Dysregulation of neural excitability replicates physiological and functional changes in aging visual cortex

Seth Talyansky[1,2,*], Braden A. W. Brinkman[2,†]

[1]Catlin Gabel School, Portland, OR, USA

[2]Department of Neurobiology and Behavior, Stony Brook University, Stony Brook, NY, 11794, USA

† Corresponding author, braden.brinkman@stonybrook.edu

* Now at Stanford University.

January 6, 2020

Abstract

The mammalian visual system has been the focus of countless experimental and theoretical studies designed to elucidate principles of sensory coding. Most theoretical work has focused on networks intended to reflect developing or mature neural circuitry, in both health and disease. Few computational studies have attempted to model changes that occur in neural circuitry as an organism ages non-pathologically. In this work we contribute to closing this gap, studying how physiological changes correlated with advanced age impact the computational performance of a spiking network model of primary visual cortex (V1). Our results demonstrate that deterioration of homeostatic regulation of excitatory firing, coupled with long-term synaptic plasticity, is a sufficient mechanism to reproduce features of observed physiological and functional changes in neural activity data, specifically declines in inhibition and in selectivity to oriented stimuli. This suggests a potential causality between dysregulation of neuron firing and age-induced changes in brain physiology and performance. While this does not rule out deeper underlying causes or other mechanisms that could give rise to these changes, our approach opens new avenues for exploring these underlying mechanisms in greater depth and making predictions for future experiments.

1 Introduction

Despite affecting approximately 100% of the human population, healthy aging, and in particular its impact on neurological performance, has been the focus of few theoretical and computational studies in neuroscience,

compared to diseases such as Alzheimer's [1–12] or Parkinson's [13–20], which often manifest in advanced age in humans [21]. Only recently has theoretical and computational work on aging in non-pathological networks begun to emerge [22]. Because advanced age is one of the most important risk factors for developing such neurological disorders [21,23], to fully understand the progression of these diseases we ought to have a baseline understanding of how the brain's circuitry changes under healthy aging, both in terms of physiological properties and functional performance. This would help dissociate disease-related changes from those caused during normal aging, and thereby allow researchers to focus their attention on treating potential causes of the disease progression not directly related to aging. On the other hand, understanding how the healthy brain ages may enable us to treat declines in performance caused solely by aging, in both healthy subjects and those with neurological disorders or diseases.

In this work we seek to advance our understanding of potential mechanisms and consequences of age-induced changes in physiology and performance in visual cortex. We do so by adapting a previously-developed spiking network model of V1 activity, E-I Net [24]. The network structure of E-I Net is learned by training it on pixel images of natural scenes, to which it develops Gabor-like receptive fields. The synaptic and homeostatic learning rules are local, and therefore biologically possible. The training period of E-I Net's predecessor, SAILnet [25], has been used to model the experiential development of visual cortex in growing mammals [26]. Here, we take a similar approach to the aging process, modifying the synaptic and homeostatic learning rules to reflect changes that occur after the network has reached a mature steady state, which subsequently drive age-induced changes in the network. The changes we implement are motivated by experimental work that has found that senescent brain tissue shows increased excitability [27], decreased GABAergic inhibition [28], and decreased selectivity to the orientation of grating stimuli [27]. In particular, we show that dysregulation of the network's target firing rate in the learning rules not only leads to increased excitability, but also decreases in synaptic strength and declines in orientation selectivity, mimicking the experimentally observed changes. Although there may be other mechanisms that could also replicate these experimental findings, our model makes several predictions for other yet-to-be-observed changes in physiology and functional performance that could be used to further test our model and its interpretation.

2 Results

Network model

We use E-I Net [24], a spiking network model, as our model of visual cortex. As shown schematically in Fig. 1, E-I Net comprises a population of excitatory and inhibitory leaky integrate-and-fire spiking neurons, which receive visual input in the form of patches from pixel images of natural scenes. E-I Net learns the lateral weights W_{ij} between neurons, inputs weights Q_{ik} from input pixels to neurons, and the firing thresholds θ_i, over repeated exposure to patches of visual pixel input X_k. These parameters develop according to the

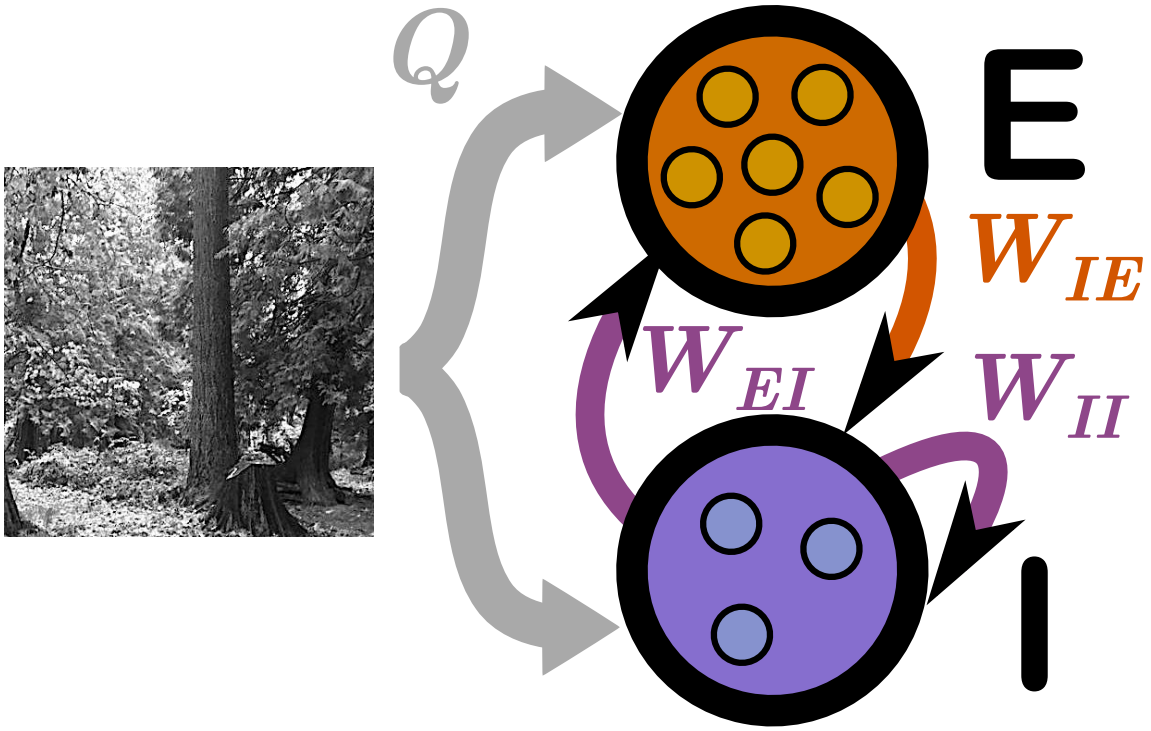

Figure 1: **Schematic of the network model**: We use E-I Net [24] as our model of visual cortex. E-I Net comprises a network of leaky integrate-and-fire neurons that receive visual input in the form of patches from pixel images, communicated to the neurons through input weights Q. Each neuron belongs to either the excitatory (E) or inhibitory (I) population. Neurons are synaptically connected by lateral weights W that obey Dale's law, such that excitatory-to-inhibitory weights W_{IE} are positive, while inhibitory-to-excitatory (W_{EI}) and inhibitory-to-inhibitory (W_{II}) weights are negative. We assume excitatory-to-excitatory connections may be omitted. The input weights Q, lateral weights W, and firing thresholds of each neuron θ, are learned by training the network on natural pixel images.

biologically-plausible learning rules

$$\Delta Q_{ik} \propto n_i X_k - n_i^2 Q_{ik} \tag{1}$$

$$\Delta W_{ij} \propto n_i n_j - \langle n_i \rangle \langle n_j \rangle (1 + W_{ij}) \tag{2}$$

$$\Delta \theta_i \propto n_i - p; \tag{3}$$

where n_i is the spike count of neuron i in response to the presented image patch over a duration T and $p = p_E$ or p_I is the desired mean number of spikes emitted over this window, with separate rates for excitatory neurons (p_E) and inhibitory neurons (p_I). The quantities $\langle n_i \rangle$ are moving-averages over time, weighted by an exponential decay factor. As in the original work using E-I Net [24] and its predecessor SAILnet [25], these learning rules were designed to enforce linear decoding of the image patches from the spike counts (giving Eq. (1)), minimal correlation in spike counts (giving Eq. (2)), and sparseness in firing (giving Eq. (3)) due to the fact that $p \ll 1$. Notably, these rules are local—changes to a synapse only involve parameters local to that synapse.

Training E-I Net on image patches from natural scenes results in a steady-state network structure in which the pixel-space arrangement of input weights Q resemble Gabor-like receptive fields typical of networks trained on natural scenes [25]. The learned lateral weights are organized such that excitatory neurons with similar RFs effectively inhibit one another [25,29]. For our simulations reported in this work, we follow [24] in using $N_E = 400$ excitatory neurons and $N_I = 49$ inhibitory neurons. We similarly do not include excitatory-excitatory lateral connections (i.e., $W_{EE} = 0$). See this reference for details of this choice and other aspects of training the model.

To study the effect of aging, we first allow our network model to reach a steady state, taking 30 training loops, and then modify the learning rates to mimic the aging process. This is similar to work which interprets the initial training phases of neural networks as a model for development of the visual system or other brain areas [26], and is an example of recently-advocated strategies for using machine learning methods to better understand neural systems [30]. In order to study how the performance of the network changes with age we pause the training every 5 loops to test the neural responses on a set of images of oriented gratings.

To guide our change to the training procedure, we considered several physiological changes observed in the literature. We identified two candidate changes that we hypothesized could be related: a decrease in GABA in senescent brain tissue [28], which we interpret in the context of our model as a decrease in inhibitory coupling strength, and an increase in excitability of neurons in elderly animals [27]. While one might anticipate that a drop in inhibition would result in increased excitability of a network, in the context of network development, increasing the excitatory target spike rate (ETSR) might also cause the network to reorganize and decrease inhibitory synaptic strengths in order to allow for increased excitability. In early versions of this work we tested both ideas, and ultimately chose to focus on the latter implementation, as it naturally suited making a simple and interpretable change to the learning rules of the network model, changing the ETSR parameter p_E, whereas producing a network with weaker inhibitory connections either

demands making by-hand changes to the weights and shutting off the lateral weight learning so that these changes are not counteracted by the normal learning process, or would require introducing new terms to the learning rules of inhibitory weights that may not be interpretable in terms of previously established learning rules such as the Hebbian-Oja or Földiák learning rules [24, 25].

Therefore, to mimic aging in our model, we modify the excitatory target spike rate (ETSR) p_E by allowing it to begin increasing with further training loops $p_E \to p_E(t)$ beyond the initial maturation age of 30 training loops. Prior to age 30, both p_E and p_I are kept constant at a steady value of 0.01. At subsequent ages, we increase p_E by 0.01 every 5 training loops. This may be interpreted as potentially arising from a breakdown in a homeostatic mechanism that regulates this target firing rate, which we do not model explicitly (see Discussion). As the ETSR increases, the network parameters will begin to depart from the previously stabilized values, and we can track their evolution with training time, or "age." Moreover, at each of these ages we can test the functional performance of the network in response to different tasks. We choose to focus on selectivity to the orientation of edges in the visual scene, tested using oriented grating stimuli. This task mimics the experimental work of [27], to test whether we observe similar changes in selectivity. Finally, because our network model gives us full access to the network parameters at every stage of "aging," we test how much features of the different parameter sets contribute to the degradation of functional performance.

Physiological parameter changes

In Fig. 2, we show the distributions of the learned physiological parameters in "mature youth" (30 training loops) compared to "old age" (80 training loops). We find the range of magnitudes of the synaptic and input weights decreases with age, while the firing thresholds expand by orders of magnitude. (Note that the network model is formulated in dimensionless units, so the parameter ranges will not correspond directly to dimensionful quantities measured in experiments).

In addition to changes in magnitude, the input weights also undergo a change in their spatial organization in pixel space—i.e., the receptive field of the neuron is not preserved with aging. Fig. 3A shows the receptive field of a particular neuron with a Gabor-like receptive field (RF) at age 30 (top) and random-looking RF at age 80 (bottom). We can quantify just how different young and old receptive fields become by computing the angular separation φ_i between the vectorized young and old receptive fields Q^{young} and Q^{old} for each neuron i:

$$\varphi_i = \cos^{-1}\left(\frac{\sum_k Q_{ik}^{\text{young}} Q_{ik}^{\text{old}}}{\sqrt{\sum_k (Q_{ik}^{\text{young}})^2 \sum_k (Q_{ik}^{\text{old}})^2}}\right). \tag{4}$$

This angle is 0 if the old RF is identical to the young and $\pi/2$ if the young and old RFs are orthogonal. We find that the mean angle steadily increases with age, to the point that in the oldest age of our network the mean angle between a neuron's young and old receptive fields tends to $\pi/2$, i.e., on average the old-age RFs of a population of neurons becomes orthogonal to their RFs in youth. However, as seen in Fig. 3, the spread of angles around the mean also increases with age, indicating that while on average the RFs are orthogonal

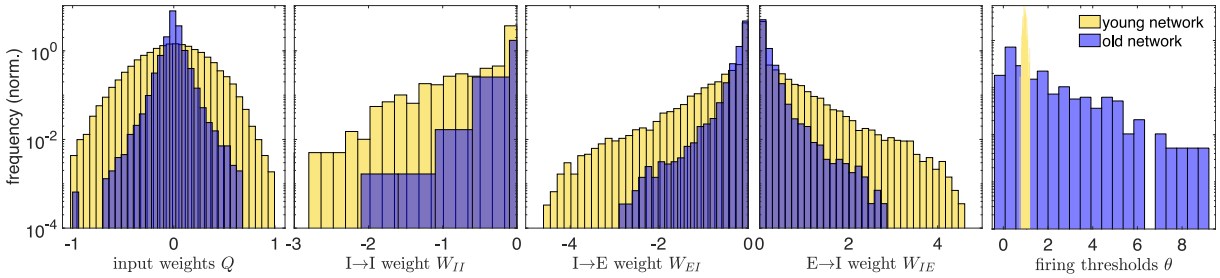

Figure 2: **Histograms of physiological network parameters in youth vs. old age**: Histograms of the magnitudes of network parameters learned during training on natural images: the input weights Q, lateral weights W, and firing thresholds θ. Yellow histograms correspond to the distributions of these parameters in the mature young network, the steady state of network development (reached by 30 training loops). After this age, the excitatory target firing rate begins to increase, "aging" the network. The blue histograms are the distributions of network parameters after 80 training loops. Aging tends to decrease input and lateral weights, while it increases the distribution of firing thresholds. A decrease in inhibitory weights is consistent with experimental work of [28]. All quantities are expressed in arbitrary dimensionless units.

in old age, there is significant variance in the mean distance between young and old RFs.

The fact that an individual neuron's RF changes drastically with age does not necessarily indicate the performance of the network as a whole changes—perhaps RFs are redistributed throughout the network, preserving the overall distribution of RFs but not each individual neuron's RF. To show that this is not the case, we categorized the receptive fields of excitatory neurons into either "Gabor-like," showing characteristics of Gabor edge detectors, and "not-Gabor," and quantified how membership in these categories changed with age in Table 1. As some RFs categorized as not-Gabor in youth became Gabor-like with age, we also recorded transitions between types (bottom table in Table 1). We assessed Gabor-ness by fitting a spatial Gabor profile to each neuron's receptive field and reject as Gabor if the goodness-of-fit relative to the RF magnitude is too large; see Methods and [25]. We used a conservative goodness-of-fit criterion; nonetheless, under this criterion close to a third of the excitatory RFs were classified as Gabor-like, while this fraction dropped to nearly 1-in-20 in our oldest network, demonstrating that the aged networks undergo a net loss in neurons with Gabor-like RFs. Given the edge-detecting properties of Gabor-like receptive fields, we thus anticipate this loss plays a prominent role in the degradation of the sensitivity of the network to oriented grating stimuli, discussed in the next section.

Performance changes

As a first basic check of network performance, we analyzed the distribution of spike counts over image presentations. The single change we made to the network—increasing the ETSR p_E with training—is intended to produce increased excitability with age, as observed in experiments [27]. Indeed, as shown in Fig. 4A, we find that not only do the mean excitatory spike counts increase with age, but the spread of the

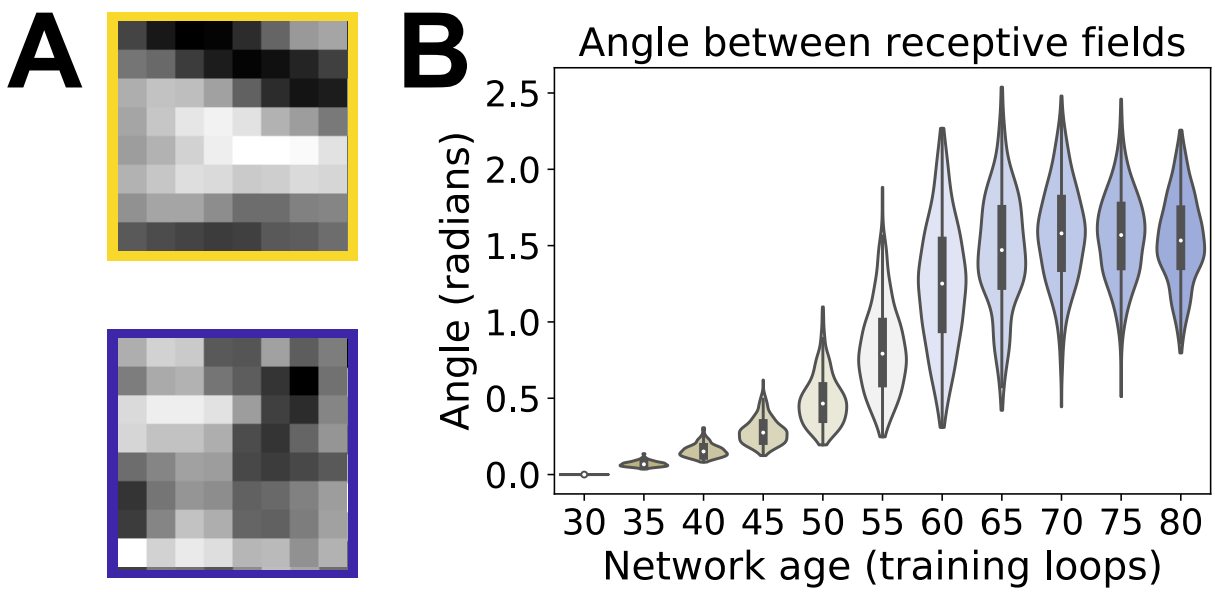

Figure 3: **Angle between receptive fields in the input weight parameter space**: **A**: A particular neuron's receptive field (RF) in youth (top, yellow border) versus old age (bottom, blue border), demonstrating that features of a neuron's RF may not be preserved in our model's aging process. **B**: Violin plot of the distribution of angles between each individual excitatory neuron's young and older vectorized receptive fields (RFs) (Eq. 4). As our network ages a neuron's RF becomes increasingly orthogonal to its mature young RF, on average, indicated by the median of the RF distributions tending to approximately $\pi/2$.

Table 1: **Quantification of Gabor-like receptive fields**: We quantify whether each excitatory neuron in our network model has Gabor-like receptive fields or not, in both our mature young network (30 loops) and old age (80 loops). We assess Gabor-ness by fitting a spatial Gabor profile to each neuron's receptive field and reject as Gabor if the goodness-of-fit relative to the RF magnitude is too large; see Methods and [25]

Gabor-like in youth (30 loops)	Gabor-like in age (80 loops)
123/400	26/400

	Gabor-like in age	not-Gabor in age
Gabor-like in youth	6/123	117/123
not-Gabor in youth	20/277	257/277

distribution does as well, confirming that our network becomes more excitable with age. We did not increase the target firing rate p_I of the inhibitory neurons, and as a result neither the mean spike counts nor variance of the spike count distribution increase with age (Fig. 4B).

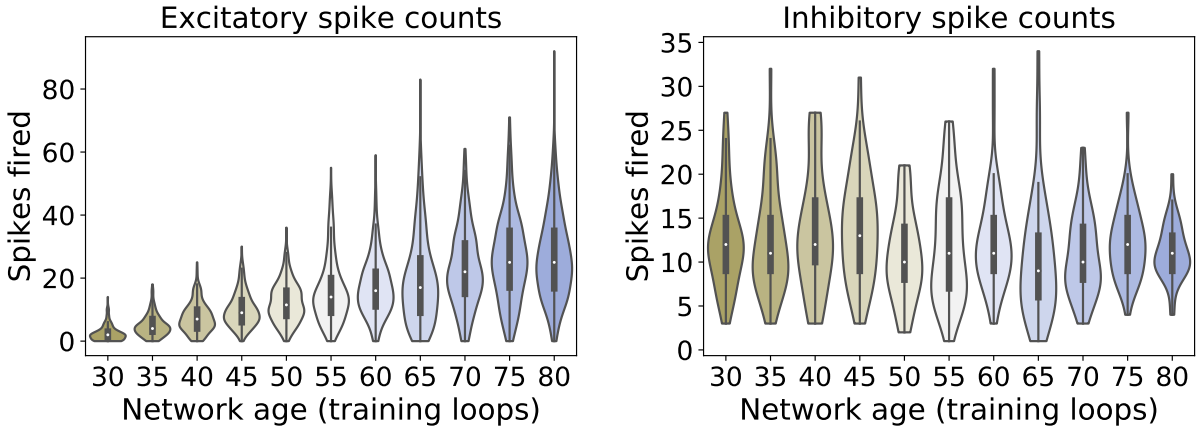

Figure 4: **Excitability of the network with age**: **Left**: Violin plot of the excitatory spike count distributions with age. Markers indicate the median. Both median and variance of the distributions increase with age, indicating increased excitability of the network with age, as expected. **Right**: Violin plot of the inhibitory spike count distributions. The inhibitory target rate does not increase with age, and thus the distributions do not change systematically during aging.

To test functional performance, every 5 training loops we paused training and tested the network on patches drawn from pixel images of oriented gratings of spatial frequency 0.1 cycles per pixel, such that one 8×8 patch contains approximately 1 cycle. We quantify a neuron's selectivity to each orientation using the same orientation selectivity index employed by Hua *et al.* [27]. As seen in Fig. 5, the cumulative distribution plots of the orientation indices (the fraction of cells with indices less than a particular value) qualitatively agree with the experimental observations of [27]: selectivity decreases with age, demonstrated by the shift of the cumulative plots toward lower orientation indices. However, our model results also suggest this process

may not be a straightforward monotonic degradation of selectivity with age. In the first several epochs of aging our network we find that the mean selectivity of the network does not actually change appreciably, but rather that changes in weakly- and strongly-selective neurons initially balance out (ages 30-50, Fig. 5C), and only later in age does the network begin to lose strongly-selective neurons, a prediction that could be tested experimentally in future studies.

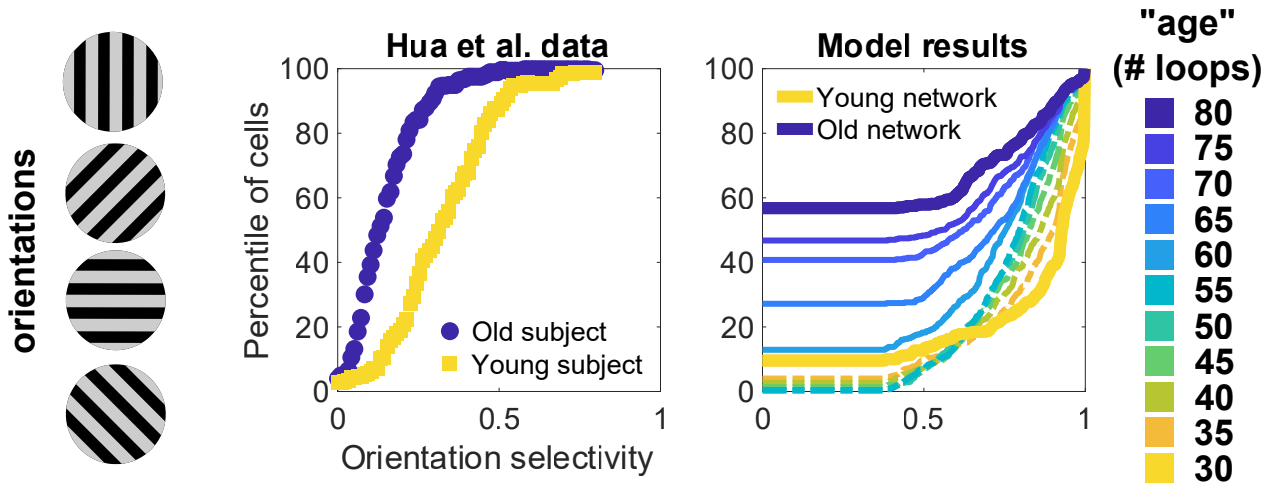

Figure 5: **Selectivity of cells to oriented gratings**: **Left**: schematic of oriented gratings presented to the network. **Middle**: Cumulative distribution plot of orientation selectivity indices of real neurons from Hua *et al.* [27] in young (yellow) and elderly (blue) animals. **Right**: Cumulative distributions of orientation selectivities at several different "ages" of our network model. Orientation selectivity decreases in old age in both. Additionally, our network model predicts that mean selectivity initially does not change much in the early stages of aging.

The quantitative differences between our simulation results and the data may be traced to some of the simplified features of the network model we use. For example, the selectivities of neurons in our model network are much more polarized, having many neurons with vanishing orientation selectivity index (due to not firing to any stimulus in our testing phase) and many with nearly maximal selectivity, which also leads to our cumulative plots having the opposite convexity to that of the data. These more polarized orientation selectivities in our simulated network are produced in part because the learning rules of E-I Net promote sparse firing yet demand decodability of inputs, resulting in highly selective neurons that are active essentially only when the feature they are selective for is presented, and are silent otherwise. Combined with the fact that the network is deterministic (no noisy current inputs), the resulting activity levels tend to be highly polarized. Intermediate selectivities are likely generated by lateral coupling between neurons with similar but distinct receptive fields [29], leading to imperfect selectivity. Another factor that allows for these more polarized selectivities is the fact that every neuron in the network responds to the same visual patch at a given time. Potentially, a very large network in which different clusters of neurons respond to many inputs

from different locations in the visual field would generate enough recurrent activity for recorded neurons to appear spontaneously active, which might smooth out the distributions of orientation selectivity. We leave training and simulating aging in such a network for future work.

Numerical experiments to test contribution of different physiological parameters to declines in performance

Given the variety of changes our model network undergoes during aging, can we distinguish which physiological changes contribute most to the observed changes in performance? We certainly expect the degradation of the Gabor-like structure of the receptive fields to play a strong role, at least later in the aging process when the RFs are no longer Gabor-like, as in the absence of these edge-detector filters the network we naturally expect selectivity to drop. However, before the RF structure has completely degraded, we may also expect lateral weights to play a role, as several theoretical and experimental studies suggest that neurons with similar response properties will effectively mutually inhibit one another and thereby sharpen the selectivity of their responses [31–37]. Thus, a loss of inhibition could also contribute significantly to decreased orientation selectivity at intermediate old ages. To probe which parameter set dominates the contribution to the loss of orientation selectivity as a function of age, we performed several numerical experiments on our networks.

In our first set of simulation experiments, we created several young-old hybrid networks by mixing together different combinations of the young (age 30) and old (age 80) learned parameters and tested these networks on our oriented grating stimuli to measure the resulting orientation selectivity. As shown in Fig. 6, the mean selectivity appears impacted most when input weights are old and thresholds are young. This makes sense: the magnitude of the old-age input weights is smaller than in youth, and therefore the current inputs to each model $\sum_k Q_{ik}X_k$ are comparatively smaller, whereas the young threshold values are adapted to correspondingly higher current inputs and are therefore typically set higher in order to achieve the target spike rates during development ($p = 0.01$). Thus, typical current inputs drive the neurons to fire less. We see that replacing the young thresholds with the old-age thresholds improves the mean selectivity back to the pure old-age values. Moreover, we see that replacing the young lateral weights with old lateral weights does not appear to have a substantial impact on the mean selectivities at all, as simulations differing only in young/old lateral weights have nearly identical selectivities. Taken all together, this suggests that the input weights are by far the primary contributors to the decline in orientation selectivity, and any sharpening of selectivity that lateral weights might contribute is a second order effect.

This conclusion is further supported by tests in which we age our network while shutting off the learning rules for one or more of the physiological parameters, freezing these parameters to their young values. We perform three such tests: i) freeze the input weight matrices Q, ii) freeze the lateral weight matrices W, and iii) freeze both Q and W. Note that we do not freeze the thresholds to their young values because our network aging process is driven by the increasing excitatory target spike rate p_E, which only appears in the threshold learning rule. Freezing this learning rule would therefore preclude aging in our network.

The results of these tests appear in Fig. 7, in which we plot the average selectivity as a function of age

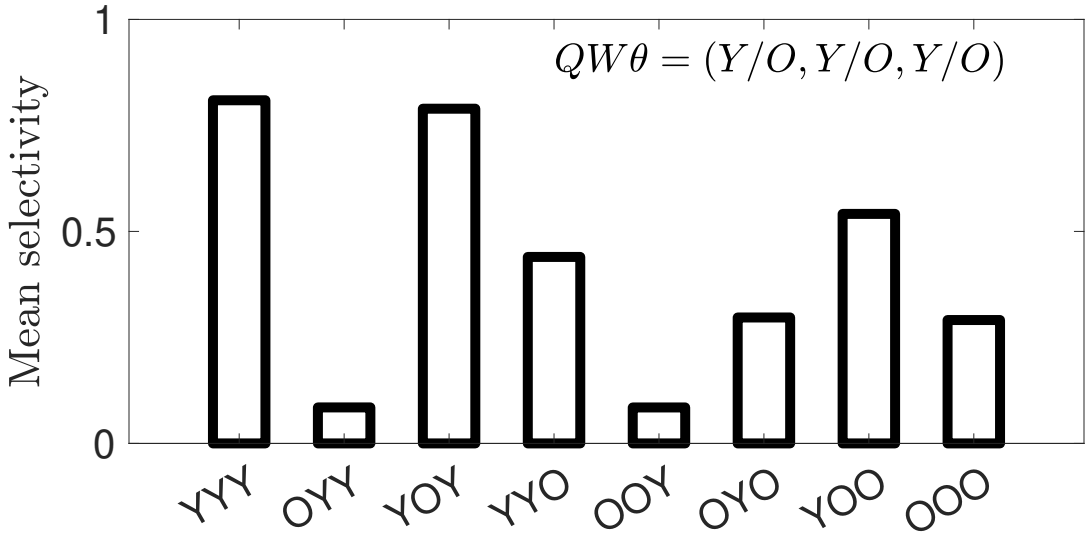

Figure 6: **Mean selectivity in networks with swapped parameter sets**: The mean orientation selectivity of neurons in network simulations in which we mix-and-match the parameter sets of young (30 loops) and old (80 loops) networks. Horizontal axis labels correspond to young (Y) or old (O) to input weights (Q), lateral weights (W), and thresholds (θ). For example, the "OOY" result corresponds to a simulation using the old-age input weights and lateral weights but the young thresholds.

for each case. We see that both cases in which the input weights Q are frozen retain a high mean orientation selectivity across ages, comparable to the mean selectivity in youth. When only the lateral weights are frozen and the input weights are allowed to adapt to the increasing ETSR, the mean selectivity maintains its youthful value for most of the aging process before suddenly dropping at the oldest ages we simulated, though not quite as low as the case in which both input and lateral weights are reorganized by ongoing plasticity. Taken together with the results of Fig. 6, these results support the conclusion that changes in input weights during aging are more important than lateral weights for maintaining the functional performance of the network. However, the fact that it is only in advanced age that selectivity falls off when lateral weights are frozen suggests that it is the coordinated reorganization of both input weights and lateral weights that leads to large degradation in functional performance.

These tests so far have only demonstrated that input weights changes are the dominant contributor to the decline in selectivity, but do not disambiguate the contribution of changes in input weight magnitude from changes in the spatial organization (the receptive field structure). As observed in Figs. 2 and 3, we know both the receptive field structure and the distribution of input weight magnitudes are altered during our aging process. With our network model we can separate the effects of input weight magnitude versus structure to test whether one of these features contributes more to the changes in neural selectivity. We might expect, for instance, that the loss of Gabor receptive field structure is primarily responsible for the loss in orientation selectivity. We test sensitivity to input weight structure and strength in two ways: i) we preserved each

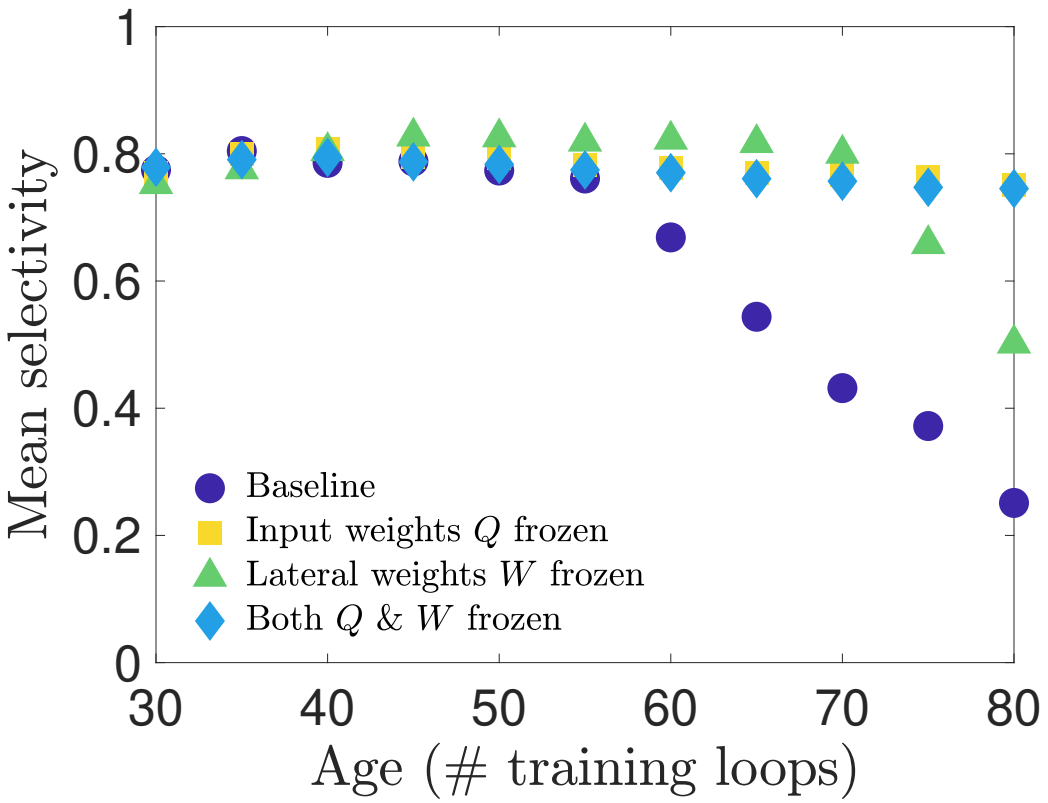

Figure 7: **Selectivity of cells when connection weights are frozen during training**: We test how network parameters contribute to orientation selectivity by shutting off learning of the input weights or lateral weights during the aging process. When input weights are frozen to their young values, the selectivity is barely impacted. However, when lateral weights are frozen to their young values, the selectivity remains high for most ages until it suddenly begins to drop off near the oldest ages in our simulations. When both Q and W are frozen the weights remain near the young selectivity values.

individual neuron's young receptive field structure while remapping the input weight magnitudes such that the distribution matches that of the old-age input weights; ii) we preserved each individual neuron's old receptive field structure while remapping the input weight magnitudes to match the young input weight distribution. We find, as seen in Fig. 8, that restoring the young RF structure gives an overall boost to the mean orientation selectivity at all ages, as might have been expected. However, we also see that at the oldest ages (70 to 80 training loops) restoring the young magnitudes yields a greater improvement in selectivity than that obtained by restoring young structure.

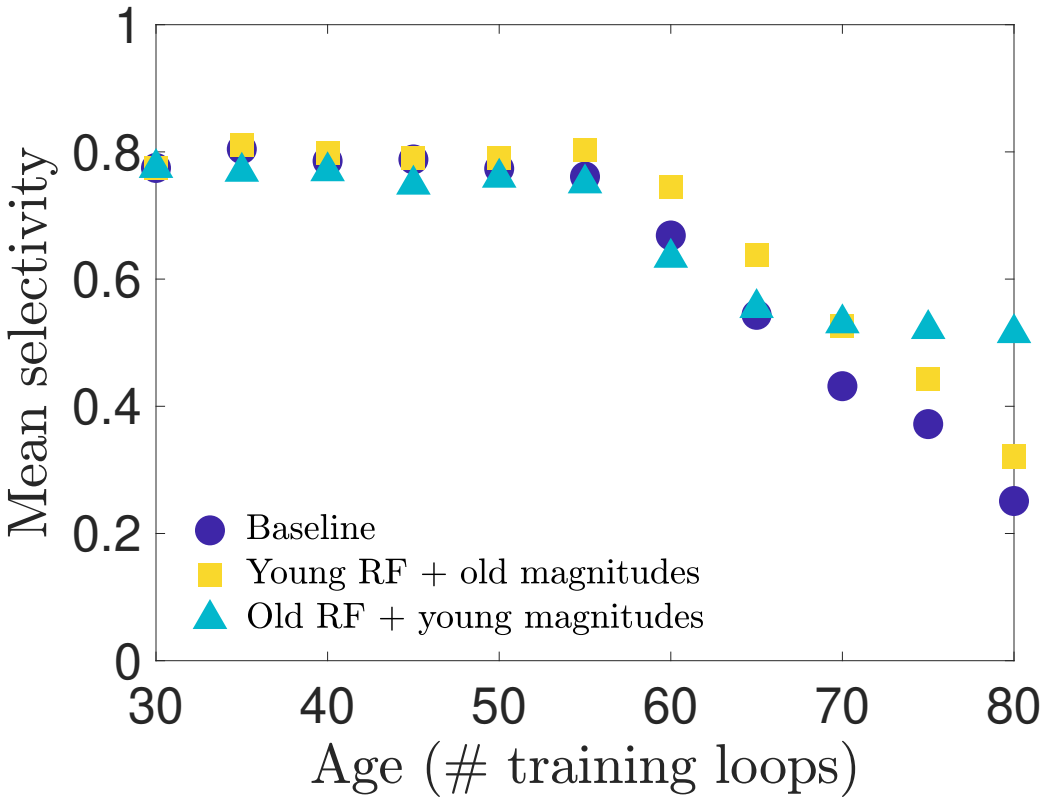

Figure 8: **Selectivity of cells with remapped receptive fields**: To disambiguate the impact of aging receptive field (RF) magnitudes versus spatial structure, we remap the magnitudes of each neuron's RF to create two types RFs: RFs with the original spatial structure they had in the young network but a magnitude distribution matching the old-age distribution, and RFs with the old-age spatial structure but young magnitude distribution. (See Methods for details). We create these hybrid RFs for each age and simulate the network model to measure the orientation selectivity of the population. We find that the young RF spatial structure improves selectivity in the early stages of aging, out-performing the baseline network (no remapping) as aging progresses, but the old RF spatial structure with young RF magnitudes maintains an overall higher level of selectivity in the latest stage of aging simulated.

3 Discussion

We have demonstrated that relatively simple changes to the homeostatic dynamics of a network can reorganize the physiological properties of a network in ways that qualitatively replicate experimentally observed changes in network properties and performance in old age: increased excitability (Fig. 4), weakened inhibitory weights (Fig. 2), and declines in orientation selectivity (Fig. 5). We also find decreases in excitatory weights (Fig. 2), despite the fact that overall excitation increases, potentially due to the drastic broadening of the firing threshold distribution. Moreover, with access to all our model parameters we can show that declines in selectivity are primarily due to the changes in input weight structure and strength (the receptive fields), rather than changes in lateral inhibition (Figs. 6, 7, and 8).

Because these physiological changes develop as the consequence of a modification to the network learning rules—interpreted here as a breakdown in the homeostatic mechanisms regulating target firing—our model allows us to test the potential causal mechanisms during aging, in contrast to an approach that simply alters network parameters to test the effects on function. Thus, this work suggests that dysregulation of excitability in a network may be one of the causal mechanisms that leads to physiological and functional changes in networks and their dynamics as an organism ages. Below, we discuss the comparison to the experimental results in detail, along with predictions for future experiments that could lend more evidence toward or against this hypothesized mechanism. To that end, we also discuss the limitations of the model—what features we expect to be reasonably realistic, albeit simplified, and which features of experimental observations we were not able to replicate with the current model.

Comparison to experimental results and implications for future experiments

Our investigation was motivated by several experimental results: i) neural excitability increases with age [27], ii) the proportion of GABAergic neurons in senescent brain tissue decreases with age (but total neuron density does not) [28], and iii) aged neurons exhibit decreased orientation and direction selectivity.

It is not clear *a priori* what the causal relationships are between these observations; most likely, there is an underlying common cause, rather than several (approximately) independent causes leading to these changes. Intuitively, we might expect that decreases in GABAergic inhibition could lead to increased excitability. In an early version of this work, we indeed explored this hypothesis by training the network model to maturity and then turning off learning and decreasing the magnitude of inhibitory connections in the network by a global fraction. We were able to reproduce loss in performance only when we added a tonic current to drive the neurons, and only for particular ranges of this tonic current and global reduction in the inhibitory weights. However, as the changes to the inhibitory weights in this early version of the model were implemented by hand, and not by changing the learning rules, they provided no plausible mechanisms for explaining how the network aged, and as learning was turned off in this approach it did not allow for the possibility that other network parameters might adjust to compensate.

We ultimately developed the hypothesis pursued in this work, that the causal relationship could be that

a breakdown in homeostatic regulation of excitability, coupled with long-term synaptic plasticity, would lead to changes in inhibition and other physiological features, as well as changes in functional performance. Therefore, the fundamental change we made to the healthy mature network was intended to produce increased excitability in it by allowing the excitatory target spike rate (ETSR) to steadily increase as training progressed past the phase at which the network parameters had reached a steady state. All other results follow from the coupling between this change and the learning dynamics of the model parameters. As expected, the excitability of excitatory neurons indeed increased with network age, as demonstrated by the increasing mean firing rate and variance of firing shown in Fig. 4.

Our interpretation that the observed decrease in inhibitory synaptic weights in our model (Fig. 2) agrees with observations is more subtle and open to interpretation. The direct experimental observation in [28] is that the density of GABAergic neurons measured by GABA-immunoreactive staining decreased in older subjects compared to younger subjects; however, the total density of neurons (measured by Nissl staining) *did not* decrease. This interpretation is supported by other work [38], which finds minimal neuron loss with age and notes that during aging genes involved with GABA-mediated neurotransmission are strongly down-regulated in human and macaque pre-frontal cortex. Thus, observed age-related changes are not caused by a loss of GABAergic neurons through cell death, but more likely a drop in GABA expression itself. In a conductance-based leaky-integrate-and-fire model this could be interpreted as a decrease in the maximum fraction of GABAergic receptors being activated by a pre-synaptic spike; as E-I Net is a current-based leaky-integrate-and-fire model, the closest equivalent effect is a change in synaptic weight. Because the reversal potential of GABAergic synapses is typically much larger in magnitude than the membrane potential, it is often valid to approximate the inhibitory conductance-based synapses as current-based synapses [39]. Therefore, we interpret the observed decrease in the range of inhibitory input and lateral synaptic weights in our model to be consistent with the experimental observations. This said, our model also predicts a decrease in the range of excitatory input and lateral weights. To our knowledge, a similar measurement has not been performed to measure changes in density of excitatory neurotransmitters. Thus, this result serves as a prediction for future experiments, with the caveat that such a measurement could be complicated by the fact that there are several principal excitatory neurotransmitters to consider, as well as the fact that the true conductance-based nature of excitatory synapses may become relevant due to the fact that excitatory reversal potentials are not significantly larger than membrane potentials, so a current-based approximation may not fully capture changes in excitatory synapses with age.

Several studies have found decreased sensitivity to oriented and directed grating stimuli in older neurons and subjects [27]. Our aged network model replicates this loss of orientation selectivity to grated stimuli compared to the mature young network; our model did not develop significant direction selectivity, as we discuss further in *Limitations of the model*, below. As our model gives us full access to network parameters, we were able to demonstrate that this loss in orientation selectivity is primarily caused by the deterioration of Gabor-like receptive fields of the model neurons as the runaway target excitability causes the network to reorganize, rather than, say, changes in lateral weights that might also have been predicted to impact

sensitivity [31–37]. In particular, we found both that the net proportion of Gabor-like receptive fields decreased as the network aged (Table 1) as well as the fact that the receptive field of a neuron was not preserved with age (Fig. 3). The net loss of Gabor-like receptive fields impacting orientation selectivity makes intuitive sense, as such receptive fields are typically interpreted to be edge-detectors [40]. This is also a new prediction that could potentially be tested experimentally in longitudinal experiments on the same individuals, measuring the receptive fields of a population of neurons in visual cortex in both youth and old age, inferring the receptive fields using, for example, generalized linear model fits [41–46], and comparing the distributions of the receptive fields at these different ages. It would likely be difficult to track a single neuron across ages to test our model's prediction that individual receptive fields change with age.

Other experimental work has also found that contrast sensitivity is diminished with age [47]. We did not perform a detailed investigation of whether our model could qualitatively replicate this result, for technical reasons: images on which the network model is trained are first "whitened" to remove the mean intensity of the images and normalize the variance across pixels in an image—operations interpreted within the context of E-I Net to be performed by the retina and lateral geniculate nucleus (LGN) en route to V1. However, this pre-processing step will also normalize the contrast of an image. To properly test contrast sensitivity would thus require a model of the aging retina and LGN, which is beyond the scope of this work, but could potentially be developed and studied with a similar setup, using, for instance, recent deep network models of retina [48, 49].

Limitations of the model

Although our model successfully replicated several experimental observations and makes predictions for new experiments, there are of course some results the model as currently implemented could not replicate. Moreover, as a simplified model of network development and dynamics, there could be features of the model that are considered biologically implausible or oversimplified, and bear discussion here.

As mentioned above, one of Ref. [27]'s results that our model did not successfully reproduce is the loss of direction selectivity. We trained our network on sequential patches from the Catcam database [50, 51] for the purpose of testing direction selectivity, but found that our network was not complex enough to learn direction selectivity from training on these natural movies alone. Our network was able to learn some direction selectivity when trained on moving grating stimuli, but this begins to overwrite the network structure learned from the natural scenes, and therefore we considered it invalid to induce direction selectivity by training on the types of stimuli on which we would later test network responses.

The difficulty of learning direction selectivity could be because directed motion is not represented frequently enough in our training movies, but could also be due to the fact that our network neurons exponentially filter inputs, all with the same single timescale. Potentially, a modified network model that allowed for different spatiotemporal history filters or even multiple membrane timescales might be better suited to learn direction selectivity from natural movies.

The biological plausibility of training neural networks has been debated extensively in the literature [52,

53], casting doubt on whether training such networks can be an appropriate model for network development. Many, if not most of these criticisms, focus on training algorithms that employ non-local learning rules (i.e., rules which require global information about the network in order to update an individual neuron's synaptic connections), such as backpropagation [54]. However, the learning rules employed by E-I Net and its predecessor SAILnet [24, 25] are local (Eqs. (1)-(3)), such that updating a particular neuron's parameters only requires information about its own incoming connections and firing threshold. Therefore, the learning rules used in our model may be considered a reasonable simplification of biologically plausible learning mechanisms, and thus can capture the essential minimal features of network development.

One subtle and potentially unintuitive feature of our model is the fact that the time-course of training need not be straightforwardly mapped onto the time-course of development. For instance, we use the same learning rates for the entire duration of our training procedure, including both the development phase (loops $1 - 30$) and our aging phase (loops $30 - 80$). However, experimental evidence indicates that there is a "critical period" in which long-term synaptic plasticity is high early in development, after which plasticity drops, allowing the nervous system to approximately lock-in learned structure and prevent sudden drastic changes in sensory experience from overwriting these structures [55]. At first, this seems at odds with our training procedure, which has no such critical period, but that is true only if one assumes a linear relationship between training time and true organism age. A nonlinear relationship would allow for our results to be consistent with a critical period. In particular, we could assume that training time and organism age are approximately linearly related during development (loops $1 - 30$) but undergo a shift to a different linear relationship after maturation tht holds during the old-age phase. We have checked that decreasing all learning rates by a factor of 10 and extending stimulus presentation time by a factor of 10 after 30 training loops indeed give results qualitatively similar to the baseline model; see Fig. 9 in Methods.

Finally, our results do not rule out other possible mechanisms for generating the observed changes in networks with age, nor do they explain what physiological mechanisms might cause a drift in target firing rate to begin with. Possible causes could be related, for example, to changes in calcium dynamics controlling synaptic plasticity [56], but a detailed study of such possible breakdowns of this firing rate regulation are beyond the scope of the current work. Nevertheless, starting from the hypothesis that there is some breakdown of homeostatic regulation of neural firing in the presence of ongoing long-term synaptic plasticity has allowed us to replicate several experimental results and make predictions for effects not yet observed experimentally, which would enable future validation or refutation of this hypothetical mechanism.

Conclusions and implications for disease

We have shown that dysregulation of neural excitability in the presence of ongoing long-term synaptic plasticity is a potential mechanism that could explain several experimentally observed changes in both physiology and functional performance in healthy aging brain networks, in addition to making predictions for future studies in aging neural tissue. Should further experimental evidence lend support to this mechanism, our results also suggest possible ways some of these performance changes could be partly mitigated, at

least in terms of visual acuity. For instance, our numerical experiment altering the structure and strength of receptive field input into visual cortex (Fig. 8) suggests that in advanced age boosting the strength of sensory inputs by increasing the gain of those neurons could be more effective for improving orientation selectivity than resorting the spatial structure of inputs, which would be a more complex and potentially invasive treatment, as it would involve a coordinated adjustment of synaptic strengths.

Beyond replicating previous experimental observations and making predictions for new experiments, modeling how neural circuitry changes with age will play an important role in understanding the progression of neurological disorders whose onset correlates with age, such as Alzheimer's, Parkinson's disease, and general age-related cognitive decline. In particular, having a baseline model for predicting changes in healthy neural circuitry is important for disambiguating which effects are symptomatic of disease from those likely to be present in a healthy aging individual, and thereby determining which aspects of the disease progression to focus on and develop treatments for.

Acknowledgments

We thank Paul King, Arianna Maffei, Andrew Merrill, Lorna Role, and Joel Zylberberg for helpful discussions. ST would like to thank Jim and Marilyn Simons and the Simons Foundation at Stony Brook University for the Simons Summer Research Program that supported this work. ST also thanks Catlin Gabel School for incorporating this research into his coursework. BAWB would like to thank Stony Brook University for financial support.

A Methods

A.1 Network Model

We use E-I Net, a spiking network model of V1 trained on natural images, as the basis for our investigations. E-I Net is written in MATLAB, and all modifications and subsequent analysis code are also in MATLAB. We will give a brief overview of the model here, and then focus on modifications of the model we implemented for use in this study. Some of these modifications were non-default options already available in the code, while others are modifications we developed ourselves. For full details on E-I Net we direct readers to [24].

E-I Net is a network of leaky integrate-and-fire neurons that spike in response to pixel-image inputs X_k. The network, partitioned into excitatory and inhibitory subpopulations, learns input weights Q_{ik}, synaptic weights W_{ij}, and firing thresholds θ_i according to learning rules that impose linear decoding of the image inputs (Eq. 5), minimize redundancy in spiking (Eq. 6), and promote sparsity in firing (Eq. 7):

$$\Delta Q_{ik} = \alpha(n_i X_k - n_i^2 Q_{ik}) \tag{5}$$

$$\Delta W_{ij} = \beta\,(n_i n_j - \langle n_i \rangle \langle n_j \rangle (1 + W_{ij})) \tag{6}$$

$$\Delta \theta_i = \gamma(n_i - p), \tag{7}$$

where, as in the main text, n_i is the spike count in response to the pixel intensity X_k, $\langle n_i \rangle$ is an exponentially-weighted moving average during the training loop, and p is the target spike count, which may be different for excitatory and inhibitory neurons. The learning rates α, β, and γ are chosen as in the original E-I Net model, except for the numerical experiments discussed in Section 2. The global learning rate, a multiplier factor applied to all of the learning rates, is set to 0.4, the value used for accelerated training provided in examples in the E-I Net code repository. The network learning rules are entirely local [24, 25] (i.e., changes to a neuron's parameters depend only on inputs that neuron receives), and hence admit a straightforward biological interpretation as plastic or homeostatic mechanisms that shape network wiring or firing. When trained on natural image inputs the network has been shown to learn realistic Gabor-like receptive fields and produce activity that can be used to accurately reconstruct input [24, 25].

A.2 Training the network

Training our network model consists of two phases: the initial development phase, during which the network forms, and the aging phase during which we implement modifications to the learning rules to cause changes in network structure that mimic the effects of old age. The first phase of training is largely the same as the original E-I Net training procedure, while the aging phase is a novel modification that we implement after the first phase reaches a steady state.

We train our network on scenes from the CatCam database [50,51]. Our original intent was to investigate whether the network would learn direction selectivity in addition to orientation selectivity, as noted in the Discussion, though we ultimately only focused on the orientation selectivity of the network. Nonetheless, to accommodate the continuous nature of the CatCam movies and gratings used for training and testing, respectively, we draw input data from frames in sequence, rather than randomly, and we switch on a built-in E-I Net toggle to make neuron membrane potential continuous across training epochs. Thus, for each epoch, we randomly select 100 starting frames from the input movie, excluding the final 50 frames. For each starting frame, we create a copy of the network and randomly choose a square patch. Each parallel network receives this patch from the associated starting frame as stimulus and then, in series, the patch at the same coordinates from each of the next 49 frames. At the end of the epoch, the parameters of the main network evolve based on the results from all one hundred network copies taken in combination.

In our model we take the fully trained network to represent our mature "young" network, trained for 30 loops at the baseline excitatory target spike rate of $p_E = 0.01$ expected spikes per image presentation and $p_I = 0.05$ expected spikes per image presentation for inhibitory neurons, comparable to the values used in [24]. The goal of this work is to modify this model to reflect observed age-related physiological changes, such as increased excitability and decreased inhibition, and thereby produce "aged" networks in which we

can test how functional performance and coding have changed. Rather than implement the physiological changes by several independent by-hand adjustments, a much more parsimonious modification is to alter only the excitatory target spiking rate (ETSR) p_E. If we increase p_E with training time past the mature young stage of the network, the network will reorganize to achieve a higher mean firing rate close to this target. We increase p_E by 0.02 every subsequent loop, culminating in the "old" network after 80 total loops. This not only results in increased excitability (both in mean firing and variance of firing), but also naturally gives rise to a decrease in inhibition (Fig. 2). The changes in selectivity to oriented grating stimuli also emerge from this single by-hand change. To check how sensitive the aging procedure is to the initial starting value of the ETSR, we also ran two additional network simulations starting from higher values, $p_E(0) = 0.05$ and $p_E(0) = 0.07$ spikes per image; see Supplementary Figure 25.

Because our method produces a sequence of aged networks that are gradual evolutions of earlier "ages," we loosely interpret the training process as the "aging" process in our network. (As noted in the Discussion, a concrete relationship between training loops and organism age is not well-defined). This is similar in spirit to previous work interpreting the initial training phase of the network as a model for the developmental period of visual cortex [26].

A.3 Testing network responses to oriented gratings

To measure orientation selectivity, every 5 loops we pause training and test the network on moving grating stimuli (see Fig. 5) of 4 possible orientations in order to measure the selectivity of model neurons, during which we set the global network learning rate to zero so that network parameters do not adapt to the statistics of the gratings. We generate the gratings at spatial and temporal frequencies of 0.1 cycles/pixel and 0.1 cycles/frame, respectively, using the `GratingStim` module of `VisualStimulusToolbox` [57]. As in the training phase, we use 8×8 pixel patches of these grating images as visual input to our network. For this patch size there is approximately one cycle of the grating per patch.

For each testing session we record the spiking activity of all neuron in response to grating patches of each orientation and use these responses to compute the orientation selectivity of each neuron. While there are various methods of determining orientation selectivity [58, 59], we choose to follow [27] and quantify the overall selectivity of each neuron to grating orientation as

$$OS_i = \frac{\left|\sum_{\text{trials}} \mathbf{r}_i^{\text{trial}}\right|}{\sum_{\text{trials}} \left|\mathbf{r}_i^{\text{trial}}\right|},$$

where $\mathbf{r}_i^{\text{trial}}$ is a vector of neuron i's responses to each orientation on a given trial.

A.4 "Gabor-ness" analysis

The receptive fields (RFs) of the neurons in the network are observed to evolve over the aging process (Fig. 3). We quantify this change by fitting a Gabor wavelet profile G to each RF in both youth and old age

and rejecting those fits for which the normalized goodness-of-fit $||G - RF||^2/||RF|| > 0.8$. This is the same procedure as performed in [25], except that we use a less strict rejection condition (0.8 vs. 0.5). The Gabor profile takes the following form:

$$G = A\cos\left(2\pi f x_p + \psi\right)\exp\left(-\frac{x_p^2}{(2\sqrt{2}\sigma_x)^2} - \frac{y_p^2}{(2\sqrt{2}\sigma_y)^2}\right),$$
$$x_p = (x - x_0)\cos\theta + (y - y_0)\sin\theta,$$
$$y_p = -(x - x_0)\sin\theta + (y - y_0)\cos\theta,$$

where A is the amplitude of the Gabor wavelet, f is the frequency of amplitude oscillation, ψ is the phase, σ_x and σ_y are the spreads of the Gabor profile in the x and y directions. The Gabor wavelet is taken to be a function of the pixel-space positions x_p and y_p, related to the raw positions x and y by a rotation angle θ and offsets x_0 and y_0. In addition to rejecting fits that do not meet our goodness-of-fit criterion, we also reject fits that are centered outside the frame of the receptive field, and wholly exclude receptive fields for which fits cannot be computed from the statistics in Table 1.

A.5 Numerical experiments

A.5.1 Learning rate freezing

To examine the contribution of each network property to the observed decline in neural orientation selectivity as a result of increase in p_E, we conduct the following series of tests. We train several instances of the network to maturity (30 loops) under normal conditions. Upon the completion of the 30^{th} loop, we set the learning rate to zero for different sets of connection weights: in one instance, we zero out the learning rates for only the input connections to the excitatory and inhibitory neurons; in another, we zero out the learning rates for only the lateral connections between neurons; and in the last instance, we zero out the learning rates for both input weights and lateral weights. There is no case in which we zero out the learning rates of the thresholds, as it is this learning rule through which our network ages. After zeroing out the learning rates, we continue training these variants until old age (80 loops), probing selectivity every five loops.

A.5.2 Receptive field remapping

To decouple the effects of the changing RF structure and magnitude on network selectivity, we examine, independent of the training procedure, the selectivity after remapping the young input weights onto a later distribution of input weight magnitudes, for the networks after $35, 40, 45...80$ loops:

$$Q_{\text{remapped}} = \Phi_{\text{old}}^{-1}(\Phi_{\text{young}}(Q_{\text{young}})), \tag{8}$$

where $\Phi_{\text{young}}(\cdot)$ and $\Phi_{\text{old}}(\cdot)$ are the cumulative distribution functions (CDFs) for the young and old input weight distributions. Eq. (8) modifies the value of every input weight such that the distribution of remapped

young input weight *values* matches the distribution of old input weight values, but the spatial pixel organization of the RF is not altered. We empirically estimate the CDFs from the data and use interpolating splines to obtain smooth estimates (simply swapping the x- and y-values for inverse CDF estimation). Thus, we can insert these remapped weights into our old networks to test how RF structure contributes to selectivity. We find that networks with these remapped input weights exhibit consistently higher selectivity than their counterparts with the original aged input weights. We can also test the inverted remapping, $\Phi^{-1}_{\text{young}}(\Phi_{\text{old}}(Q_{\text{old}}))$, in which aged input weights are remapped onto the young distribution of magnitudes, retaining the old spatial pixel structure but with younger magnitudes. Networks with these remapped weights are comparable in selectivity to their corresponding original networks during maturation and the early stages of the aging process, and increasingly more selective late in aging.

A.6 Supporting analyses

A.6.1 Critical learning periods and time rescaling

The critical learning period hypothesis posits that there is a window early in development in which learning rates are high, but eventually taper off to lower levels of plasticity [55]. This is unlike our network model, in which the learning rates are the same throughout the training period. To demonstrate that our training procedure does not preclude interpretation in terms of a critical learning period, we run a training case in which, after 30 loops, we reduce the learning rate by a factor of 10 and extend the length of natural image stimulus presentation by a factor of 10, and then train the network for 50 more loops under these conditions. We find that average neural selectivity follows largely the same trend as in the case of normal aging. This suggests that networks with time-varying learning rates can be mapped to networks with constant learning rates by adjusting the length of training loops during these periods. The implication for our results is that we can interpret the development phase of our network as a critical learning period with higher rates and a shorter learning interval compared to the aging phase.

A.6.2 Excitatory neurons with similar receptive fields effectively mutually inhibit each other

To gauge the extent to which excitatory neurons with similar receptive fields inhibit one another, we take element (i, j) of the product of the lateral connection matrices W_{EI} and W_{IE} to approximate the net charge transfer capacity from pre-synaptic excitatory neuron i to post-synaptic excitatory neuron j, by way of all disynaptic pathways through inhibitory neurons. (Recall that in E-I Net there are no direct E-to-E connections; i.e., $W_{EE} = 0$). The greater the magnitude of $(W_{EI}W_{IE})_{ij}$, the greater the inhibition of E neuron i by E neuron j. As in Fig. 6B of [25], we plot the RF overlap, computed as the cosine similarity of the vectorized input weight matrices, against the elements of $W_{EI}W_{IE}$ for all pairs of excitatory neurons. In Fig. 10 below we see, similar to [25], that a larger overlap is generally associated with a larger $W_{EI}W_{IE}$ value and thus, per our interpretation, stronger inhibition of one excitatory neuron in the pair by the other.

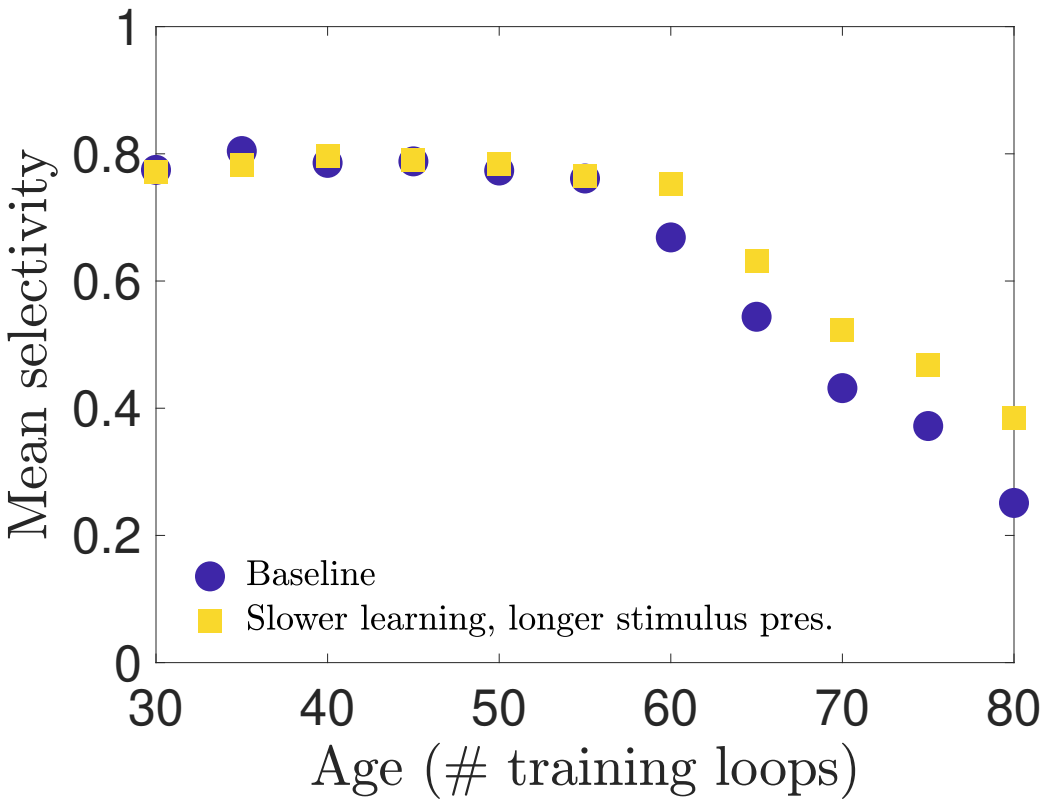

Figure 9: **Testing critical period learning**: To show that one could in principle implement a critical learning period in the training phase of our network model, we show that if the learning rates are reduced by a factor of 10 and the duration of training is increased by 10 the results are qualitatively similar. Thus, the training procedure employed in our model is expected to produce approximately the same results as a model with a critical period during which learning rates vary nonlinearly during training, so long as those changes are compensated by reciprocal changes in the duration of training during the critical period.

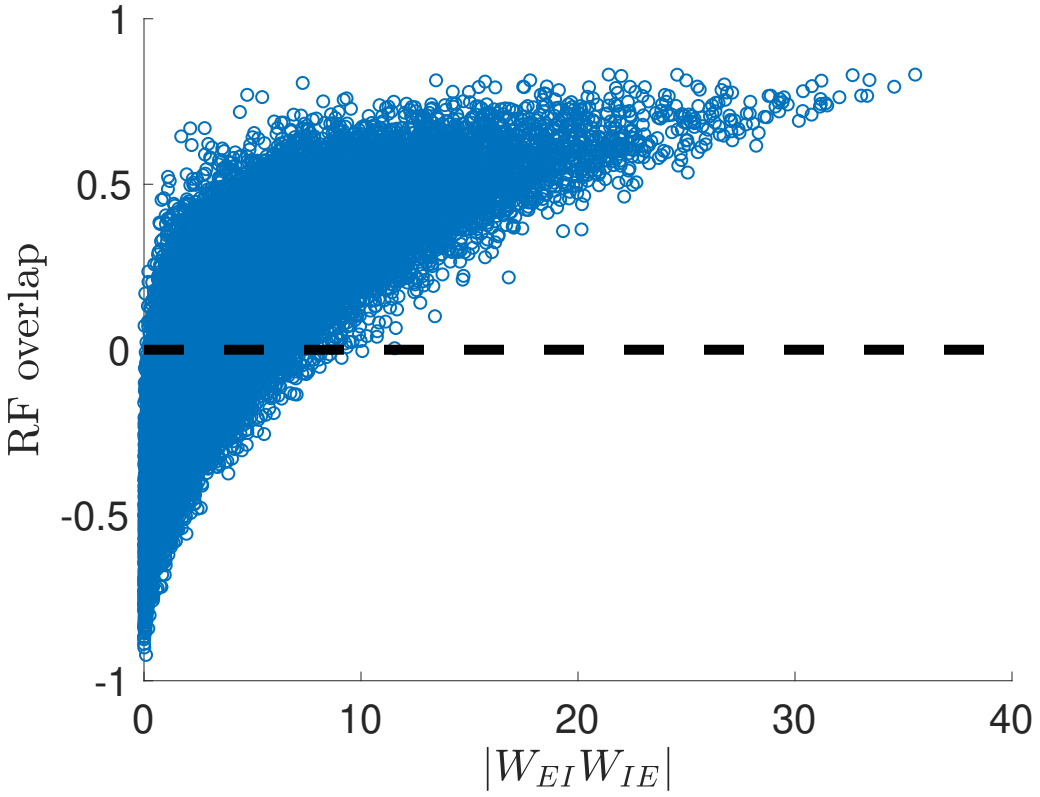

Figure 10: **Excitatory neurons with similar receptive fields effectively inhibit one another**: As a proxy for the total amount of charge excitatory neurons trasmit between each other, we use the elements of the matrix $W_{EI}W_{IE}$, which correspond to the weight of disynaptic connections between excitatory neurons, through a single inhibitory interneuron. (There are no direct E-to-E connections in our network). We plot the magnitude of $(W_{EI}W_{IE})_{ij}$ for each pair of neurons $i \neq j$ against the overlap of the receptive fields of those same two neurons. As seen in the plot, neurons with largest magnitude of inhibitory charge transfer tend to have relatively large receptive field overlaps; i.e., a large fraction of neurons with similar RFs tend to effectively inhibit one another.

A.7 Figures

Figures created in MATLAB were saved using the `export_fig` script, available on GitHub [60].

B Supplementary Information

We check the robustness of our results by training the model on three different movies from the CatCam repository (`movie01.tar`, `movie07.tar`, and `movie16.tar`) and resetting all random number generation before each run. We find no appreciable differences across the results on the three image sets. See Supplementary Figures 11-24.

movie07.tar

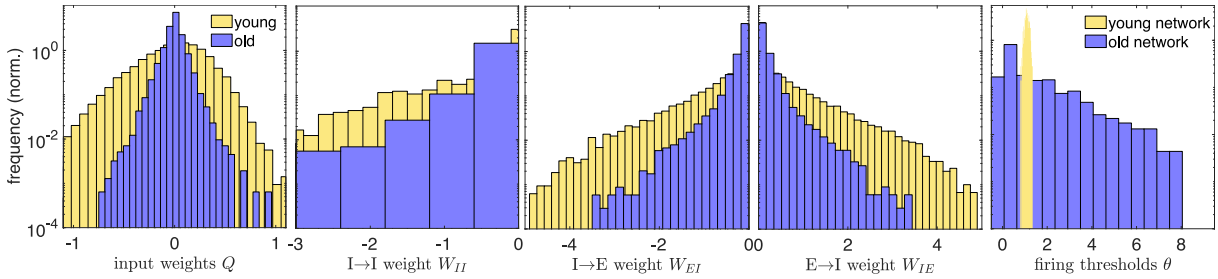

Figure 11: **Supplementary Figure: Physiological parameter changes, initial seed 2**: Same as Fig. 2 for a different initial seed and selection of training frames from `movie07` in the CatCam database.

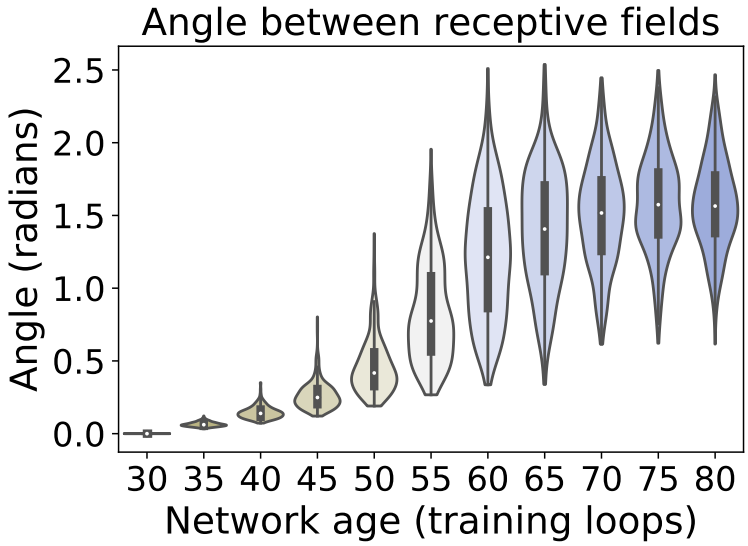

Figure 12: **Angle between receptive fields in the input weight parameter space, initial seed 2**: Same as Fig. 3B for a different initial seed and selection of training frames from `movie07` in the CatCam database.

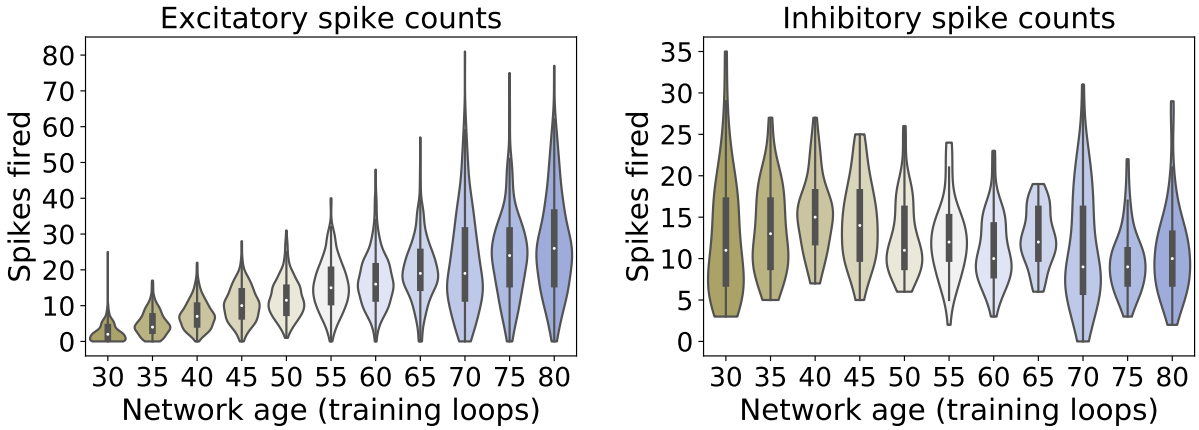

Figure 13: **Excitability of the network with age, initial seed 2**: Same as Fig. 4 for a different initial seed and selection of training frames from `movie07` in the CatCam database.

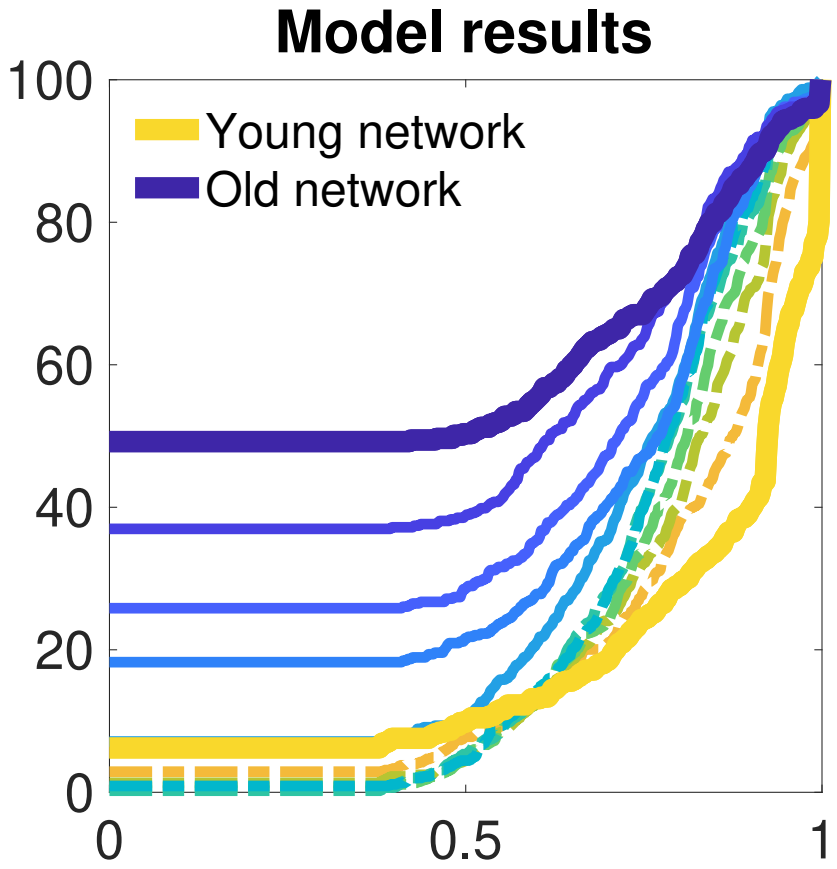

Figure 14: **Supplementary Figure: Orientation selectivity changes, initial seed 2**: Same as the model CDFs in Fig. 5 for a different initial seed and selection of training frames from `movie07` in the CatCam database.

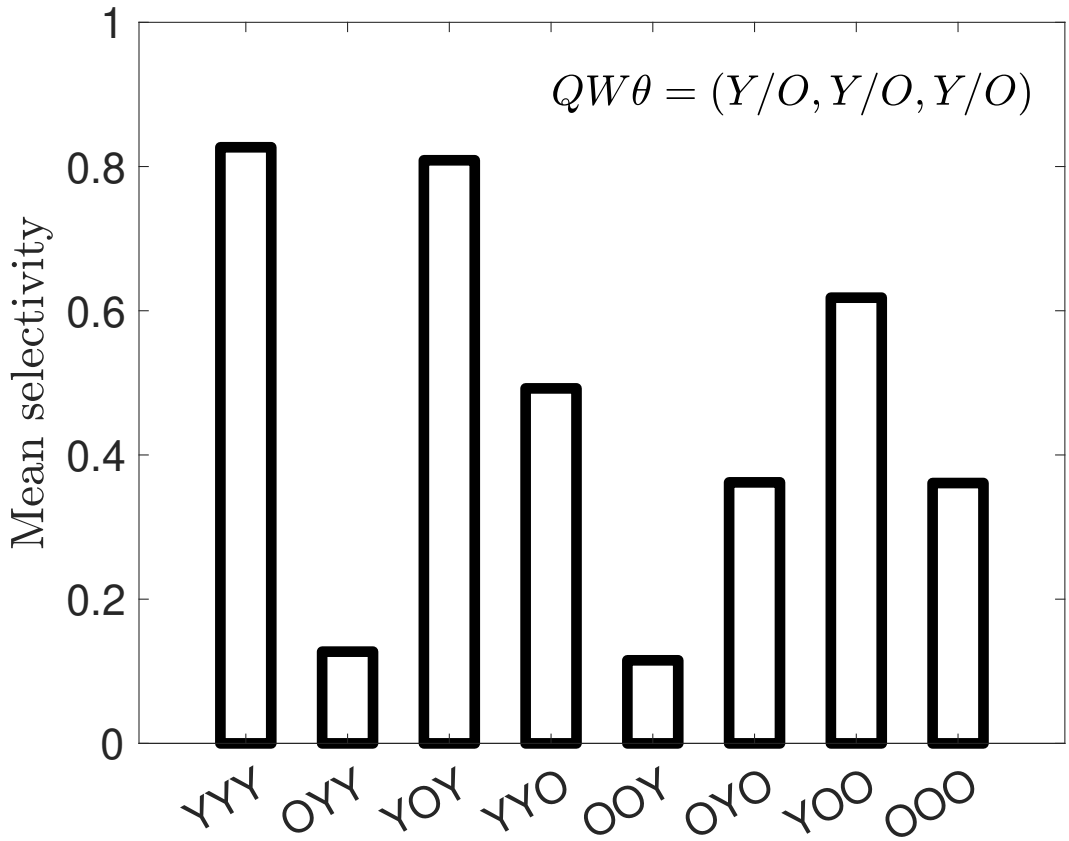

Figure 15: **Supplementary Figure: Parameter swap tests, initial seed 2**: Same as Fig. 6 for a different initial seed and selection of training frames from `movie07` in the CatCam database.

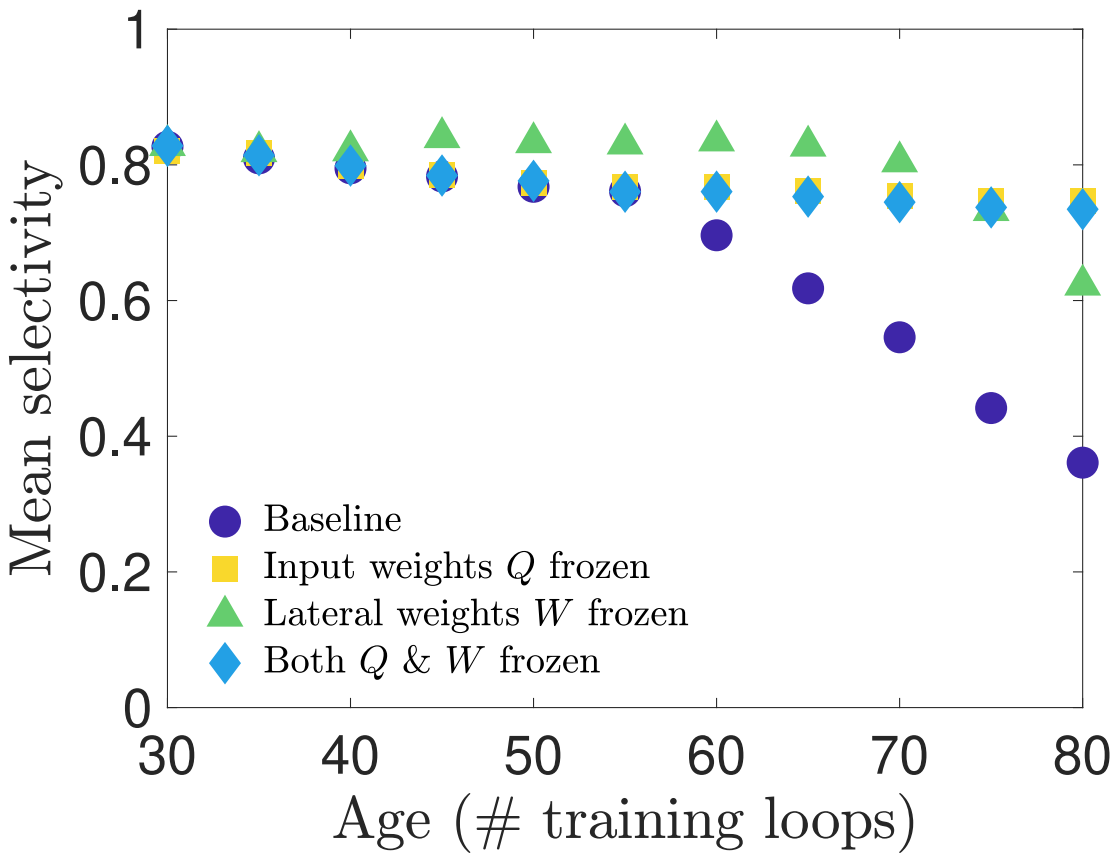

Figure 16: **Supplementary Figure: Parameter swap tests, initial seed 2**: Same as Fig. 7 for a different initial seed and selection of training frames from `movie07` in the CatCam database.

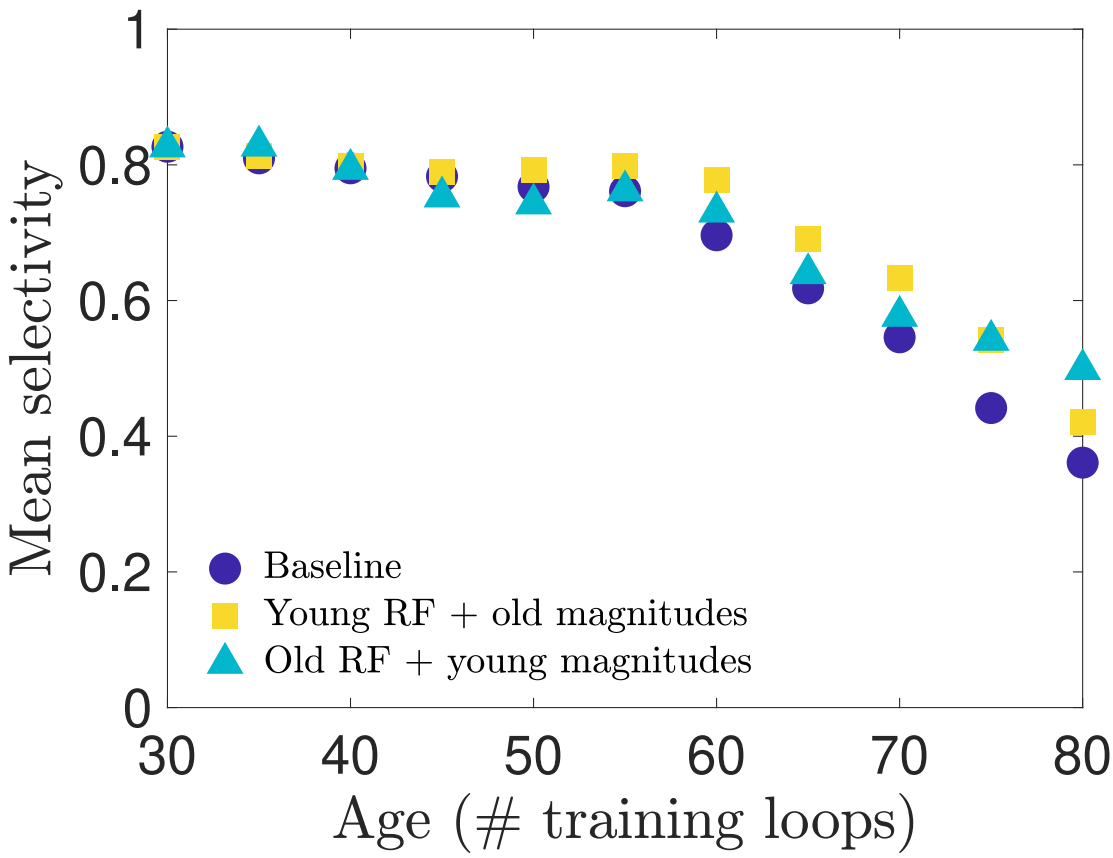

Figure 17: **Supplementary Figure: Parameter swap tests, initial seed 2**: Same as Fig. 8 for a different initial seed and selection of training frames from `movie07` in the CatCam database.

movie16.tar

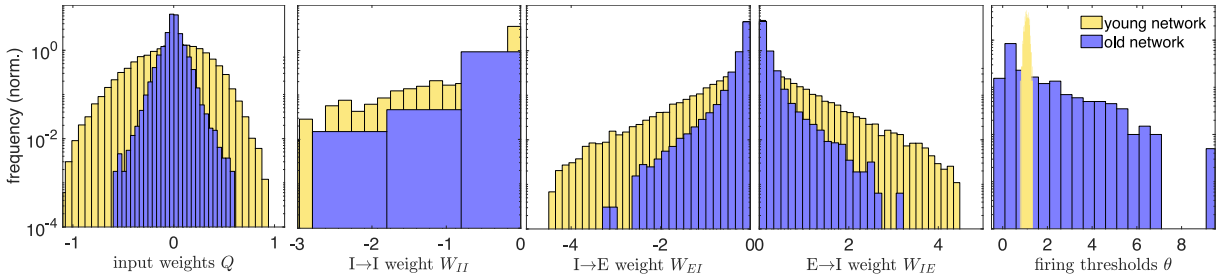

Figure 18: **Supplementary Figure: Physiological parameter changes, initial seed 3**: Same as Fig. 2 for a different initial seed and selection of training frames from `movie16` in the CatCam database.

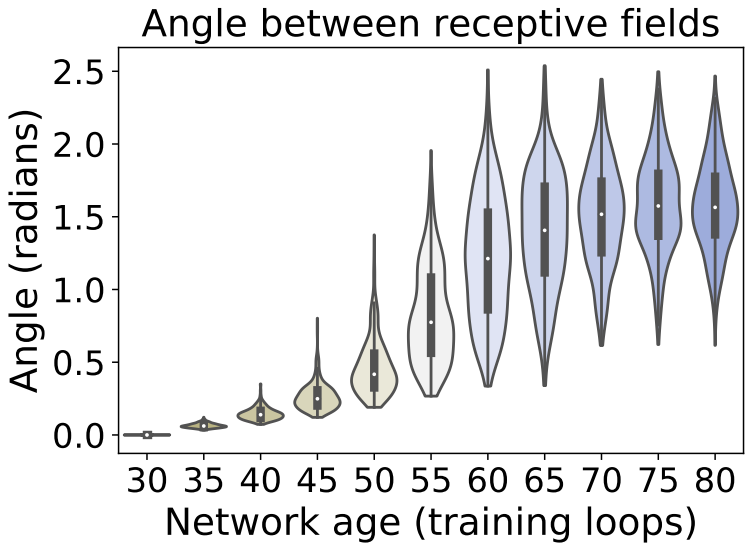

Figure 19: **Angle between receptive fields in the input weight parameter space, initial seed 3**: Same as Fig. 3B for a different initial seed and selection of training frames from `movie16` in the CatCam database.

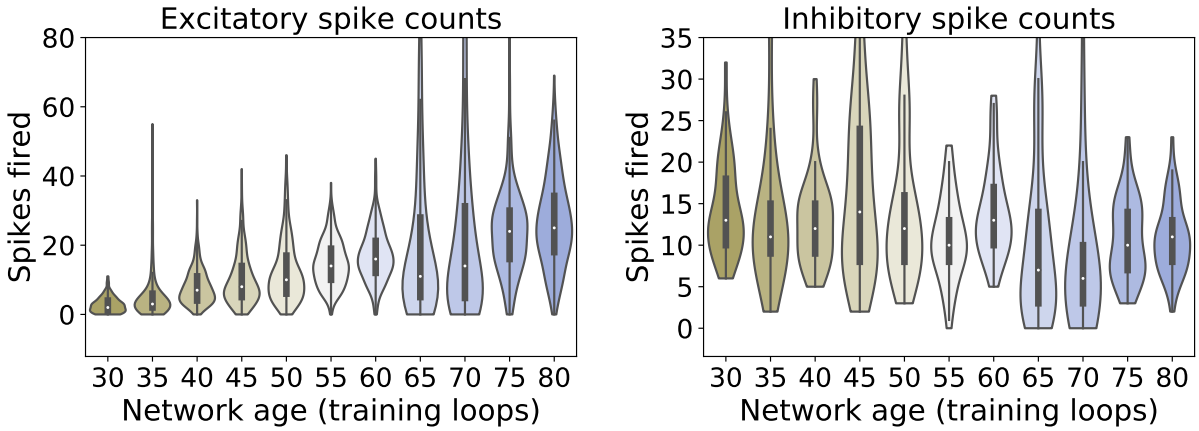

Figure 20: **Excitability of the network with age, initial seed 3**: Same as Fig. 4 for a different initial seed and selection of training frames from `movie16` in the CatCam database. In these particular simulations there were some outlier spike counts (e.g., the largest excitatory spike count was 243 at loop 65), so we have restricted the vertical axis to match the range observed in Figs. 4 and 13.

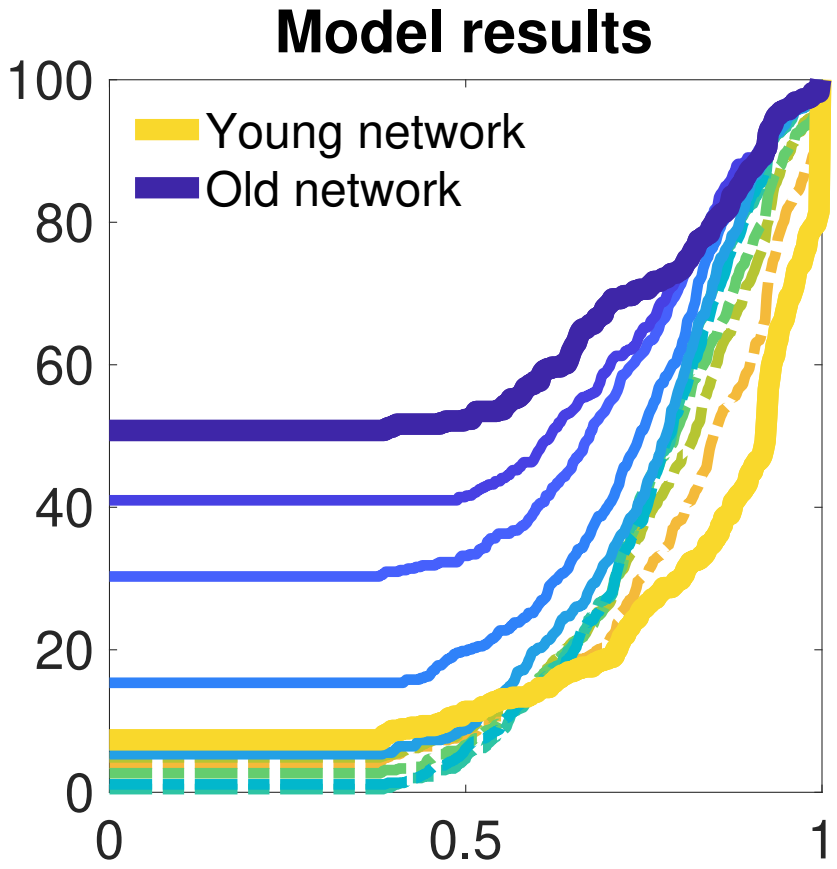

Figure 21: **Supplementary Figure: Orientation selectivity changes, initial seed 3**: Same as the model CDFs in Fig. 5 for a different initial seed and selection of training frames from `movie16` in the CatCam database.

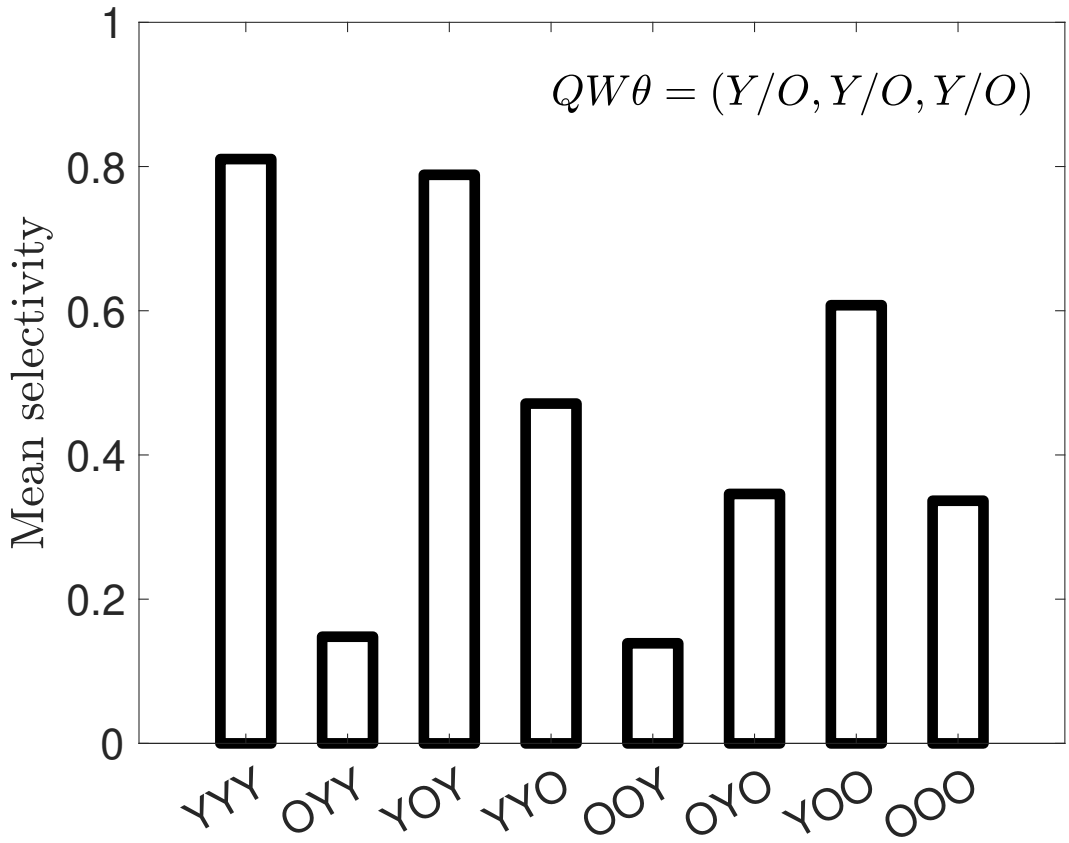

Figure 22: **Supplementary Figure: Parameter swap tests, initial seed 3**: Same as Fig. 6 for a different initial seed and selection of training frames from `movie16` in the CatCam database.

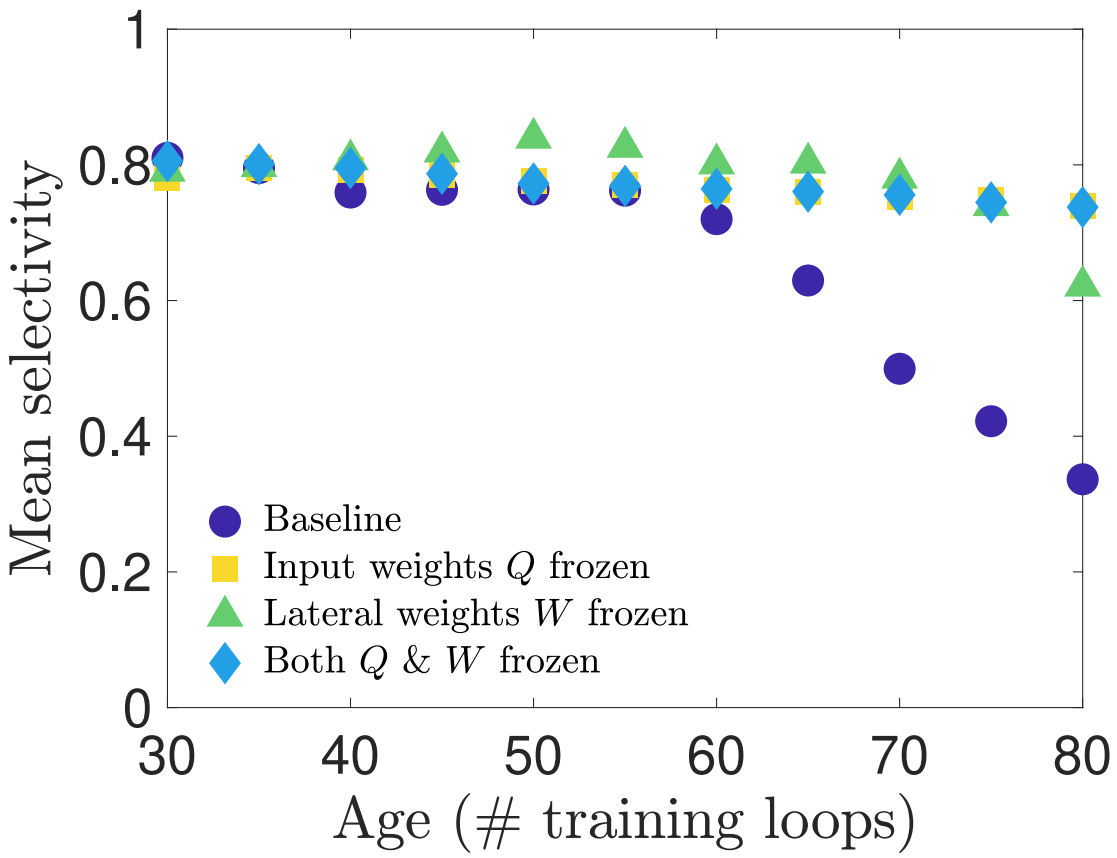

Figure 23: **Supplementary Figure: Parameter swap tests, initial seed 3**: Same as Fig. 7 for a different initial seed and selection of training frames from `movie16` in the CatCam database.

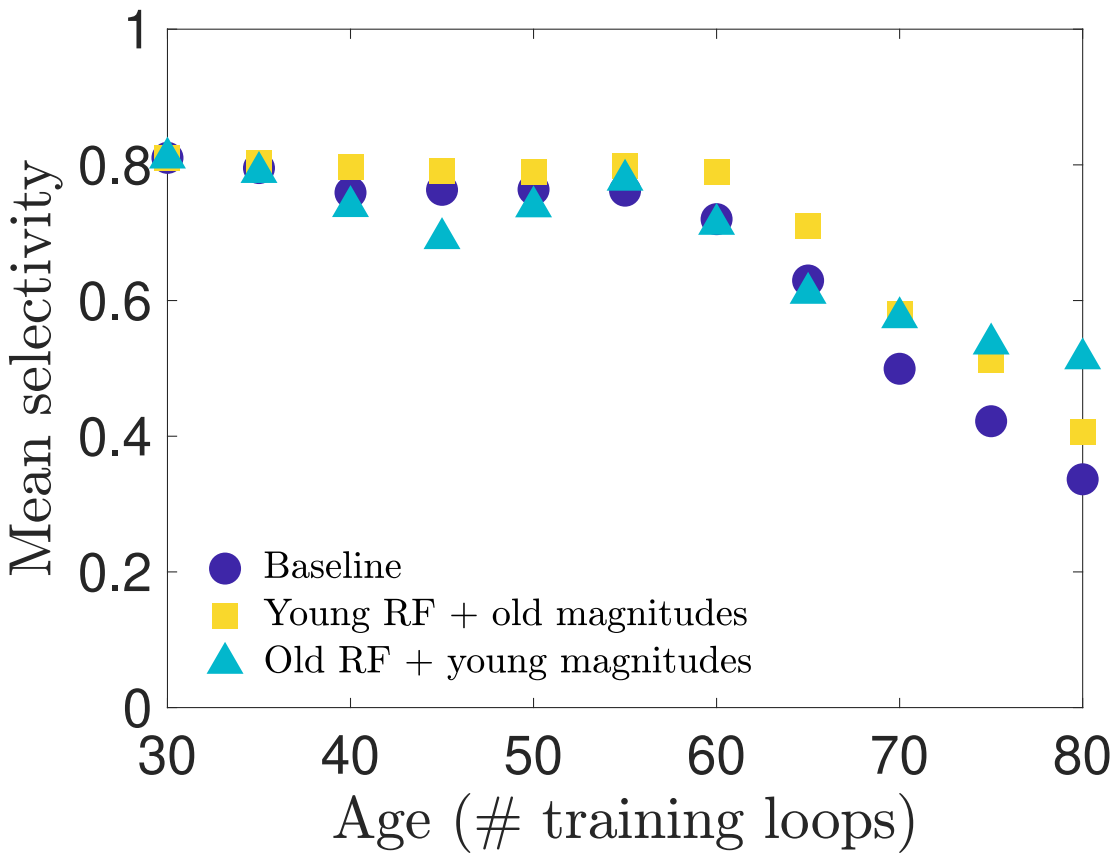

Figure 24: **Supplementary Figure: Parameter swap tests, initial seed 3**: Same as Fig. 8 for a different initial seed and selection of training frames from `movie16` in the CatCam database.

Selectivity at larger initial ETSR (Supplementary Figure)

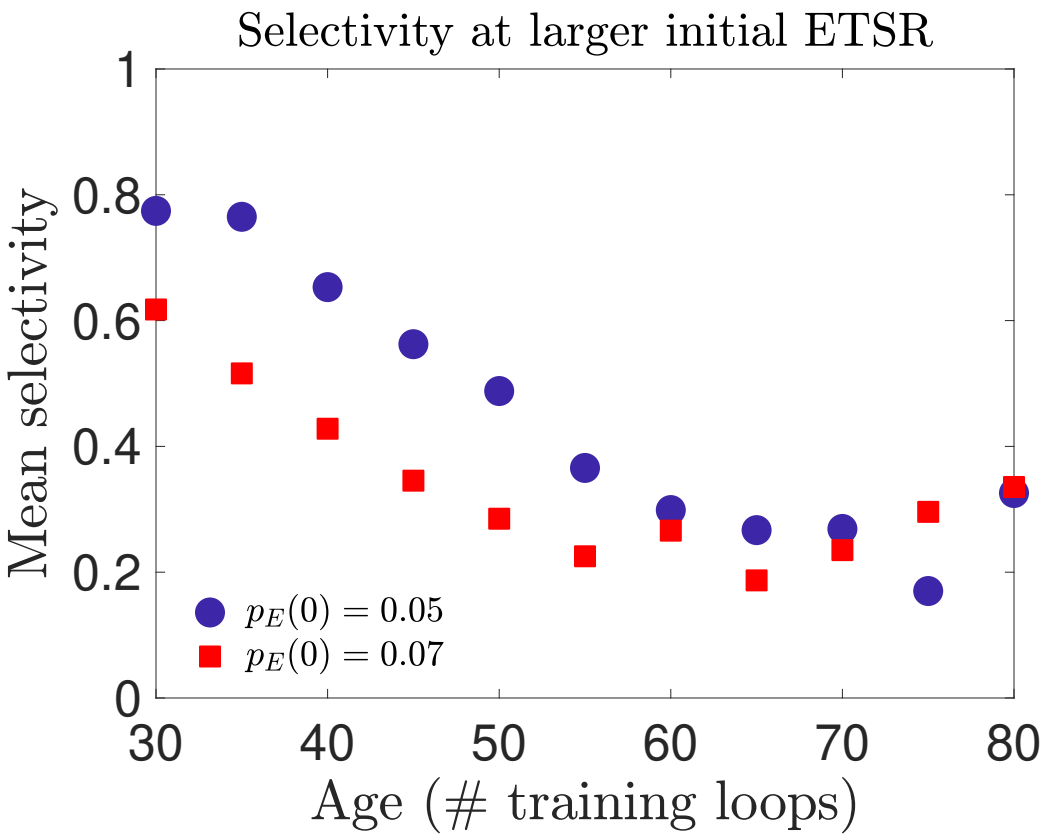

Figure 25: **Mean selectivity for larger initial target firing rates**: If the initial excitatory target spike rate is increased from the value of $p_E(0) = 0.01$ used in the main text, the mean selectivity that develops in maturity is comparable or diminished compared to the baseline results. We still observe an overall decreasing trend, although there is some increased variability and non-monotonicity evident at the later age for the less sparse case of $p_E(0) = 0.07$.

References

[1] D. Horn, E. Ruppin, M. Usher, and M. Herrmann. Neural network modeling of memory deterioration in alzheimer's disease. *Neural Computation*, 5(5):736–749, 1993.

[2] Michael E. Hasselmo. Runaway synaptic modification in models of cortex: Implications for alzheimer's disease. *Neural Networks*, 7(1):13 – 40, 1994.

[3] L. J. Tippett and M. J. Farah. A computational model of naming in alzheimer's disease: Unitary or multiple impairments? *Neuropsychology*, 8(1):3–13, 1994.

[4] Eytan Ruppin and James A. Reggia. A neural model of memory impairment in diffuse cerebral atrophy. *British Journal of Psychiatry*, 166(1):1928, 1995.

[5] ME Hasselmo. A computational model of the progression of alzheimer's disease. *M.D. computing : computers in medical practice*, 14(3):181191, 1997.

[6] ED Menschik and LH Finkel. Neuromodulatory control of hippocampal function: towards a model of alzheimer's disease. *Artificial intelligence in medicine*, 13(1-2):99121, May 1998.

[7] Martijn Meeter and Jaap M. J. Murre. Tracelink: A model of consolidation and amnesia. *Cognitive Neuropsychology*, 22(5):559–587, 2005. PMID: 21038266.

[8] Ahmed A. Moustafa, Szabolcs Keri, Mohammad M. Herzallah, Catherine E. Myers, and Mark A. Gluck. A neural model of hippocampalstriatal interactions in associative learning and transfer generalization in various neurological and psychiatric patients. *Brain and Cognition*, 74(2):132 – 144, 2010.

[9] Basabdatta Sen Bhattacharya, Damien Coyle, and Liam P. Maguire. A thalamocorticothalamic neural mass model to study alpha rhythms in alzheimers disease. *Neural Networks*, 24(6):631 – 645, 2011. Special Issue: Neurocomputational Models of Brain Disorders.

[10] K. Abuhassan, D. Coyle, and L. P. Maguire. Investigating the neural correlates of pathological cortical networks in alzheimer's disease using heterogeneous neuronal models. *IEEE Transactions on Biomedical Engineering*, 59(3):890–896, March 2012.

[11] Mark Rowan, Sam Neymotin, and William Lytton. Electrostimulation to reduce synaptic scaling driven progression of alzheimer's disease. *Frontiers in Computational Neuroscience*, 8:39, 2014.

[12] Daniela Bianchi, Pasquale De Michele, Cristina Marchetti, Brunello Tirozzi, Salvatore Cuomo, Hlne Marie, and Michele Migliore. Effects of increasing creb-dependent transcription on the storage and recall processes in a hippocampal ca1 microcircuit. *Hippocampus*, 24(2):165–177, 2014.

[13] R. Edwards, A. Beuter, and L. Glass. Parkinsonian tremor and simplification in network dynamics. *Bulletin of Mathematical Biology*, 61(1):157 – 177, 1999.

[14] Vassilis Cutsuridis and Stavros Perantonis. A neural network model of parkinson's disease bradykinesia. *Neural Netw.*, 19(4):354–374, May 2006.

[15] Steven J. Schiff. Towards model-based control of parkinson's disease. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 368(1918):2269–2308, 2010.

[16] Cliff Kerr, Sacha Van Albada, Samuel Neymotin, George Chadderdon, P. Robinson, and William Lytton. Cortical information flow in parkinson's disease: a composite network/field model. *Frontiers in Computational Neuroscience*, 7:39, 2013.

[17] Andrey Dovzhenok, Choongseok Park, Robert M. Worth, and Leonid L. Rubchinsky. Failure of delayed feedback deep brain stimulation for intermittent pathological synchronization in parkinsons disease. *PLOS ONE*, 8(3):1–10, 03 2013.

[18] Alex Pavlides, S. John Hogan, and Rafal Bogacz. Computational models describing possible mechanisms for generation of excessive beta oscillations in parkinsons disease. *PLOS Computational Biology*, 11(12):1–29, 12 2015.

[19] Karthik Kumaravelu, David T. Brocker, and Warren M. Grill. A biophysical model of the cortex-basal ganglia-thalamus network in the 6-ohda lesioned rat model of parkinson's disease. *Journal of Computational Neuroscience*, 40(2):207–229, 2016.

[20] Mark D Humphries, Jose Angel Obeso, and Jakob Kisbye Dreyer. Insights into parkinson's disease from computational models of the basal ganglia. *Journal of Neurology, Neurosurgery & Psychiatry*, 89(11):1181–1188, 2018.

[21] Ian Johnson. Age-related neurodegenerative disease research needs aging models. *Frontiers in Aging Neuroscience*, 7:168, 2015.

[22] Aina Oll-Vila, Lus F. Seoane, and Ricard Sol. Aging, computation, and the evolution of neural regeneration processes, 2019.

[23] Geoffrey A. Kerchner and Tony Wyss-Coray. *The Role of Aging in Alzheimer's Disease*, pages 197–227. Springer International Publishing, Cham, 2016.

[24] Paul D. King, Joel Zylberberg, and Michael R. DeWeese. Inhibitory interneurons decorrelate excitatory cells to drive sparse code formation in a spiking model of v1. *Journal of Neuroscience*, 33(13):5475–5485, 2013.

[25] Joel Zylberberg, Jason Timothy Murphy, and Michael Robert DeWeese. A sparse coding model with synaptically local plasticity and spiking neurons can account for the diverse shapes of v1 simple cell receptive fields. *PLOS Computational Biology*, 7(10):1–12, 10 2011.

[26] Joel Zylberberg and Michael Robert DeWeese. Sparse coding models can exhibit decreasing sparseness while learning sparse codes for natural images. *PLOS Computational Biology*, 9(8):1–10, 08 2013.

[27] Tianmiao Hua, Xiangrui Li, Lihua He, Yifeng Zhou, Yongchang Wang, and Audie G. Leventhal. Functional degradation of visual cortical cells in old cats. *Neurobiology of Aging*, 27(1):155 – 162, 2006.

[28] Tianmiao Hua, Chuanchao Kao, Qingyan Sun, Xiangrui Li, and Yifeng Zhou. Decreased proportion of gaba neurons accompanies age-related degradation of neuronal function in cat striate cortex. *Brain Research Bulletin*, 75(1):119 – 125, 2008.

[29] Supplementary information.

[30] Blake A. Richards, Timothy P. Lillicrap, Philippe Beaudoin, Yoshua Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, Rui Ponte Costa, Archy de Berker, Surya Ganguli, Colleen J. Gillon, Danijar Hafner, Adam Kepecs, Nikolaus Kriegeskorte, Peter Latham, Grace W. Lindsay, Kenneth D. Miller, Richard Naud, Christopher C. Pack, Panayiota Poirazi, Pieter Roelfsema, João Sacramento, Andrew Saxe, Benjamin Scellier, Anna C. Schapiro, Walter Senn, Greg Wayne, Daniel Yamins, Friedemann Zenke, Joel Zylberberg, Denis Therien, and Konrad P. Kording. A deep learning framework for neuroscience. *Nature Neuroscience*, 22(11):1761–1770, 2019.

[31] R. H. S. Carpenter and Colin Blakemore. Interactions between orientations in human vision. *Experimental Brain Research*, 18(3):287–303, 1973.

[32] Michael Wehr and Anthony M. Zador. Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nature*, 426(6965):442–446, 2003.

[33] Jeffry S. Isaacson and Massimo Scanziani. How inhibition shapes cortical activity. *Neuron*, 72(2):231–243, 2019/12/11 2011.

[34] Kenneth D. Harris and Thomas D. Mrsic-Flogel. Cortical connectivity and sensory coding. *Nature*, 503(7474):51–58, 2013.

[35] Martin Boerlin, Christian K. Machens, and Sophie Denve. Predictive coding of dynamical variables in balanced spiking networks. *PLOS Computational Biology*, 9(11):1–16, 11 2013.

[36] Michael A. Schwemmer, Adrienne L. Fairhall, Sophie Denéve, and Eric T. Shea-Brown. Constructing precisely computing networks with biophysical spiking neurons. *Journal of Neuroscience*, 35(28):10112–10134, 2015.

[37] Gabrielle J Gutierrez and Sophie Denve. Population adaptation in efficient balanced networks. *eLife*, 8:e46926, sep 2019.

[38] Nicholas Bishop, Tao Lu, and Bruce Yankner. Neural mechanisms of ageing and cognitive decline. *Nature*, 464:529–35, 03 2010.

[39] Peter Dayan and L. F. Abbott. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. The MIT Press, 2005.

[40] Charles F. Stevens. Novel neural circuit mechanism for visual edge detection. *Proceedings of the National Academy of Sciences*, 112(3):875–880, 2015.

[41] EP Simoncelli, L Paninski, JW Pillow, and O Schwartz. Characterization of Neural Responses with Stochastic Stimuli. In M Gazzaniga, editor, *The Cognitive Neurosciences*, pages 327–338. MIT Press, 3rd edition, 2004.

[42] L Paninski, JW Pillow, and EP Simoncelli. Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model. *Neural computation*, 16(12):2533–61, December 2004.

[43] Jayant E. Kulkarni and Liam Paninski. Common-input models for multiple neural spike-train data. *Network: Computation in Neural Systems*, 18(4):375–407, 2007. PMID: 17943613.

[44] JW Pillow, Jonathon Shlens, L Paninski, and Alexander Sher. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, pages 1–13, 2008.

[45] M Vidne, Y Ahmadian, J Shlens, JW Pillow, JE Kulkarni, AM Litke, EJ Chichilnisky, EP Simoncelli, and L Paninski. Modeling the impact of common noise inputs on the network activity of retinal ganglion cells. *J Comput Neurosci*, 33(1):97–121, 2012.

[46] Liam Paninski. Maximum likelihood estimation of cascade point- process neural encoding models. *Network: Computation in Neural Systems*, 6536(October), 2015.

[47] Zhengchun Wang, Zhimo Yao, Nini Yuan, Zhen Liang, Guangxing Li, and Yifeng Zhou. Declined contrast sensitivity of neurons along the visual pathway in aging cats. *Frontiers in Aging Neuroscience*, 6:163, 2014.

[48] Niru Maheswaranathan, Lane T. McIntosh, David B. Kastner, Josh B. Melander, Luke Brezovec, Aran Nayebi, Julia Wang, Surya Ganguli, and Stephen A. Baccus. Deep learning models reveal internal structure and diverse computations in the retina under natural scenes. *bioRxiv*, 2018.

[49] Hidenori Tanaka, Aran Nayebi, Niru Maheswaranathan, Lane McIntosh, Stephen Baccus, and Surya Ganguli. From deep learning to mechanistic understanding in neuroscience: the structure of retinal prediction. In H. Wallach, H. Larochelle, A. Beygelzimer, F. dAlché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8535–8545. Curran Associates, Inc., 2019.

[50] Belinda Y. Betsch, Wolfgang Einhäuser, Konrad P. Körding, and Peter König. The world from a cat's perspective – statistics of natural videos. *Biological Cybernetics*, 90(1):41–50, Jan 2004.

[51] Belinda Y. Betsch, Wolfgang Einhuser, Konrad P. Körding, and Peter König. Catcam dataset [biol. cybern. 90(1):41-50], January 2004.

[52] Bernd Illing, Wulfram Gerstner, and Johanni Brea. Biologically plausible deep learning but how far can we go with shallow networks? *Neural Networks*, 118:90 – 101, 2019.

[53] Aboozar Taherkhani, Ammar Belatreche, Yuhua Li, Georgina Cosma, Liam P. Maguire, and T.M. McGinnity. A review of learning in biologically plausible spiking neural networks. *Neural Networks*, 122:253 – 272, 2020.

[54] Pietro Mazzoni, Richard A. Andersen, and Michael I. Jordan. A More Biologically Plausible Learning Rule Than Backpropagation Applied to a Network Model of Cortical Area 7a. *Cerebral Cortex*, 1(4):293–307, 07 1991.

[55] Christiaan N. Levelt and Mark Hbener. Critical-period plasticity in the visual cortex. *Annual Review of Neuroscience*, 35(1):309–330, 2012. PMID: 22462544.

[56] Moritz Helias, Stefan Rotter, Marc-Oliver Gewaltig, and Markus Diesmann. Structural plasticity controlled by calcium based correlation detection. *Frontiers in Computational Neuroscience*, 2:7, 2008.

[57] Michael Beyeler. Visual stimulus toolbox: v1.0.0, June 2016.

[58] Mark Mazurek, Marisa Kager, and Stephen D. Van Hooser. Robust quantification of orientation selectivity and direction selectivity. *Frontiers in Neural Circuits*, 8:92, 2014.

[59] EL Bienenstock, LN Cooper, and PW Munro. Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2(1):32–48, 1982.

[60] Yair Altman. `export_fig`, GitHub. `https://www.github.com/altmany/export_fig`, Retrieved December 13, 2019.