The scale bar shows a 1-pixel distance.

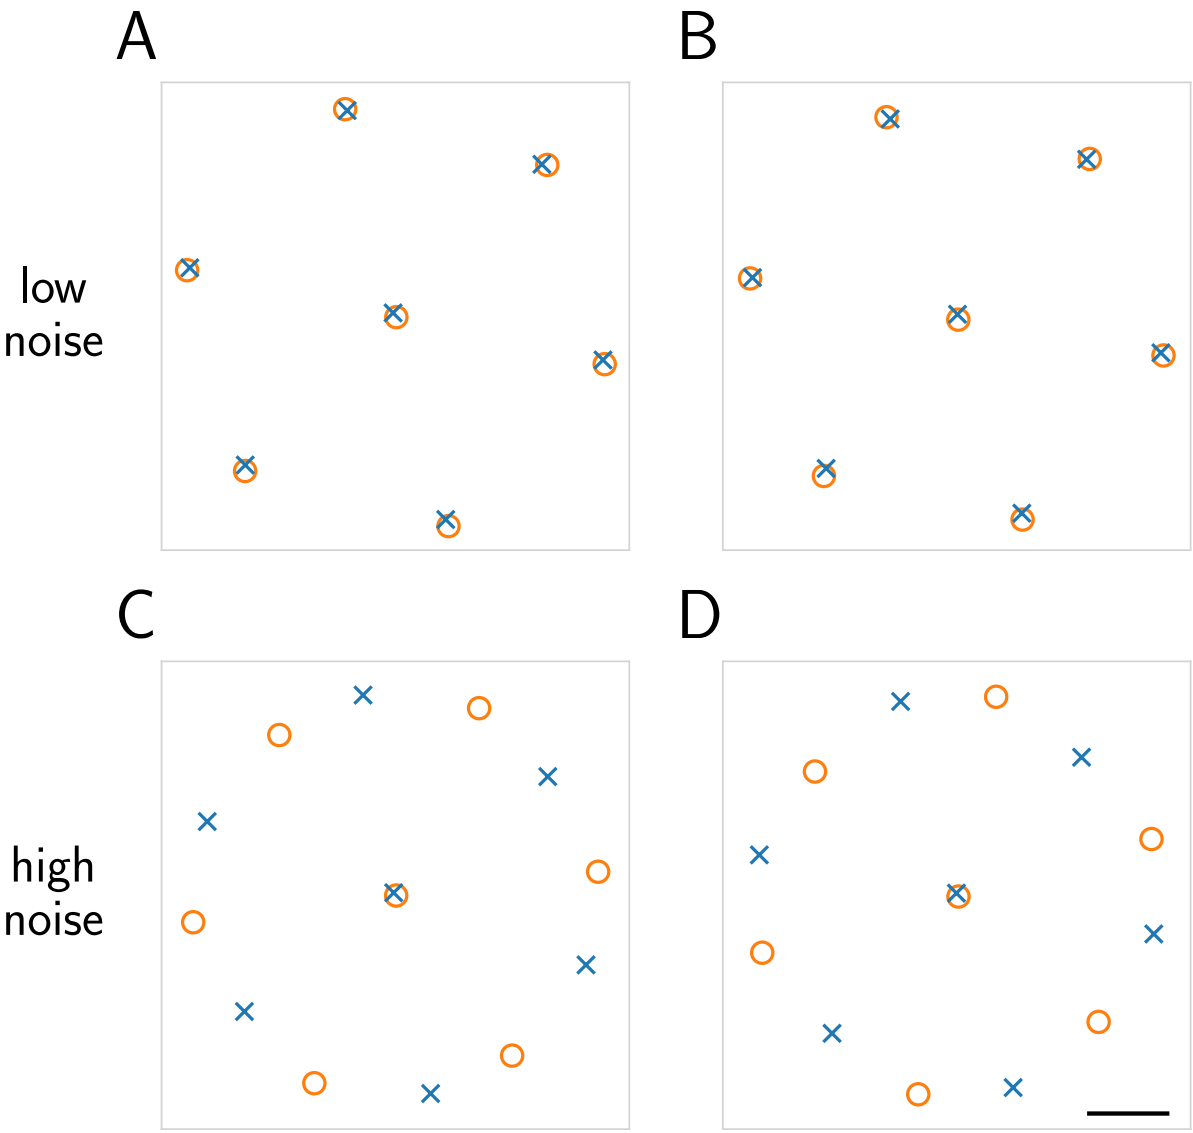

Supplementary Figure 3. **A one-shape model also captures the spatial coordination of receptive field mosaics depending on input and output noise level.**
Receptive field centers for ON (orange circle) and OFF (blue x) cells under differing sets of noise parameters. **A:** (σ_{in}=0.02, σ_{out}=1.0), **B:** (σ_{in}=0.1, σ_{out}=1.0), **C:** (σ_{in}=0.4, σ_{out}=2.0), **D:** (σ_{in}=0.4, σ_{out}=3.0). The first two parameter sets result in aligned mosaics, while the latter two, at higher levels of output noise, are anti-aligned. The scale bar shows a 1-pixel distance.

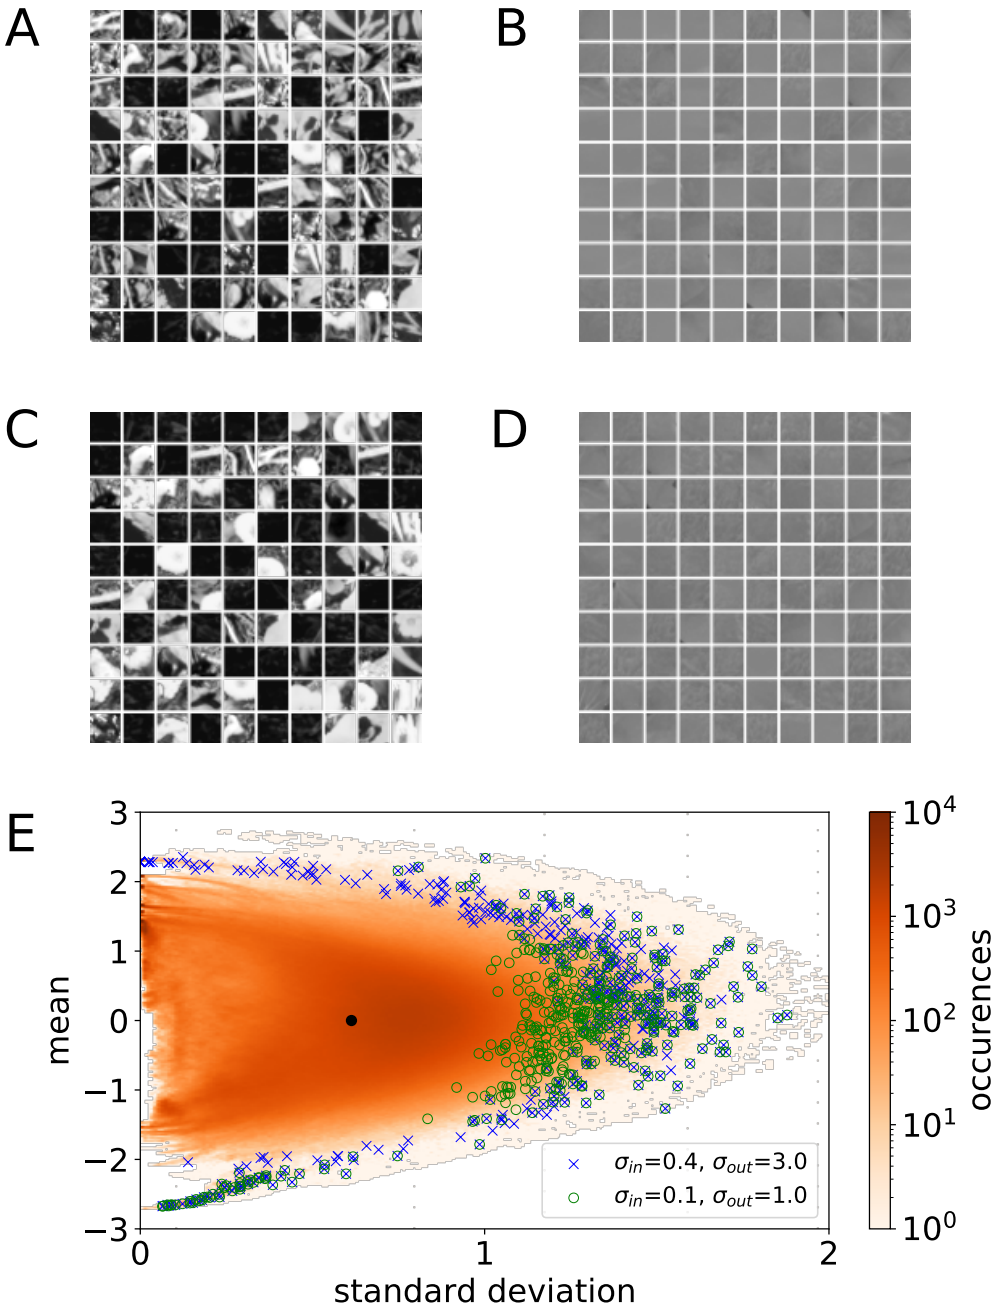

Supplementary Figure 4. **Image patches with the highest firing rates and lowest firing rates (among 100,000 randomly sampled image patches) are distinct in both the high noise ($\sigma_{\mathbf{in}}$=0.4, $\sigma_{\mathbf{out}}$=3.0) and the low noise ($\sigma_{\mathbf{in}}$=0.1, $\sigma_{\mathbf{out}}$=1.0) regime.**
Most of the image patches with the highest firing rates are outliers of the overall image patch distribution, containing high contrast and edges. This is also evident in firing rate outputs: Using the kernels optimized under the two noise regimes, as shown in **Fig 2B** and **Fig 2D**, we calculate the firing rates of 100,000 randomly sampled image patches and overlay the top 100 patches with the highest firing rates over the z-score histogram introduced in **Fig 6A**. **A.** 100 image patches with the highest firing rates under σ_{in}=0.1, σ_{out}=1.0. Images are either extremely dark or contain distinct edges, including those with more complicated edges compared to **C**, and had firing rates of 3.07±0.23. **B.** 100 image patches with the lowest firing rates under σ_{in}=0.1, σ_{out}=1.0. The 100 images with the lowest firing rates were also visually similar and had firing rates of 0.161±0.003. **C.** 100 image patches with the highest firing rates under σ_{in}=0.4, σ_{out}=3.0. Images are either dark or contain distinct edges, with firing rates of 4.39±0.54. **D.** 100 image patches with the lowest firing rates under σ_{in}=0.4, σ_{out}=3.0. Patches were almost entirely gray, with firing rates of 0.036±0.001. **E.** A 2-D histogram showing the mean and standard deviation of all of the available sample image patches (in orange, occurrences in log scale, same as **Fig 6A**) with the image patches in **A** (blue x) and the image patches in **B** (green o). In the low-noise condition, the outliers are a greater portion of the tail (green o).

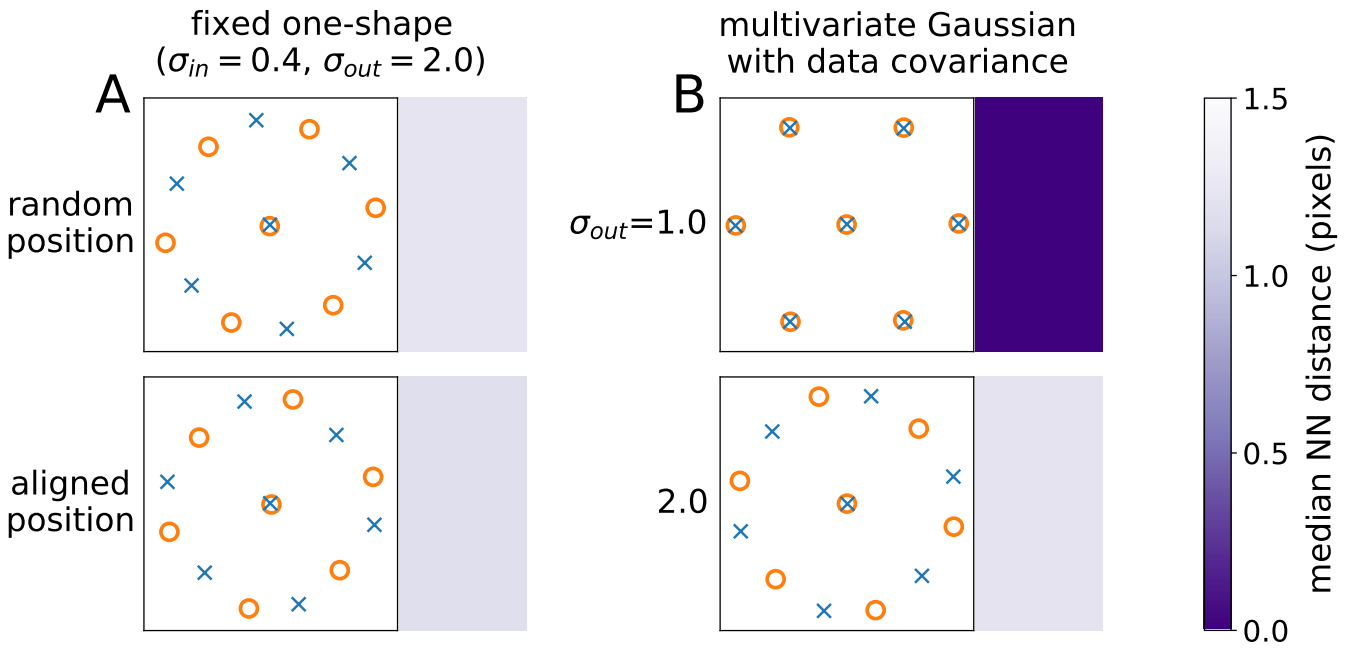

Supplementary Figure 5. **The phase transition is still observed with a fixed kernel shape and Gaussian image data.**
A. Even when kernel parameters are not learnable (fixed to a, b, c = [0.258, 0.162, 0.648]), the optimal configuration for the mosaics is still anti-alignment. This is true regardless of whether the receptive fields are initialized at random (top) or aligned (bottom) initial positions (cf. **Fig 5B** top row of "all patches"). Here $\sigma_{\rm in}$=0.4, $\sigma_{\rm out}$=1.0. **B.** The transition to anti-alignment still occurs when images are not natural images but drawn from a covariance-matched Gaussian distribution. This matches the assumptions of our analytical model and demonstrates that heaver-than-Gaussian tails are not necessary for anti-alignment to be favored.

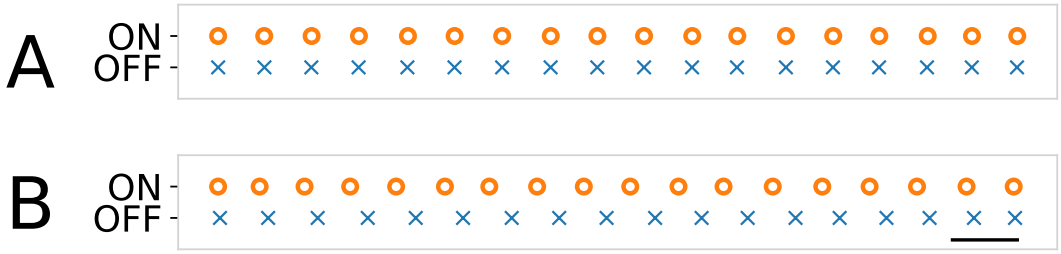

Supplementary Figure 6. The phase transition is still observed with a 1-dimensional model. Optimal receptive field centers of 18 ON and 18 OFF RFs in a 1-by-64 pixel input. Receptive field centers for ON (orange circle) and OFF (blue x) cells are optimized under two sets of noise parameters, A: ($\sigma_{\rm in}$=0.4, $\sigma_{\rm out}$=1.0), B: ($\sigma_{\rm in}$=0.4, $\sigma_{\rm out}$=3.0). **A** results in aligned mosaics, while **B** results anti-aligned mosaics. The scale bar shows a 5-pixel distance.

II. MODELING DETAILS

A. Mutual information in the linear-nonlinear model

We follow the conventions of Karklin and Simoncelli (2011). Let $\mathbf{x} \in \mathbb{R}^D$ be a vectorized representation of an input image. We assume J neurons characterized by linear filters $\mathbf{w}_j$ ($\|\mathbf{w}_j\| = 1$) for $j = 1 \ldots N$, which form the columns of a matrix $\mathbf{W} \in \mathbb{R}^{D \times J}$. We assume pixelwise input noise to each image $\mathbf{n}_x \sim \mathcal{N}(0, \sigma_{\text{in}}^2)$ and output noise in firing rates $\mathbf{n}_r \sim \mathcal{N}(0, \sigma_{\text{out}}^2)$ so that the random input image and firing rate are $\mathbf{X} = \mathbf{x} + \mathbf{n}_x$ and $\mathbf{R} = \mathbf{r} + \mathbf{n}_r$, respectively. With these conventions, the firing rate for the jth neuron given input $\mathbf{x}$ is

$$r_j = \gamma_j h(\mathbf{w}_j^{\mathsf{T}} (\mathbf{x} + \mathbf{n}_x) - \theta_j) + n_{r,j}, \tag{7}$$

with γ_j the neuron gain, θ_j its response threshold, and h the softplus nonlinearity described in Methods. Note that, in principle, this firing can be negative for sufficiently large and negative $n_{r,j}$. As in Karklin and Simoncelli (2011), we notheless treat **Eq 7** as a useful approximation.

Under these assumptions, the covariance matrix of firing rates given input image is

$$\mathbf{C}_{r|x} = \mathbf{G}\mathbf{W}^{\mathsf{T}}\mathbf{C}_{\text{in}}\mathbf{W}\mathbf{G} + \mathbf{C}_{\text{out}} \tag{8}$$

with $\mathbf{C}_{\text{in}} = \sigma_{\text{in}}^2 \mathbb{1}$, $\mathbf{C}_{\text{out}} = \sigma_{\text{out}}^2 \mathbb{1}$, and $\mathbf{G}(\mathbf{x}) = \operatorname{diag}(\gamma_j f'(\mathbf{w}^{\mathsf{T}}(\mathbf{x} + \mathbf{n}_x) - \theta))$ a matrix of derivatives of the nonlinearities. As in Karklin and Simoncelli (2011), we make a local Gaussian approximation to the image distribution with covariance $\mathbf{C}_x$, allowing us write the posterior $p(\mathbf{x}|\mathbf{r})$ as Gaussian with covariance

$$\begin{aligned}\mathbf{C}_{x|r} &= \left(\mathbf{C}_x^{-1} + \mathbf{W}\mathbf{G}\mathbf{C}_{r|x}^{-1}\mathbf{G}\mathbf{W}^{\mathsf{T}}\right)^{-1} \\ &= \left(\mathbf{C}_x^{-1} + \mathbf{W}\mathbf{G}\left(\mathbf{G}\mathbf{W}^{\mathsf{T}}\mathbf{C}_{\text{in}}\mathbf{W}\mathbf{G} + \mathbf{C}_{\text{out}}\right)^{-1}\mathbf{G}\mathbf{W}^{\mathsf{T}}\right)^{-1}\end{aligned} \tag{9}$$

and the conditional (differential) entropy as

$$H(\mathbf{X}|\mathbf{R}) = \mathbb{E}_{\mathbf{x}}\left[\frac{1}{2}\log\det\left(2\pi e \mathbf{C}_{x|r}\right)\right]. \tag{10}$$

Since $H(\mathbf{X})$ is an unknown constant, maximizing the mutual information $\mathcal{I}(\mathbf{X};\mathbf{R}) = H(\mathbf{X}) - H(\mathbf{X}|\mathbf{R})$ involves *minimizing* **Eq 10** subject to a cost per spike that sets constrains the mean firing rate of each neuron:

$$\begin{aligned}\underset{\mathbf{W},\boldsymbol{\gamma},\boldsymbol{\theta}}{\text{maximize}} \quad & \mathbb{E}_{\mathbf{x}}\left[\log\frac{\det\left(\mathbf{G}\mathbf{W}^{\mathsf{T}}(\mathbf{C}_x + \mathbf{C}_{\text{in}})\mathbf{W}\mathbf{G} + \mathbf{C}_{\text{out}}\right)}{\det\left(\mathbf{G}\mathbf{W}^{\mathsf{T}}\mathbf{C}_{\text{in}}\mathbf{W}\mathbf{G} + \mathbf{C}_{\text{out}}\right)}\right] \\ \text{subject to} \quad & \mathbb{E}\left[\mathbf{R}\right] = 1,\end{aligned} \tag{11}$$

where the optimization objective follows from **Eqs 9** and **10** via the Matrix Inversion Lemma and we implement the constraint using an augmented Lagrangian method as detailed in Methods. Finally, the following rewriting of **Eq 11** will be useful, particularly in our analysis of the single-neuron case:

$$\begin{aligned}&\frac{\det\left(\mathbf{G}\mathbf{W}^{\mathsf{T}}(\mathbf{C}_x + \mathbf{C}_{\text{in}})\mathbf{W}\mathbf{G} + \mathbf{C}_{\text{out}}\right)}{\det\left(\mathbf{G}\mathbf{W}^{\mathsf{T}}\mathbf{C}_{\text{in}}\mathbf{W}\mathbf{G} + \mathbf{C}_{\text{out}}\right)} \\ &= \det\left(\mathbb{I} + \mathbf{G}\mathbf{W}^{\mathsf{T}}\mathbf{C}_x\mathbf{W}\mathbf{G}\left(\mathbf{G}\mathbf{W}^{\mathsf{T}}\mathbf{C}_{\text{in}}\mathbf{W}\mathbf{G} + \mathbf{C}_{\text{out}}\right)^{-1}\right).\end{aligned} \tag{12}$$

B. Analytic model of alignment in 1-D

To better understand the optimization problem posed in **Eq 11**, particularly the contributions of kernel shape, image statistics, noise, and response thresholds to the transition from aligned to anti-aligned mosaics, we consider a simplified model that nonetheless exhibits the same behavior. In particular:

1. We restrict ourselves to a one-dimensional space in which ON and OFF mosaics are assumed to be an equidistant array of RF centers (distance = 1) for both mosaics. The ON cells are located at positions $z_j = j$, $j \in \mathbb{Z}$, and the mosaics are allowed to rigidly shift relative to each other, so that OFF cells are located at $\bar{z}_j = j + \phi$, $\phi \in [0, 0.5]$.

2. We assume a fixed kernel shape for all neurons, meaning that receptive fields differ only in their center locations: $w_j(z) = \pm\kappa(|z - z_j|)$ for a kernel shape function $\kappa(z)$, where z_j is the location of j-th neuron's kernel center.

3. We likewise assume identical nonlinearities for all units: $f_j(y) = \gamma(y - \theta)_+$, where γ is the gain, θ the response threshold parameter, and $(\cdot)_+ = \max(\cdot, 0)$ denotes positive rectification.

4. Images $\mathbf{x}$ are assumed to follow a multivariate Gaussian distribution $\mathcal{N}(0, \mathbf{C}_x)$, where the covariance between locations z_i and z_j only depends on $|z_i - z_j|$, defined as a function $C_x(z)$. In analogy with the finding that the amplitudes of natural images decrease as $\sim 1/k$ with k the spatial frequency (power spectral density $\sim 1/k^2$), for our numerical examples, we likewise assume a *one-dimensional* scale-free distribution $\sim 1/\sqrt{k}$ (power spectrum $\sim 1/k$). As shown in **Fig S5B**, the phase transition still takes place even with this Gaussian image assumption.

Now, for a pair of neurons j and k, the scalars $\mathbf{w}_j^{\mathsf{T}}\mathbf{w}_k$ and $\mathbf{w}_j^{\mathsf{T}}\mathbf{C}_x\mathbf{w}_k$ depend only on $|z_j - z_k|$, the distance between their centers, which we write in the frequency domain as:

$$\begin{aligned}\mathbf{w}_j^{\mathsf{T}}\mathbf{w}_k &= \frac{1}{2\pi}\int e^{i(z_j - z_k)\omega}|\kappa(\omega)|^2 d\omega \\ &\overset{\text{def}}{=} f(|z_j - z_k|)\end{aligned} \tag{13}$$

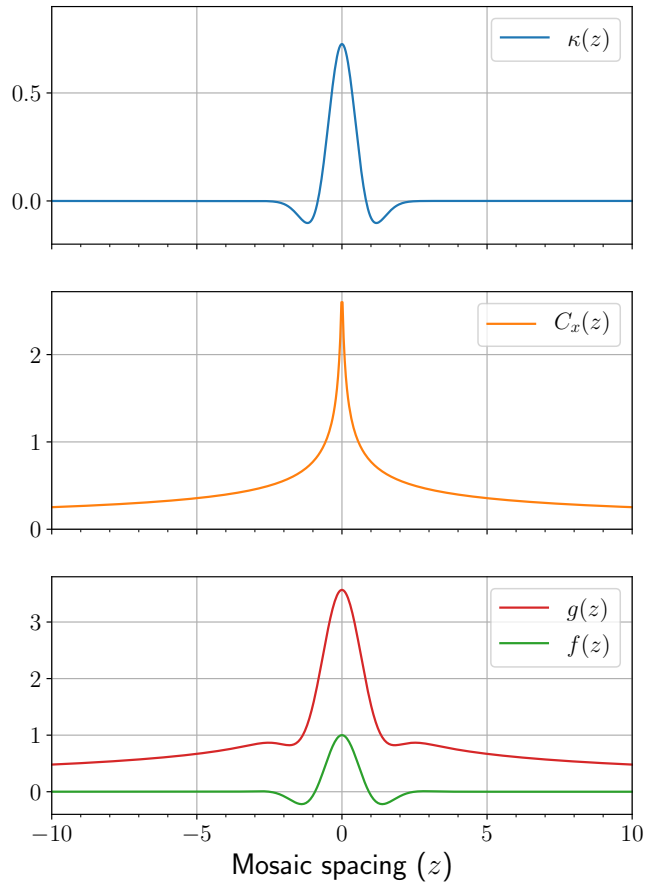

Supplementary Figure 7. Plots of $\kappa(z)$, $C_x(z)$, $f(z)$, and $g(z)$ using trained parameters from a one-shape model and an approximation of the data covariance: $\kappa(z) \propto e^{-1.76z^2} - 0.65e^{-1.12z^2}$ and $C_x(z) = \frac{0.8}{\sqrt{|z|+0.53}}$.

and

$$\begin{aligned} \mathbf{w}_j^\mathsf{T}\mathbf{C}_x\mathbf{w}_k &= \frac{1}{2\pi}\int e^{i(z_j-z_k)\omega}C_x(\omega)|\kappa(\omega)|^2 d\omega \\ &\stackrel{\text{def}}{=} g(|z_j - z_k|), \end{aligned} \tag{14}$$

where $C_x(\omega)$ and $\kappa(\omega)$ are Fourier transforms of $C_x(z)$ and $\kappa(z)$, respectively. Using any $\mathbf{w}$, we can define the constants f_0 and g_0 for the special case where $z_j = z_k$:

$$f_0 \stackrel{\text{def}}{=} f(0) = \mathbf{w}^\mathsf{T}\mathbf{w} = \|\mathbf{w}\| = 1 \tag{15}$$

$$g_0 \stackrel{\text{def}}{=} g(0) = \mathbf{w}^\mathsf{T}\mathbf{C}_x\mathbf{w} \tag{16}$$

where we used our assumption that the kernel $\mathbf{w}$ is a unit vector in **Eq 15**. An example of these functions is shown in **Sup Fig 7**.

The single-neuron case: $\mathcal{I}_1$

We first consider the mutual information capacity of a single neuron in isolation. Since we have assumed all units to be identical, the same expression holds for all ON and OFF cells, and we let $\mathcal{I}_1 \stackrel{\text{def}}{=} -H(X|R)$ in this single-channel case. From **Eq 12**, this quantity is

$$\begin{aligned} \mathcal{I}_1 &= p_1 \log\left(1 + \frac{\gamma^2\mathbf{w}^\mathsf{T}\mathbf{C}_x\mathbf{w}}{\gamma^2\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2}\right) \\ &= p_1 \log\left(1 + \frac{g_0}{\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2\gamma^{-2}}\right), \end{aligned} \tag{17}$$

where $p_1 = \mathbb{E}_{\mathbf{x}}[\mathbb{I}\{\mathbf{w}^\mathsf{T}(\mathbf{x}+\mathbf{n}_x) > \theta\}]$ is the probability that the neuron is active, i.e., that $\mathbf{w}^\mathsf{T}(\mathbf{x}+\mathbf{n}_x) > \theta$. Now, since $\mathbf{x}+\mathbf{n}_x$ is Gaussian with mean 0 and covariance $\mathbf{C}_x + \mathbf{C}_{\text{in}}$, we have $\mathbf{w}^\mathsf{T}(\mathbf{x}+\mathbf{n}_x) \sim \mathcal{N}(0, g_0 + \sigma_{\text{in}}^2)$ and

$$p_1 = 1 - \Phi\left(\frac{\theta}{\sqrt{g_0 + \sigma_{\text{in}}^2}}\right) = S\left(\frac{\theta}{\sqrt{g_0 + \sigma_{\text{in}}^2}}\right), \tag{18}$$

where Φ and S are the cumulative distribution function and the survival function of the standard normal distribution, respectively. Note that if we define $\varsigma^2 \stackrel{\text{def}}{=} g_0 + \sigma_{\text{in}}^2$ to denote the signal variance and $\tilde{\theta} \stackrel{\text{def}}{=} \theta/\varsigma$ to denote the effective threshold, we see that the effective threshold increases when signal variance decreases. We can calculate γ as:

$$\begin{aligned} \gamma^{-1} &= \int_\theta^\infty (y-\theta)\frac{1}{\sqrt{2\pi(g_0+\sigma_{\text{in}}^2)}}e^{-\frac{y^2}{2\left(g_0+\sigma_{\text{in}}^2\right)}}dy \\ &= \sqrt{\frac{g_0+\sigma_{\text{in}}^2}{2\pi}}e^{-\frac{\theta^2}{2\left(g_0+\sigma_{\text{in}}^2\right)}} - \theta p_1 \\ &= \varsigma\left(\frac{1}{\sqrt{2\pi}\varsigma}e^{-\frac{\theta^2}{2\varsigma^2}} - \frac{\theta}{\varsigma}S\left(\frac{\theta}{\varsigma}\right)\right) \\ &= \varsigma\left(\varphi(\tilde{\theta}) - \tilde{\theta}S(\tilde{\theta})\right) \end{aligned} \tag{19}$$

where φ is the probability density function of the standard normal distribution. Substituting this to **Eq 17**, we get:

$$\mathcal{I}_1 = S(\tilde{\theta})\log\underbrace{\left(1 + \frac{g_0}{\sigma_{\text{in}}^2 + \varsigma^2\sigma_{\text{out}}^2\left(\varphi(\tilde{\theta}) - \tilde{\theta}S(\tilde{\theta})\right)^2}\right)}_{\stackrel{\text{def}}{=}A_1}, \tag{20}$$

which was used to plot the single-neuron mutual information contributions in **Fig 4A**.

A single-neuron case with heavy tails

To investigate the relationship between threshold and outliers in the image distribution, we performed the

analysis in **Fig 4B**, which replaced the Gaussian image distribution with a Student's t-distribution. This allowed us to adjust the heaviness of the tails of the distribution by changing the degrees of freedom from $\nu = 4$ (heavy tails) to $\nu = \infty$ (Gaussian). In this case, a formula similar to **Eq 20** can be derived in a similar manner to **Eq 19**.

More explicitly, if $\mathbf{x} \sim t_\nu(0, \mathbf{C}_x)$, then $\mathbf{w}^\mathsf{T}(\mathbf{x} + \mathbf{n}_x)/\varsigma$ is likewise t_ν-distributed, and we have:

$$p_1 = 1 - \Phi_\nu\left(\frac{\theta}{\varsigma}\right). \tag{21}$$

where Φ_ν is the cumulative distribution function of the t-distribution. Moreover,

$$\begin{aligned}
&\gamma^{-1} \\
&= \int_{\frac{\theta}{\varsigma}}^{\infty} (\varsigma t - \theta) \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\Gamma\left(\frac{\nu}{2}\right)} \left(1 + \frac{t^2}{\nu}\right)^{-\frac{\nu+1}{2}} dt \\
&= \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\Gamma\left(\frac{\nu}{2}\right)} \frac{\varsigma\nu}{2} \int_{\frac{\theta^2}{\varsigma^2\nu}}^{\infty} (1+u)^{-\frac{\nu+1}{2}} du - \theta p_1 \\
&= \frac{\Gamma\left(\frac{\nu+1}{2}\right)}{\sqrt{\nu\pi}\Gamma\left(\frac{\nu}{2}\right)} \frac{\varsigma\nu}{\nu-1} \left(1 + \frac{\theta^2}{\varsigma^2\nu}\right)^{-\frac{\nu-1}{2}} - \theta p_1 \\
&= \varsigma\left(\frac{\nu + \tilde{\theta}^2}{\nu - 1}\varphi_\nu\left(\tilde{\theta}\right) - \tilde{\theta} S_\nu\left(\tilde{\theta}\right)\right)
\end{aligned} \tag{22}$$

where φ_ν and S_ν are respectively the probability density function and the survival function of Student's t-distribution with ν degrees of freedom. We see that $\frac{\nu+\tilde{\theta}^2}{\nu-1} \to 1$ as $\nu \to \infty$, in which case **Eq 22** is identical to **Eq 19**. These results can then be substituted into **Eq 17** to derive an expression analogous to **Eq 20**.

Decomposition of the mutual information

In previous sections, we have considered the encoding properties of a single neuron in isolation. Here, we return to the full optimization objective in **Eq 11**, showing how we can rewrite this as a sum of independent, single-neuron contributions plus pairwise (and higher-order) correction terms.

We begin by noting that the expectation in **Eq 11** over images $\mathbf{x}$ involves a sum over patches that only activate subsets of neurons. More specifically, with the hard rectifying nonlinearity assumed in the previous section, we have $\mathbf{G}(\mathbf{x}) = \gamma\mathbb{I}\{\mathbf{w}^\mathsf{T}(\mathbf{x} + \mathbf{n}_x) > \theta\}$ for ON cells and the same for OFF cells with $\theta \to -\theta$. Here, $\mathbb{I}$ is the indicator function, which is 1 when its argument is true and 0 otherwise. Thus, the gain matrix is diagonal, with entries equal to either 0 or γ, and the matrix products involving $\mathbf{G}(\mathbf{x})$ in **Eq 11** are thus low-rank for most $\mathbf{x}$. As a result, we can rewrite the expectation as a sum over log determinants of low-rank matrices, each of which corresponds to a specific collection of active neurons.

For example, consider a three-neuron system, whose activity can be coded as a binary triple, (e.g., 010). Let A be the determinant in **Eq 12**. Letting $p_{\text{(triplet)}}$ denote the probability that each subset is activer over the image set and $A_{\text{(triplet)}}$ be the corresponding determinant ratio, the expectation can be expanded as:

$$\begin{aligned}
\mathbb{E}\left[\log A\right] = \; & p_{000}\log A_{000} + p_{001}\log A_{001} \\
& + p_{010}\log A_{010} + p_{011}\log A_{011} \\
& + p_{100}\log A_{100} + p_{101}\log A_{101} \\
& + p_{110}\log A_{110} + p_{111}\log A_{111}.
\end{aligned} \tag{23}$$

Unfortunately, even for a Gaussian image distribution, calculating p for a specific triple involves a high-dimensional Gaussian integral bounded by a triple of hyperplanes in image space. However, we can rewrite this collection of terms as:

$$\begin{aligned}
\mathbb{E}\left[\log A\right] = \; & p_{000}\log A_{000} \\
& + (p_{100} + p_{110} + p_{101} + p_{111})\log A_{100} \\
& + (p_{010} + p_{110} + p_{011} + p_{111})\log A_{010} \\
& + (p_{001} + p_{101} + p_{011} + p_{111})\log A_{001} \\
& + (p_{110} + p_{111})\log\frac{A_{110}}{A_{100}A_{010}} \\
& + (p_{101} + p_{111})\log\frac{A_{101}}{A_{100}A_{001}} \\
& + (p_{011} + p_{111})\log\frac{A_{011}}{A_{010}A_{001}} \\
& + p_{111}\log\left(A_{111}\frac{A_{100}A_{010}A_{001}}{A_{110}A_{101}A_{011}}\right)
\end{aligned} \tag{24}$$

The first term corresponds to the case when no neuron is active, for which $\log A_{000} = 0$. For the next three terms, each of the probabilities in parentheses is a *marginal* probability of a neuron being active. That is, it is equal to p_1 as calculated above. Similarly, the next three terms constitute pairwise corrections to the first set of terms: they involve marginal probabilities of pairwise activation, and their log ratios vanish when the individual neurons are independent (e.g., $A_{110} = A_{100}A_{010}$). Thus, as the distance between a pair of receptive field centers increases beyond the image correlation length scale, these correction terms should vanish.

In our analysis, we focus on the limit in which only first- and second-order interactions are non-negligible. This is an assumption that does not, in fact, hold for natural images (or even our Gaussian images with long-range correlations), though as we shall see, these interactions are sufficient to demonstrate the existence of a phase transition. Moreover, as Figures 5 and 6 show, the same phase transition occurs in simulation when images with long-range correlations are used, suggesting that the principles we identify in the pairwise limit generalize to the full model.

Pairwise interactions: $\mathcal{I}_{2'}$

Returning to our 1-D model with N ON and N OFF neurons, we note that, as per **Eq 5**, we can further decompose the pairwise corrections into two terms, one corresponding to ON-ON and OFF-OFF corrections, the other to ON-OFF pairs. More explicitly, if we denote by a_j and $\bar{a}_j$ the binary variables corresponding to whether the ON cell at location j and the OFF cell at location $j+\phi$ are active, respectively, we can write the mutual information up to pairwise corrections as:

$$2Np_1\log A_1 + \underbrace{2N\sum_{j\neq 0}\mathbb{E}_{\mathbf{x}}[a_0 a_j]\log\frac{A_{2j}}{A_1^2}}_{\mathcal{I}_2} + \underbrace{N\sum_{j}\mathbb{E}_{\mathbf{x}}[a_0 \bar{a}_j]\log\frac{\bar{A}_{2j}}{A_1^2}}_{\mathcal{I}_{2'}} \tag{25}$$

where the first term contains $\mathcal{I}_1$ from **Eq 20**, and the next two terms correspond to pairwise interactions between neurons of the same and different polarities, respectively. Note here that we have used translational invariance to write the pairwise corrections as the interaction between a single ON cell at 0 and all other cells. Likewise, since the ON and OFF cells are identical up to a polarity flip, the $\mathcal{I}_2$ term includes both mosaics' self-interactions.

Clearly, only the last term in **Eq 25** is a function of ϕ, and we only need to consider the value ϕ that maximizes:

$$\mathcal{I}_{2'}(\phi) \overset{\text{def}}{=} N\sum_j \mathbb{E}_{\mathbf{x}}[a_0\bar{a}_j]\log\frac{\bar{A}_{2j}}{A_1^2}, \tag{26}$$

to determine whether alignment ($\phi = 0$) or anti-alignment ($\phi = 0.5$) is preferred. We know from **Eq 17** that

$$A_1 = 1 + \frac{g_0}{\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2\gamma^{-2}}, \tag{27}$$

and we can calculate the two-neuron determinant ratio $\bar{A}_{2j}$ as follows: First define matrices $\mathbf{P}$ and $\mathbf{Q}$ from **Eq 11** to write

$$\begin{aligned} A &= \frac{\det(\overbrace{\mathbf{G}\mathbf{W}^\mathsf{T}\mathbf{C}_x\mathbf{W}\mathbf{G}}^{\overset{\text{def}}{=}\mathbf{Q}} + \overbrace{\mathbf{G}\mathbf{W}^\mathsf{T}\mathbf{C}_{in}\mathbf{W}\mathbf{G} + \mathbf{C}_{out}}^{\overset{\text{def}}{=}\mathbf{P}})}{\det(\underbrace{\mathbf{G}\mathbf{W}^\mathsf{T}\mathbf{C}_{in}\mathbf{W}\mathbf{G} + \mathbf{C}_{out}}_{\overset{\text{def}}{=}\mathbf{P}})} \\ &= \frac{\det(\mathbf{Q}+\mathbf{P})}{\det(\mathbf{P})}, \end{aligned} \tag{28}$$

then substitute kernels $\mathbf{w}_0$ and $-\mathbf{w}_j$ (negative because of the opposite polarity):

$$\begin{aligned} \mathbf{P} &= \mathbf{G}\mathbf{W}^\mathsf{T}\mathbf{C}_{in}\mathbf{W}\mathbf{G} + \mathbf{C}_{out} \\ &= \gamma^2\sigma_{\text{in}}^2\begin{bmatrix}\mathbf{w}_0^\mathsf{T}\mathbf{w}_0 & -\mathbf{w}_0^\mathsf{T}\mathbf{w}_j \\ -\mathbf{w}_j^\mathsf{T}\mathbf{w}_0 & \mathbf{w}_j^\mathsf{T}\mathbf{w}_j\end{bmatrix} + \begin{bmatrix}\sigma_{\text{out}}^2 & 0 \\ 0 & \sigma_{\text{out}}^2\end{bmatrix} \\ &= \begin{bmatrix}\gamma^2\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2 & -\gamma^2\sigma_{\text{in}}^2 f_j \\ -\gamma^2\sigma_{\text{in}}^2 f_j & \gamma^2\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2\end{bmatrix} \end{aligned} \tag{29}$$

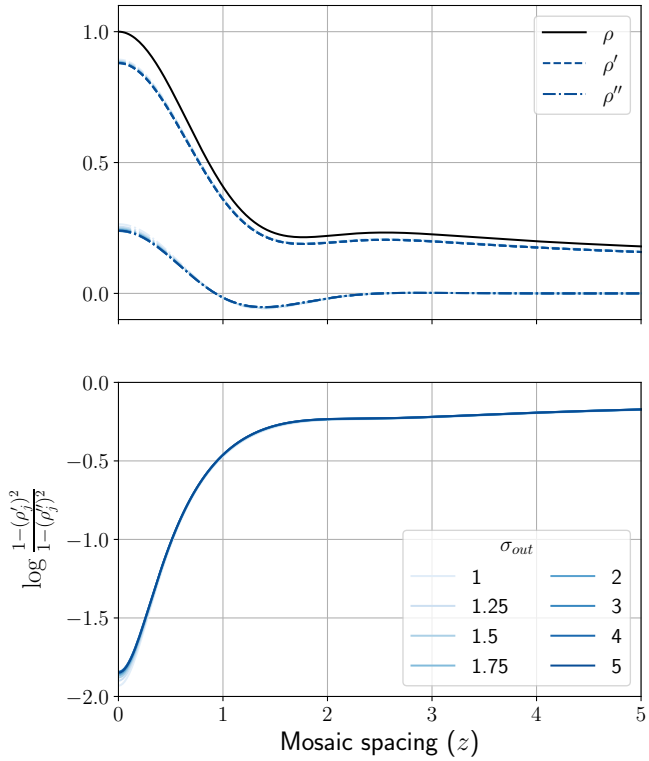

Supplementary Figure 8. Plots of the correlation-like quantities ρ, ρ', and ρ'' as defined in **Eqs 37** and **33** and the log ratio $\log\frac{\bar{A}_{2j}}{A_1^2} = \log\frac{1-(\rho')^2}{1-(\rho'')^2}$, with respect to color-coded σ_{out} values. These depend very weakly on σ_{out}, suggesting the $\mathbb{E}_{\mathbf{x}}[a_0\bar{a}_j]$ term in **Eq 26** plays a more crucial role in the phase transition.

and

$$\begin{aligned} \mathbf{Q} &= \mathbf{G}\mathbf{W}^\mathsf{T}\mathbf{C}_x\mathbf{W}\mathbf{G} \\ &= \gamma^2\begin{bmatrix}\mathbf{w}_0^\mathsf{T}\mathbf{C}_x\mathbf{w}_0 & -\mathbf{w}_0^\mathsf{T}\mathbf{C}_x\mathbf{w}_j \\ -\mathbf{w}_j^\mathsf{T}\mathbf{C}_x\mathbf{w}_0 & \mathbf{w}_j^\mathsf{T}\mathbf{C}_x\mathbf{w}_j\end{bmatrix} = \gamma^2\begin{bmatrix}g_0 & -g_j \\ -g_j & g_0\end{bmatrix}, \end{aligned} \tag{30}$$

where we define $g_j \overset{\text{def}}{=} g(|\bar{z}_j|)$, $f_j \overset{\text{def}}{=} f(|\bar{z}_j|)$, and thus $f_0 = 1$. We then have

$$\det\mathbf{P} = (\gamma^2\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2)^2\left[1 - \left(\frac{\gamma^2\sigma_{\text{in}}^2 f_j}{\gamma^2\sigma_{\text{in}}^2 + \sigma_{\text{out}}^2}\right)^2\right] \tag{31}$$

and

$$\begin{aligned} &\det(\mathbf{P}+\mathbf{Q}) \\ &= \det\begin{bmatrix}\gamma^2(g_0+\sigma_{\text{in}}^2) + \sigma_{\text{out}}^2 & -\gamma^2(g_j + \sigma_{\text{in}}^2 f_j) \\ -\gamma^2(g_j+\sigma_{\text{in}}^2 f_j) & \gamma^2(g_0+\sigma_{\text{in}}^2) + \sigma_{\text{out}}^2\end{bmatrix} \\ &= (\gamma^2(g_0+\sigma_{\text{in}}^2) + \sigma_{\text{out}}^2)^2\left[1 - \left(\frac{\gamma^2(g_j+\sigma_{\text{in}}^2 f_j)}{\gamma^2(g_0+\sigma_{\text{in}}^2) + \sigma_{\text{out}}^2}\right)^2\right]. \end{aligned} \tag{32}$$

Collecting these with the expression for A_1 in **Eq 27**,

$$\begin{aligned}\frac{\bar{A}_{2j}}{A_1^2} &= \frac{\dfrac{\det(\mathbf{Q}+\mathbf{P})}{\det(\mathbf{P})}}{\left(1+\dfrac{g_0}{\sigma_{\text{in}}^2+\sigma_{\text{out}}^2\gamma^{-2}}\right)^2} \\ &= \frac{1-\Big(\overbrace{\dfrac{g_j+\sigma_{\text{in}}^2 f_j}{g_0+\sigma_{\text{in}}^2+\sigma_{\text{out}}^2\gamma^{-2}}}^{\stackrel{\text{def}}{=}\rho_j'}\Big)^2}{1-\Big(\underbrace{\dfrac{\sigma_{\text{in}}^2 f_j}{\sigma_{\text{in}}^2+\sigma_{\text{out}}^2\gamma^{-2}}}_{\stackrel{\text{def}}{=}\rho_j''}\Big)^2} = \frac{1-(\rho_j')^2}{1-(\rho_j'')^2}\end{aligned} \tag{33}$$

Note that since $f_j \leq 1$ by definition and $g_j \leq g_0$ for most reasonable image distributions, both terms in brackets are less than 1. Also, as $j \to \infty$, we have $g_j, f_j \to 0$, so the term vanishes asymptotically. In fact, f_j vanishes very quickly (as soon as the RFs have negligible overlap), while g_j diminishes much more slowly, since $\mathbf{C}_x$ contains long-range correlations. Note that f_j and g_j are the only terms here that depend on ϕ. More importantly, this term only varies weakly with σ_{out} as shown in **Sup Fig 8**, suggesting it does not play a crucial role in the phase transition.

Now, as for the coactivation probability $\mathbb{E}_{\mathbf{x}}[a_0\bar{a}_j]$, we consider $y_0 = \mathbf{w}_0^{\mathsf{T}}(\mathbf{x}+\mathbf{n}_x)$ and $y_j = -\mathbf{w}_j^{\mathsf{T}}(\mathbf{x}+\mathbf{n}_x)$ whose joint distribution follows

$$(y_0, y_j) \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} g_0+\sigma_{\text{in}}^2 & -g_j-\sigma_{\text{in}}^2 f_j \\ -g_j-\sigma_{\text{in}}^2 f_j & g_0+\sigma_{\text{in}}^2 \end{bmatrix}\right), \tag{34}$$

which means that $\mathbb{E}_{\mathbf{x}}[a_0\bar{a}_j]$ can be calculated via a Gaussian integral over the region $y_0 > \theta$ and $y_j > \theta$:

$$\begin{aligned}\mathbb{E}[a_0\bar{a}_j] &= \int_\theta^\infty \int_\theta^\infty p(y_0, y_j)\, dy_0\, dy_j \\ &= \int_\theta^\infty \frac{1}{\sqrt{2\pi}} e^{-\frac{y^2}{2}} \left(1-\Phi\left(\frac{\tilde{\theta}-\rho_j y}{\sqrt{1-\rho_j^2}}\right)\right) dy \\ &\stackrel{\text{def}}{=} p_2(\tilde{\theta}, \rho_j)\end{aligned} \tag{35}$$

where $\tilde{\theta}$ and ρ_j are defined as

$$\tilde{\theta} = \theta/\sqrt{g_0+\sigma_{\text{in}}^2} \tag{36}$$

$$\rho_j \stackrel{\text{def}}{=} \frac{g_j+\sigma_{\text{in}}^2 f_j}{g_0+\sigma_{\text{in}}^2}. \tag{37}$$

Putting all this together, we can write $\mathcal{I}_{2'}$ as

$$\begin{aligned}\mathcal{I}_{2'}(\phi) &= N\sum_j p_2(\tilde{\theta}, \rho_j) \log \frac{1-(\rho_j')^2}{1-(\rho_j'')^2} \\ &= \sum_{i,j} p_2(d_{ij}) \log R(d_{ij}),\end{aligned} \tag{38}$$

where again, the dependence on ϕ is solely through $f_j = f(|j+\phi|)$ and $g_j = g(|j+\phi|)$ and we have used translation invariance to rewrite these arguments as d_{ij}, the distance between ON-OFF pairs. Finally, since $f_j \leq 1$ by definition, $\rho_j > \rho_j'$ always, and $\rho_j' > \rho_j'' \iff \frac{g_j}{g_0} > \rho_j''$ (which holds empirically with our image distribution and optimized kernels), we have $0 < \frac{1-(\rho_j')^2}{1-(\rho_j'')^2} < 1$ and thus $\mathcal{I}_{2'} \leq 0$.

For the results plotted in **Fig 5A**, we numerically integrated **Eq 35** to calculate p_2 and used kernels and covariance parameters taken from a trained one-shape model, shown in **Sup Fig 7**, to calculate $\mathbf{C}_x$, f_j, and g_j. To make these quantities comparable as σ_{out} increases (and thus p_2 decreases), **Fig 5A** plots the normalized contribution $\mathcal{I}_{2'}/(Np_1^2)$.

C. Redundancy in aligned and independent ON-OFF pairs

We have argued in **Fig 4D-E** that redundancy is reduced by aligning a single OFF cell (in the case of low θ) and anti-aligning multiple OFF neighbors (for high θ) inside a "reduced redundancy" region of small p_2 around each ON cell. However, this characterization may be objected to on the grounds that, in fact, aligned response fields are highly redundant, since the activation of one by a stimulus excludes activation of the other.

In a limited sense, this is true. Consider replacing our monotonically increasing nonlinearity h with a simple threshold nonlinearity, such that neurons fire at a constant rate when active, similar to Gjorgjieva et al. (2014). In this case, we can calculate for a single neuron with $\theta = 0$

$$H(\mathbf{X}|\mathbf{R}) = -p(0)\log p(0) - p(1)\log p(1) = \log 2 = 1 \text{ bit} \tag{39}$$

since the neuron is expected to be active half the time. Clearly, for two independent cells (e.g., those with RF centers separated by a distance much larger than the image correlation scale), we then expect 2 bits of conditional entropy. However, for an aligned ON-OFF pair, we would have

$$H(\mathbf{X}|\mathbf{R}) = -p(0,1)\log p(0,1) - p(1,0)\log p(1,0) = \log 2, \tag{40}$$

just as before. Thus, when each neuron has a capacity of one bit, two aligned RFs are completely redundant.

However, the assumption of a graded firing rate conveys additional information about the stimulus beyond its sign, and this information is not fully shared even by aligned RFs. For instance, in the case of a single neuron, we have from **Eqs 20** and **27**

$$H(\mathbf{X}|\mathbf{R}) \approx \frac{1}{2}\log\left(1+\frac{g_0}{\sigma_{\text{in}}^2+\sigma_{\text{out}}^2\frac{g_0}{2\pi}}\right) \approx 0.5\log 5.7 \tag{41}$$

for $\theta = 0$, $\sigma_{\text{in}} = 0.4$, $\sigma_{\text{out}} = 1$, $g_0 \approx 3.8$. However, this is *differential* entropy. To relate this to the answer in the 1-bit case, we need to consider a quantization of the signal. If we assume that the signal can be quantized using 1 bit for its sign and n bits for its magnitude, we have, for n sufficiently large (Cover 1999)

$$H(\mathbf{X}^{\Delta}|\mathbf{R}) \approx \tilde{H}(\mathbf{X}|\mathbf{R}) + (n+1)\log 2 \tag{42}$$

where $\mathbf{X}^{\Delta}$ is the quantized variable, and we denote the differential entropy by $\tilde{H}$.

From this result, it is clear that, while far-separated RFs encode precisely $2H(\mathbf{X}^{\Delta}|\mathbf{R})$, perfectly aligned RFs encode only $2H(\mathbf{X}^{\Delta}|\mathbf{R}) - \log 2$, where the correction is due to the shared sign bit. Thus, while it is correct that aligned RFs do exhibit some redundancy (1 bit for $\theta = 0$, less as θ increases), this is typically much smaller than the information about stimulus magnitude. More importantly, when considering how to align ON and OFF cells, informational losses due to shared sign bits are dwarfed by the size of $\mathcal{I}_{2'}$ for neighbors located outside the gray zone in **Fig 4A**.