

The Development of Active Binocular Vision under Normal and Alternate Rearing Conditions

Lukas Klimmasch¹, Johann Schneider¹, Alexander Lelais¹, Bertram E. Shi², Jochen Triesch¹

*For correspondence:

klimmasch@fias.uni-frankfurt.de
(FMS); triesch@fias.uni-frankfurt.de
(FS)

¹Frankfurt Institute for Advanced Studies, Frankfurt am Main, Germany; ²Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Hong Kong, China

Abstract The development of binocular vision is an active learning process comprising the development of disparity tuned neurons in visual cortex and the establishment of precise vergence control of the eyes. We present a computational model for the learning and self-calibration of active binocular vision based on the Active Efficient Coding framework, an extension of classic efficient coding ideas to active perception. Under normal rearing conditions, the model develops disparity tuned neurons and precise vergence control, allowing it to correctly interpret random dot stereograms. Under altered rearing conditions modeled after neurophysiological experiments, the model qualitatively reproduces key experimental findings on changes in binocularity and disparity tuning. Furthermore, the model makes testable predictions regarding how altered rearing conditions impede the learning of precise vergence control. Finally, the model predicts a surprising new effect that impaired vergence control affects the statistics of orientation tuning in visual cortical neurons.

Introduction

Humans and other species learn to perceive the world largely autonomously. This is in sharp contrast to today's machine learning approaches (*Kotsiantis et al., 2007; Jordan and Mitchell, 2015*), which typically use millions of carefully labeled training images in order to learn to, say, recognize an object or perceive its three-dimensional structure. How can biological vision systems learn so much more autonomously? The development of binocular vision presents a paradigmatic case for studying this question. This development is an active process that includes the learning of appropriate sensory representations and the learning of precise motor behavior. Species with two forward facing eyes learn to register small differences between the images projected onto the left and right retinas. These differences are called binocular disparities and are detected by populations of neurons in visual cortex (*Kandel et al., 2000; Blake and Wilson, 2011*) that have receptive subfields in both eyes. Frequently, they are modeled using separate Gabor-shaped filters for each eye, where the disparity is encoded by a shift in the centers of the filters, a difference between their phases, or by a combination of both (*Fleet et al., 1996; Chen and Qian, 2004*). The responses of such disparity tuned neurons can be used to infer the three-dimensional structure of the world. At the same time, we also learn to align our eyes such that the optical axes of our two eyes converge on the same point of interest. These so-called vergence eye movements are also learned and fine-tuned during development (*Held et al., 1980; Fox et al., 1980; Stidwill and*

41 *Fletcher, 2017*). Again, this learning does not require any obvious help from outside, but must rely
42 on some form of self-calibration.

43 While it has long been argued that the development of disparity tuning and vergence eye move-
44 ments are interdependent (*Hubel and Wiesel, 1965*), it has been only recently that computational
45 models have tried to explain how the learning of disparity tuning and vergence eye movements
46 are coupled and allow the visual system to self-calibrate (*Franz and Triesch, 2007; Zhao et al.,*
47 *2012; Klimmasch et al., 2017; Eckmann et al., 2019*). These models have been developed in the
48 framework of Active Efficient Coding (AEC), which is an extension of Barlow's classic efficient coding
49 hypothesis to active perception (*Barlow, 1961*). In a nutshell, classic efficient coding argues that
50 sensory systems should use representations that remove redundancies from sensory signals to
51 encode them more efficiently. Therefore, sensory representations should be adapted to the statis-
52 tics of sensory signals. Based on this idea, a wide range of data on tuning properties of sensory
53 neurons in different modalities have been explained from a unified theoretical framework (*Dan*
54 *et al., 1996; Vinje and Gallant, 2000; Simoncelli, 2003; Smith and Lewicki, 2006; Doi et al., 2012*).
55 AEC goes beyond classic efficient coding by acknowledging that developing sensory systems shape
56 the statistics of sensory signals through their own behavior. This gives them a second route for
57 optimizing the encoding of sensory signals by adapting their behavior. In the case of binocular
58 vision, for example, the control of vergence eye movements is shaping the statistics of binocular
59 disparities. By simultaneously optimizing neural tuning properties and behavior, AEC models have
60 provided the first comprehensive account of how humans and other binocular species may self-
61 calibrate their binocular vision through the simultaneous learning of disparity tuning and vergence
62 control.

63 A generic AEC model has two components. The first component is an efficient coding model that
64 learns to encode sensory signals by adapting the tuning properties of a population of simulated
65 sensory neurons (*Olshausen et al., 1996; Olshausen and Field, 1997*). In the case of binocular
66 vision, this is a population of visual cortical neurons receiving input from the two eyes that learns to
67 encode the visual signals via an efficient code. The second component is a reinforcement learning
68 (RL) model that learns to control the behavior. In the case of binocular vision, this component
69 will learn to control eye vergence. For this, it receives as input the population activity of the visual
70 neurons and learns to map it onto vergence commands. This learning is guided by an internally
71 generated reward signal, which reinforces movements that lead to a more efficient encoding of
72 the current visual scene. For example, when the eyes are aligned on the same point, the left and
73 right images become largely redundant. The efficient coding model can exploit this redundant
74 structure in both eyes, by developing neurons tuned to small or zero disparities. Conversely, such
75 binocular neurons tuned to small disparities will represent any remaining misalignments of the
76 eyes, providing informative input for vergence control. In this way, learning of vergence control
77 supports the development of neurons tuned to small disparities and this developing population of
78 neurons in turn facilitates the learning of fine vergence control (*Zhao et al., 2012*).

79 Importantly, however, this normal development of binocular vision is impaired in a range of
80 alternate rearing conditions. In fact, already since the days of Hubel and Wiesel, alternate rearing
81 conditions have been used to improve our understanding of visual cortex plasticity and function.
82 Manipulating the input to the visual system during development and observing how the system
83 reacts to such manipulations has shaped our understanding of visual development until today.
84 For example, artificially inducing a strabismus leads to drastic changes in the tuning properties
85 of neurons in visual cortex (*Hubel and Wiesel, 1965*). A comprehensive theoretical account of the
86 development of binocular vision must therefore also be able to explain the experimentally observed
87 differences in alternate rearing conditions. Therefore, we here test if a recently proposed AEC model
88 of the development of binocular vision can reproduce and explain the large range of experimental
89 findings from different alternate rearing conditions. Indeed, we show that the model qualitatively
90 captures findings on how different alternate rearing conditions alter the statistics of disparity tuning
91 and binocularity. Furthermore, the model makes specific novel and testable predictions about




Figure 1. Model overview. **A** The agent looking at the object plane in the simulation environment. **B** Processing steps of the active efficient coding model. One image is generated per eye. We convolve them with different kernels, such as those in the inbox, to simulate alternate rearing conditions (N: normal, V: vertical, H: horizontal, O: orthogonal, M: monocular). Binocular patches are extracted in a coarse and a fine scale (green and orange boxes) with different resolutions. These patches are encoded by activations of basis functions via sparse coding and combined with the muscle activations to generate a state vector. While this vector is fed into the reinforcement learning architecture, the sparse coding step also generates a reconstruction error that indicates the efficiency of encoding. We use this signal as reward (purple arrow) to train the critic, which in turn evaluates states to teach the actor. Finally, the actor generates changes in muscle activations, which result in rotations of the eyeballs and a new iteration of the perception-action cycle.

92 differences in vergence behavior under the different rearing conditions. Surprisingly, it also predicts
93 systematic differences in the statistics of orientation tuning of visual cortical neurons depending
94 on the fidelity of vergence eye movements. Overall, our results support AEC as a parsimonious
95 account of the development of binocular vision, highlighting the active nature of this development.

96 Results

97 A model for the development of active binocular vision

98 The model comprises a virtual agent situated in a simulated environment. The agent looks at a
99 textured plane, on which images from the man-made section of the McGill Database (*Olmos and*
100 *Kingdom, 2004a*) are rendered. The plane is positioned in front the agent at variable distances
101 (Fig. 1A). An image is rendered for the left eye and a second image is rendered for the right eye.
102 Binocular patches are extracted from these images and encoded by a sparse coding algorithm. The
103 activation levels of the learned binocular basis functions (BFs) can be thought of as firing rates of
104 binocular simple cells in primary visual cortex. The basis functions themselves roughly describe their
105 receptive fields and are optimized through learning (*Olshausen and Field, 1997*). These activations
106 are then squared and pooled across the image to obtain a more position-invariant representations
107 mimicking the behavior of complex cells. From this state representation a reinforcement learner
108 generates vergence commands that symmetrically rotate the eyeballs inwards or outwards. This
109 results in two new images being rendered and a new simulation iteration starts. The complete
110 process is depicted in Fig. 1B (see Methods for details).

111 In the human retina, the receptive field (RF) size of ganglion cells increases towards the periphery
112 (*Curcio et al., 1990*). We incorporate this idea by extracting patches from an input image at two
113 different spacial scales: A high-resolution fine scale is extracted from the central part and a low-
114 resolution coarse scale is extracted from a larger area (orange and turquoise boxes in Fig. 1 and 2).
115 Covering a visual angle of 8.3° in total, the fine scale corresponds to the central/para-central region
116 (including the fovea) and the coarse scale to the near-peripheral region with a diameter of 26.6° .
117 On the one hand, this two-scale architecture is more biologically plausible than using just a single
118 scale, on the other hand it also increases the resulting verging performance (*Lonini et al., 2013*).
119 One input patch (or subfield) in the coarse scale can detect a disparity of up to 8.8° while one patch
120 in the fine scale covers 1.6° . The coarse scale can therefore be used to detect big disparities, while
121 the fine scale represents small disparities.

122 We simulate altered rearing conditions by convolving the input images for the two eyes with two-

123 dimensional Gaussian kernels to blur certain oriented edges, or to simulate monocular deprivation.
124 To mimic strabismus, the right eyeball is rotated inwards while the left eye remains unchanged to
125 enforce non-overlapping input to corresponding positions of the left and right retina (see *Methods*
126 for details).

127 The adaptation of the neural representation and the learning of appropriate motor commands
128 occur simultaneously: While the sparse coder updates the BFs to minimize the reconstruction error,
129 the RL agent generates vergence eye movements to minimize the reconstruction error of the sparse
130 coder. Since the sparse coder has a fixed capacity, minimizing its reconstruction error is equivalent
131 to maximizing its coding efficiency. Thus, both the sparse coder and the reinforcement learner aim
132 to maximize the overall coding efficiency of the model.

133 **Normal rearing conditions lead to the autonomous learning of accurate vergence** 134 **control for natural input and random dot stereograms**

135 Under normal rearing conditions the joint learning of the neural representation and motor behavior
136 results in an agent that accurately verges the eyes on the plane in front of it (*Klimmasch et al., 2017*).
137 This behavior emerges in an autonomous fashion, since both the sparse coder and the RL agent
138 only strive to improve the neural encoding by reducing the reconstruction error. We demonstrate
139 this behavior in Video 1 (videos/vergence_movements.mp4) and will analyse it in greater detail in
140 the following sections.

141 A critical test of any model of the development of stereoscopic vision is whether it can handle
142 *random-dot stereograms* (RDSs), which represent the most challenging stimuli for stereopsis. Since
143 their introduction by *Julesz (1971)* RDSs have been used extensively to investigate the human ability
144 for stereoscopic vision. Nowadays they are used in ophthalmological examinations to assess stereo
145 acuity as well as to detect malfunctions in the visual system, such as strabismus or amblyopia
146 (*Walraven, 1975; Okuda et al., 1977; Ruttum, 1988*). In these experiments participants view a grid
147 of random dots through a stereoscope or another form of dichoptic presentation. Typically, the
148 central part is shifted in one of the two images which results in the perception of stereoscopic depth
149 in healthy subjects. The advantage of this form of examination is that there are no monocular
150 depth cues (such as occlusion, relative size, or perspective). The impression of depths arises solely
151 because of the brain's ability to integrate information coming from the two eyes.

152 To show that our model is able to perceive depth in RDS, we generate various RDS and ren-
153 der the shifted images for the left and right eye separately. We expose the model that was
154 trained on natural input stimuli to a range of RDS with different spatial frequencies, window
155 sizes, disparities, and object distances. A video of the performance can be found in Video 2
156 (videos/performance_on_RDS.mp4). The model is clearly able to detect the differences in the
157 images and align the eyes on the virtual plane that will appear in front or behind the actual object
158 plane in the RDS. Averaged over all trials, the model achieves a vergence error of 0.2° . This is
159 comparable with our results on natural images (see Fig. 6) and indicates that the model generalizes
160 well to artificial images it has never seen before.

161 **Altered rearing conditions cause changes in neural representation matching ex-** 162 **perimental findings**

163 A second critical test of any model of the development of binocular vision is whether it can account
164 for the effects of alternate rearing conditions observed in biological experiments. We simulate such
165 alternate rearing conditions by filtering the input images for the left and right eyes with Gaussian
166 filters. Figure 2 shows illustrative examples of the filtered images that were used to train our model
167 and the respective learned BFs.

168 When the model is trained with unaltered natural visual input, the resulting RFs resemble Gabor
169 wavelets (*Daugman, 1985*), as shown in the first row in Fig. 2. They appear similar in the coarse
170 and the fine scale, but tend to be more localized in the latter. The changes that are applied to the
171 input images in the alternate rearing conditions are reflected in the RFs that are learned: Among

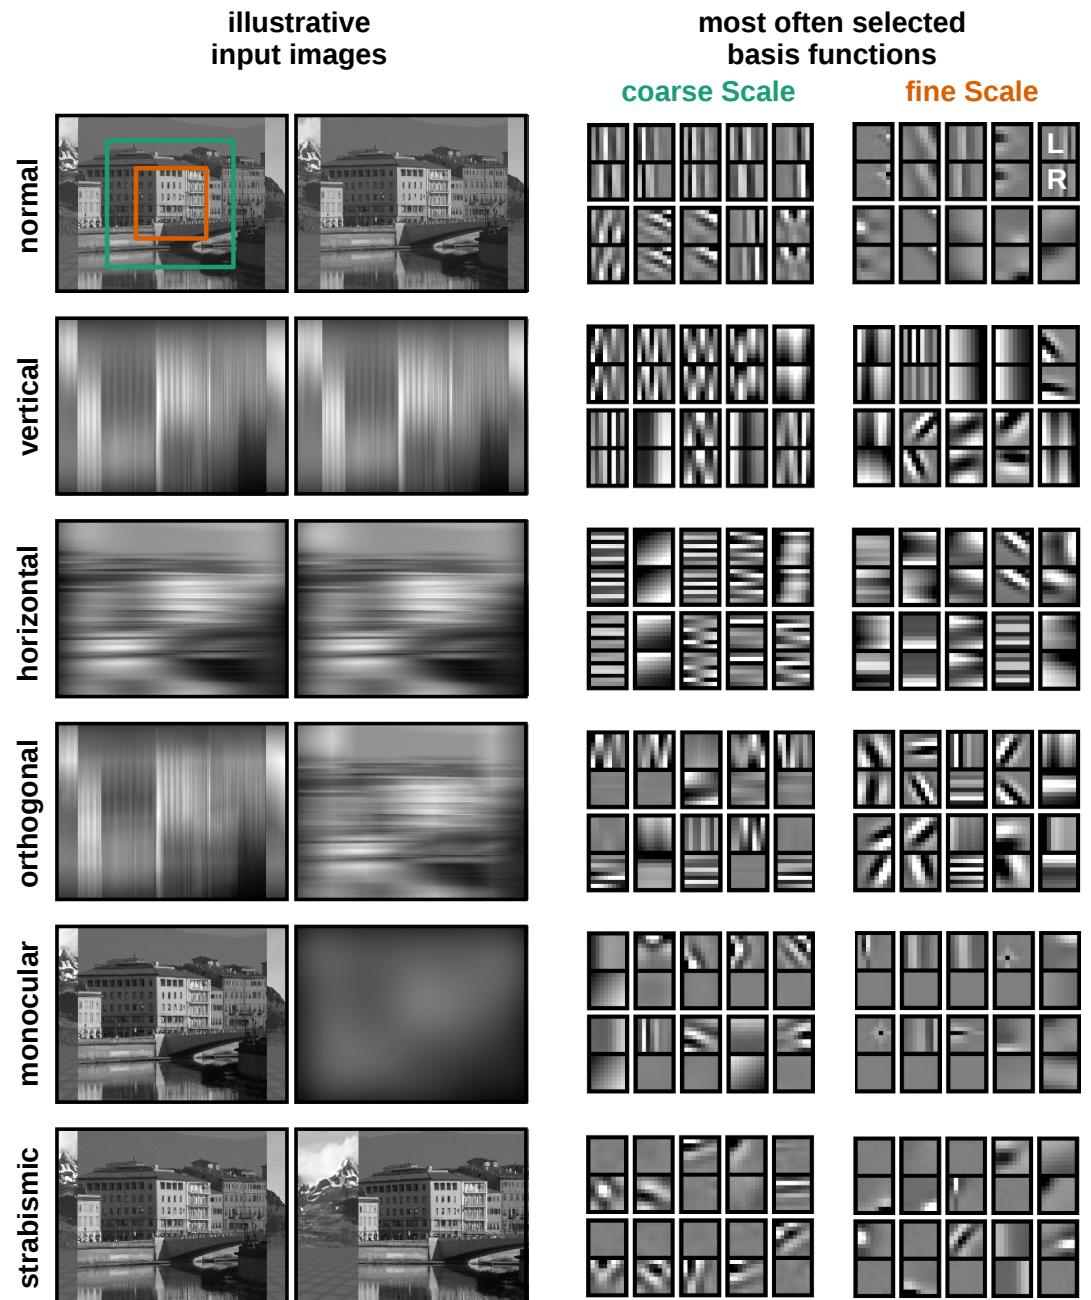


Figure 2. Input scenarios and learned receptive fields. Left: Illustration of the input under different rearing conditions. Except for the normal scenario, the images are convolved with different Gaussian filters to blur out certain orientations or simulate monocular deprivation. To simulate strabism the right eye is rotated inward by 10° , so that binocular neurons receive non-corresponding inputs to their left and right eye receptive fields. **Right:** Representative examples of binocular basis functions (BFs) for the fine and coarse scale learned under the different rearing conditions after 0.5 million iterations. For each BF the left eye and right eye patch are aligned vertically. In each case, the 10 BFs selected most frequently by the sparse coding algorithm are shown.

172 the 10 most often selected BFs there are no vertically (horizontally) oriented RFs, when the model
 173 is trained on images that are deprived of vertical (horizontal) edges. Orthogonal RFs emerge as
 174 a result of training on orthogonal input. When one eye is deprived of input, the RFs will become
 175 *monocular* and encode information coming from the “healthy” eye only. Strabismic rearing results
 176 in the development of monocular RFs without a preference for one or the other eye (*Hunt et al.,*

177 **2013).**

178 The full set of all BFs (coarse and fine scale) for all the rearing conditions can be found in
179 supplemental Fig. 1. The following sections will analyse them in a more quantitative fashion.

180

181 **Neurons' orientation tuning reflects input statistics**

182 To analyze the statistics of the developing RFs in greater detail, we fit oriented two-dimensional
183 Gabor wavelets to each BF. For this part of the analysis the left and right parts of the binocular
184 BFs are studied separately, so we look at the *monocular* BF fits only. We combine the results from
185 coarse and fine scale, since a two-sample Kolmogorov-Smirnov test (*Young, 1977*) did not reveal a
186 statistically significant difference between the distributions. Only those BFs which met a criterion
187 for a sufficiently good fit (98% of all bases) are considered for further analysis (see *Methods*).

188 Figure 3 shows how the altered input changes the distribution of preferred orientations of the
189 BFs. In the *normal* case we can observe a clear over-representation of vertically (0°) and horizontally
190 (90°) tuned BFs. This is known as the *oblique effect* and has been frequently observed in animals
191 (*Appelle, 1972; Li et al., 2003*) and humans (*Furmanski and Engel, 2000*). It has been argued that it
192 stems from the over-representation of vertical and horizontal edges in natural images (*Coppola
193 et al., 1998*). Additionally, we cannot exclude the possibility that it is related to the rectangular pixel
194 grid representing the input to our model.

195 While the distribution of orientations does not change much in the *monocular* and *strabismic*
196 rearing case, we observe a marked difference to the normal case when certain orientations are
197 attenuated in the input. The models trained on *vertical* input are missing the peak at horizontal
198 orientations and vice versa for the *horizontal* case. Additionally, we find an increased number of
199 neurons tuned to the dominant orientation in the input. These observations are consistent with
200 animal studies (*Stryker et al., 1978; Tanaka et al., 2006*)

201 The separate analysis of the RFs in the left and right eye for the models that were trained on
202 *orthogonal* input reveals the adaptation of each eye to its input statistics. Furthermore, we find that
203 orthogonal RFs developed (also see fourth row in Fig. 2) that have been observed in an orthogonal
204 rearing study in cats (*Leventhal and Hirsch, 1975*).

205

206 **The development of binocular receptive fields requires congruent input to the two 207 eyes**

208 Another interesting feature of the neural representation that has been studied extensively in
209 the context of alternate rearing is the *binocularity*. The binocularity index (BI) is used to assess how
210 *responsive* a neuron is to the inputs from the two eyes. A *binocular* neuron requires input from both
211 eyes to respond maximally, while a *monocular* neuron is mostly driven by just one eye.

212 To determine the binocularity indices for the neurons in our model we use the original method
213 from *Hubel and Wiesel (1962)*. They determined a stimulus that maximizes the monocular response,
214 and applying this stimulus separately in left or right eye to get the (monocular) neural responses L
215 and R (see *Methods* for details). For each neuron the binocularity index is then calculated as $\frac{R-L}{R+L}$.
216 Like Hubel and Wiesel we sort the binocularity indices into seven bins. The values range from -1
217 (monocular left) over 0 (binocular) to +1 (monocular right).

218 Figure 4 depicts the binocularity distributions for the coarse and the fine scale for all rearing
219 conditions. The models that were trained on input that is coherent between the left and right eye
220 (top row) exhibit the majority of neurons falling in the bin with binocularity index 0. Neurons in this
221 category receive about the same drive from the left and the right eye. In the *normal* case more
222 neurons fall into that bin than in the *vertical* and *horizontal* case. This is due to the ability of the
223 model to perform precise vergence control: Since left and right image are almost identical most of
224 the time, the great majority of basis functions will develop to encode the exact same input from
225 both eyes. This, in turn, will result in the cells being completely binocular with a binocularity index
226 of 0. In the vertical and horizontal case, we observe a reduction in the number of cells that have a




Figure 3. Orientation distributions for different rearing conditions. Displayed are the orientations of the Gabor wavelets that were fitted to the learned BFs of the left eye. Shown are the best fits from coarse and fine scale combined (800 in total). The error bars indicate the standard deviation over 5 different simulations. \bar{N} describes the average number of BFs that passed the selection criterion for their fits (see Methods).

227 binocularity index of 0. We attribute this to the limited vergence performance in these cases, that
 228 we will analyse in the next sections.

229 If, on the other hand, the input differs qualitatively for the two eyes (Fig. 4, bottom row) the
 230 receptive fields will also differ in their monocular sub-parts. This can also be observed in Fig. 2 for
 231 the orthogonal, monocular and strabismic case. Looking at the binocularity index, we find that
 232 most of the cells become monocular, with a symmetric distribution for orthogonal and strabismic
 233 rearing. Monocular deprivation of the right eye leads to a distribution of binocularity indices that is
 234 mostly monocular for the left eye.

235 We also find differences between coarse and fine scale, with slightly fewer binocular and slightly
 236 more monocular cells in the latter one. This indicates that left and right part of the BFs in the fine
 237 scale tend to be marginally more different than in the coarse scale. Patches that serve as input to
 238 this scale are not down-sampled and have a high resolution. Small differences in the input patches
 239 will therefore not be blurred out and lead to small differences in the learned BFs since the sparse
 240 coder strives for reconstructing the input as accurately as possible.

241 Looking into the biological data, we find the pronounced peak at binocular neurons in the
 242 normal case (*Wiesel and Hubel (1963)*, Fig. 1, and *Hubel and Wiesel (1965)*, Fig. 5). When trained on
 243 inputs deprived of certain orientations (*Stryker et al. (1978)*, Fig. 6B), the neurons become relatively
 244 more monocular, but most of the neurons remain binocular. This is in good agreement with our
 245 model.

246 *Stryker et al. (1978)* reared kittens on orthogonal input and report an increase in monocular
 247 neurons (Fig. 6A) when compared to the normal rearing data from Hubel and Wiesel. In comparison
 248 to the rearing on stripes, there are fewer binocular cells. The loss of binocular neurons that we see
 249 in our data is also reported in *Hirsch and Spinelli (1970)*, who reared kittens on orthogonal stripes.

250 Monocular rearing and the analysis of binocularity was performed in *Wiesel and Hubel (1963)*.




Figure 4. Binocularity distributions for different rearing conditions. The binocularity index is calculated by comparing the neuron's responses to monocular stimuli. The values range from -1 (monocular left) over 0 (binocular) to 1 (monocular right). Results for coarse and fine scale are presented next to each other. Error bars indicate the standard deviation over 5 different simulations. \bar{N}_c and \bar{N}_f are the average number of basis functions (out of a total of 400) that pass the selection criteria for their fits (see Methods).

251 In Fig. 3 and 5 we see the development of completely monocular cells after visual deprivation of the
 252 other eye. The strabismic case was studied a few years later in *Hubel and Wiesel (1965)* (Fig. 5A)
 253 and revealed a division of the neural population in monocular neurons for either left or right eye, in
 254 agreement with our model.

255 **Alternate rearing conditions reduce the number of disparity tuned cells**

256 A central aspect of the development of binocular vision is the emergence of neurons which are
 257 tuned to binocular disparities. We therefore investigate how alternate rearing affects the number
 258 of neurons with disparity tuning in the model and the distribution of their preferred disparities. We
 259 estimate disparity tuning by considering phase shifts between left and right RFs in the following
 260 way: We fit binocular Gabor wavelets to the BFs, where all parameters, except for the phase shift,
 261 are enforced to be identical for the left and right monocular BF. The disparity for one neuron can
 262 then be calculated as described in *Analysis of receptive fields*. The distribution of disparity tuning of
 263 the coarse scale neurons is shown in Fig. 5 for the different rearing conditions. Results for the fine
 264 scale are comparable and presented in supplemental Fig. 2. First, there is a noticeable difference in
 265 the number of cells that are disparity tuned between the different rearing conditions: In the normal
 266 case we find the highest number of disparity tuned cells, rearing in a striped environment reduces
 267 the number, and uncorrelated input results in the smallest number of disparity tuned cells. In every
 268 case, the distribution of preferred disparities is peaked at zero. The height of this peak is reduced
 269 for rearing conditions with in-congruent input to the two eyes.
 270

271 Comparing the normal with the vertical and horizontal case, there is an increase in the number
 272 of cells that are tuned to non-zero disparities. This indicates that under these alternate rearing
 273 condition, the agents are exposed to non-zero disparities more often. This is in good agreement




Figure 5. Disparity distributions for different rearing conditions. The neuron's preferred disparities are extracted from the binocular Gabor fits. All neurons with a disparity bigger than the maximally detectable one are removed from the analysis. Presented are the averaged data of the coarse scale from 5 random seeds. \bar{N} describes the average number of neurons that met the selection criteria (see Methods).

274 with the results from the next section (also see Fig. 6), where we will see that those models perform
 275 less accurate vergence movements compared to the normal case.

276 In the strabismic case, a neuron's receptive fields in left and right eye are driven by un-
 277 corresponding input. This results in very few disparity tuned cells that exhibit a much broader
 278 distribution of preferred disparities.

279 To investigate the effect of a less severe strabism we conduct an additional experiment similar
 280 to *Shlaer (1971)* (see Fig. 2). Here, we fix the strabismic angle to 3° , which results in a corresponding
 281 image in the two eyes because one input patch in the coarse scale covers an angle of 6.4° . *Sup-*
 282 *plemental Fig. 3* shows that this leads to an increased amount of disparity tuned cells and a shift
 283 of their preferred disparity to 3° . Exactly as in *Shlaer (1971)*, the constant exposure to a certain
 284 disparity leads to a preference to that disparity for the majority of cells.

285 **Model predicts how alternate rearing conditions affect vergence learning**

286 While the effect of alternate rearing conditions on receptive fields of visual cortical neurons is well
 287 studied, there has been little research on the effect of alternate rearing conditions on vergence
 288 behavior. To quantify vergence behavior in the model, we define the absolute vergence error. It
 289 measures by how much the vergence angle between the eyes deviates from the ideal position,
 290 which would make the two eyes fixate the same center of the object. This measurement is taken at
 291 the end of a fixation (corresponding to the last of 10 time steps), to give the model sufficient time
 292 to fixate the object.

293 Figure 6A shows the evolution of the absolute vergence error over training time for the different
 294 rearing conditions. The models with *normal* or *vertical* rearing learn to verge the eyes on the same
 295 point on the object, resulting in the reduction of the vergence error to small values of around 0.3
 296 degrees. The model that learns on images without vertical edges (*horizontal* case) does manage to
 297 verge the eyes slightly, but does not reach the accuracy of the former models. The *orthogonally*,




Figure 6. Vergence performance of models raised under different rearing conditions. A Moving average of the vergence error over the duration of the training period. During the training, a textured object plane is positioned in front of the agent at varying distances. The vergence error is defined as the difference between the angle that is desired to fixate exactly on the object plane and the actual angle between the eyes. The shadows indicate the standard deviation over 5 different random seeds. **B** Vergence errors after 20 perception-action-cycles on unknown input stimuli starting from various initial errors. This testing was done without the visual aberrations encountered during training. Displayed are conventional box plots without outliers. The gray bar indicates a vergence error of 0.2° which presents the resolution boundary of our system.

298 *monocularly* and *strabismically* reared models do not improve much in comparison to random
299 behavior in the beginning of training. The main difference to the models that were able to learn
300 vergence is that under these conditions the left and right eye are provided with in-congruent input.
301 The orthogonal model receives two monocular images that retain different orientations. The right
302 monocular image of the monocularly deprived model contains little information at all, and the two
303 eyes are physically prevented from looking at the same object in the strabismic case. In these cases,
304 very few neurons with disparity tuning emerge (compare previous section) that could drive accurate
305 vergence eye movements.

306

307 **Behavior after normal visual input is reinstated**

308 Alterations of the visual input during the critical period of visual development lead to lasting
309 visual deficits. To simulate the effect of a transient alteration of visual input during the critical
310 period, we first train the model under alternate rearing conditions as described above and then
311 reinstate normal visual input. For this, we freeze all weights after the training phase and test all
312 models on the same, un-altered input images. By doing so, we simulate a situation where the visual
313 aberrations present during development (such as astigmatism or a cataract) are corrected *after* the
314 critical period.

315 The object plane is put to a distance $\in [0.5, 1, \dots, 6]$ m, the initial vergence error is chosen
316 randomly from -2 to 2° , and 40 stimuli that were not seen during training are applied on the object
317 plane. From these initial conditions we run the simulation for 20 iterations and record the vergence
318 error at the end of each fixation. The results of this testing procedure are displayed in Fig. 6B. Here,
319 the gray shaded area indicates a vergence error of 1 pixel. We observe that the *normally* trained
320 model exhibits the best performance and actually achieves sub-pixel accuracy in the great majority
321 of trials. Interestingly, the performance declines for the *vertical* model. One could expect the model
322 that was trained solely on vertical edges to be better at aligning those edges. We attribute this to
323 mis-alignments (or false matches) between the two images that happen more frequently, when the
324 world is made up only of vertical edges. Additionally, the neural representation that was learned
325 during the exposure to vertical edges only might not be utilized as efficiently as before, now that all
326 orientations are present in the input.

327 Even though the performance of the model trained on only *horizontal* orientations is quite poor

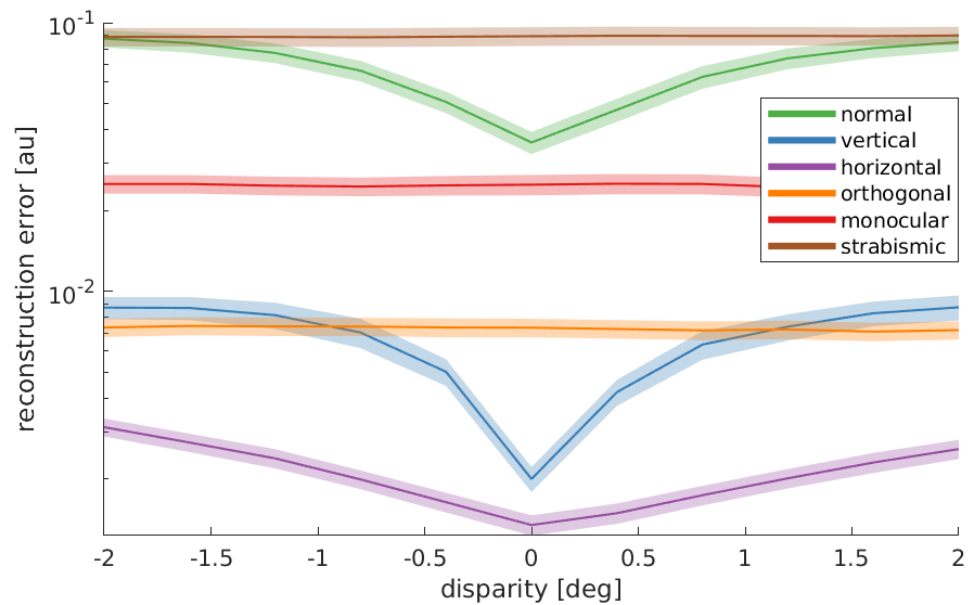


Figure 7. Reward functions for the different rearing conditions. The reward function is what drives the reinforcement learner to move the eyes in a useful fashion. For all different conditions, we plot the rewards that the models will receive at different disparities. Notice the log-scale on the y-axis. The data are averaged over 10 stimuli that were not encountered during training, three different object distances (0.5, 3, and 6 m), and 5 random seeds for every condition. The shaded area represents the standard error. Only those models that receive corresponding input in left and right eye display a reconstruction error that is minimal at zero disparity. These are the only models that learn to verge the eyes.

328 during training, after applying the correction it clearly displays a verging behavior. In comparison to
329 the *orthogonal*, *monocular* and *strabismic* models, it reduces the vergence error, though being less
330 accurate than the other two cases.

331 The main difference between the conditions under which vergence could or could not be learned
332 is the correspondence between the input images. When the inputs to the two eyes are in-congruent
333 — as in the *orthogonal*, *monocular* and *strabismic* cases — we could not observe any improvement
334 in the vergence error. Matching input, on the other hand, always led to the learning of vergence
335 behavior. This becomes apparent especially after testing the learned models on un-altered inputs.

336 Since this is the first study to investigate the quality of learned vergence movements after expo-
337 sure to alternate rearing conditions (to the best of our knowledge), the differences in performance
338 are a genuine prediction of our model.

339 To explain this phenomenon we consider the reward function, which shapes the model's behav-
340 ior during training. Figure 7 shows the averaged reconstruction error over three different object
341 distances and ten stimuli for the different rearing conditions. We defined the reward as the negative
342 reconstruction error of the sparse coders. In the normal case, we clearly see an optimum of the
343 reconstruction error at zero disparity. This also holds for the vertical and horizontal condition,
344 whereas those are at least one magnitude smaller. We argue that the differences in the rewards
345 lead to the differences in vergence performance, since all models that could not verge display a
346 reward function that is rather flat for different disparity values. The models with a negative peak
347 at zero disparity, on the other hand, all learn to verge and the difference in the magnitude of the
348 reward seems to be reflected in the vergence performance after training.

349 **Model predicts how vergence movements influence the statistics of orientation preference** 350 preference

351 Our model also allows us to investigate, for the first time, how the quality of the vergence control
352 influences the neural representation. As a baseline, we consider the orientation tuning of a

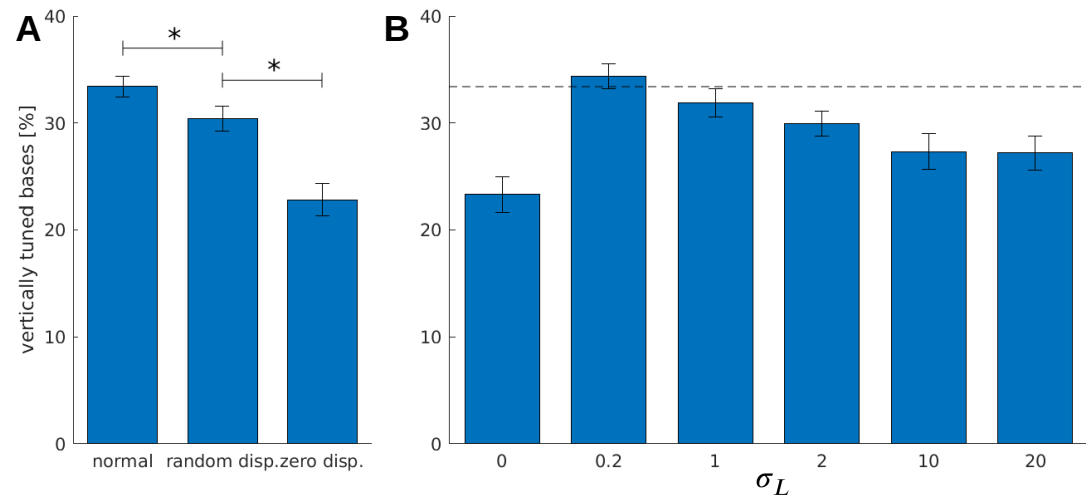


Figure 8. The effect on the learning of vergence and disparities on the number of neurons tuned to vertical edges. **A** Here we compare the relative amount of cells that are tuned to vertical orientations for three different types of models: The first is the normal model, where the sparse coder and the RL agent are trained. In the second case, only the sparse coder is trained while the RL agent is removed. During the training procedure, this model is exposed to random disparities in the input images. In the last case, only the sparse coder is trained, but additionally, we artificially set the eyes to always verge perfectly on the objects in front of it. Like that, this model does not learn vergence movements and is not exposed to disparities as well. Asterisks indicate a statistically significant difference between the samples as revealed by the students t-test (p-values are 0.007 and 0.001). **B** These models were trained with a Laplacian distribution of different disparities. Depicted are the relative amount of BFs tuned to vertical orientations in dependence of σ_L , the standard deviation of the Laplacian. $\sigma_L = 0$ corresponds to 0 disparity all the time, while $\sigma_L = 20$ is an almost uniform disparity distribution. Error bars indicate the standard deviation over 5 different seeds. The black dotted line indicates the amount of vertically tuned neurons in the *normal* model.

353 reference model which was trained on normal images and learned an appropriate vergence policy.
354 We compare this model to a version that was trained on the same input images, but could not
355 verge the eyes. Specifically, the sparse coder was trained normally, but the RL part was removed.
356 This model saw different disparities during training by looking at objects at different depths, but
357 was not able to change this distribution of disparities to facilitate the encoding. We refer to this
358 model as the “random disparity” model. In another version of the model, we artificially always
359 set the vergence angle to correctly fixate the objects. In this way, this model was never exposed
360 to disparities (except for very small ones in the periphery that arise because of slightly different
361 perspectives in the left and right eye). We refer to this version as the “zero disparity” model.

362 Figure 8A shows the fraction of neurons that are tuned to vertical orientations ($0 \pm 15^\circ$) for these
363 three models. When the influence of the RL agent is removed, we observe a significant decrease in
364 the number of vertically tuned neurons. This change must be caused by the different distributions
365 of disparities that the models experience due to their different motor behavior. In the model that
366 was trained without disparities, we find the least amount of neurons tuned to vertical edges.

367 To study the role of the distribution of experienced disparities more systematically, we train
368 the sparse coder on different truncated Laplacian distributions of disparities. The distributions
369 are heavy-tailed and centered around zero. The spread in this distribution is determined by the
370 standard deviation σ_L . $\sigma_L = 0$ means zero disparity all the time (corresponding to the zero disparity
371 case), while the distribution becomes almost uniform for big values of σ_L . Figure 8B shows how
372 the number of vertically tuned neurons changes in response to different values of σ_L . We find
373 the smallest number of vertically tuned cells when the disparity is zero throughout the whole training.
374 For very large σ_L there are more vertical cells, but not as many as for smaller values which are
375 different from zero. In fact for $\sigma_L = 0.2$, which corresponds to a standard deviation of one pixel in




Figure 9. Intuition for the over-representation of vertical edges when disparities have to be encoded. We show the location of two RFs (yellow and cyan circles) on a patch in the visual field and present them with three different disparities. The inputs are depicted as anaglyphs, compositions of two images where the left image goes into the green channel and the right into a magenta channel. When the two images are corresponding, the anaglyph will appear in black and white, while un-corresponding parts will appear in green and magenta. For each disparity and RF we show the BF that is selected by the sparse coder to encode the input. While the BF that encodes the input in the cyan RF is the same for all disparities, the input inside the yellow RF can best be reconstructed by BFs that are tuned to that exact disparity.

376 the input image, the number of vertically tuned neurons is maximized.

377 An intuitive explanation for this over-representation of cells tuned to vertical orientations is given
378 in Fig. 9. Here, we depict a part of an input image at three different disparities. While the horizontal
379 edge can be encoded by the same BF for all disparity values, the vertical edge demands three
380 different basis functions to represent the input pattern faithfully. A system that experiences these
381 disparities in its inputs, needs to devote neural resources to represent them all. If the distribution
382 of disparities becomes too wide, however, individual neurons will receive close to independent
383 input from both eyes and disparities that lie in the range that can be represented by a single basis
384 function will be rare.

385 Discussion

386 A major goal of Computational Neuroscience is the development of models that explain how the
387 tuning properties of neurons develop and how they contribute to the behavior of the organism.
388 Over the last decades, the dominant theoretical framework for understanding the development of
389 tuning properties of sensory neurons has been the *efficient coding hypothesis*. It states that sensory
390 tuning properties adapt to the statistics of the sensory signals. In this framework, the behavior of

391 the organism has been largely neglected, however. Specifically, there has been hardly any work on
392 how developing neural tuning properties shape behavior, how the developing behavior affects the
393 statistics of sensory signals, and how these changing statistics feed back on neural tuning properties.
394 We argue that understanding the development of sensory systems requires understanding this
395 feedback cycle between the statistics of sensory signals, neural tuning properties and behavior.

396 The *active efficient coding* (AEC) approach offered here extends classic theories of efficient coding
397 by a behavior component to study this feedback cycle in detail. Here we have focused on active
398 binocular vision, where a simulated agent autonomously learns to fixate a target object with both
399 eyes via vergence eye movements. All parts of our model self-organize in tandem to optimize
400 overall coding efficiency. We have shown that that our model can autonomously self-calibrate
401 and even perform accurate vergence on random dot stereograms, despite having never been
402 exposed to such stimuli. In addition, we have reproduced various animal studies on alternate
403 rearing conditions, which often show dramatic effects on neural representations and behavior.
404 Our simulation results are in qualitative agreement with experimental findings, lending additional
405 support to our model. Beyond explaining a range of experimental findings, our model also predicts
406 systematic changes in the learned vergence behavior in response to altered regarding conditions.
407 In addition, the model predicts that the learning of accurate vergence behavior systematically
408 influences the neural representation and offers a novel explanation for why vertical orientations
409 tend to be over-represented in visual cortex compared to horizontal ones, at least in primates
410 (*De Valois et al., 1982b*) and humans (*Yacoub et al., 2008; Sun et al., 2012*). These predictions
411 should be tested in future experiments. For example, in enucleated animals, a bias in favor of
412 vertical orientations over horizontal ones may be reduced or completely absent (*Fregnac et al.,*
413 *1981*).

414 By freezing the neural network after the training period, we also simulated the state of the
415 brain after the critical period. Even after fixing the optical aberrations present during training we
416 observed a reduced vergence performance for all alternate rearing regimes. This finding is in line
417 with a large body of evidence suggesting that optical aberrations should be corrected as early as
418 possible to facilitate healthy development of binocular vision (e.g. *Daw (1998); Fawcett et al. (2005),*
419 *but also see Ding and Levi (2011)*).

420 While our results qualitatively match experimental findings, there are some quantitative dif-
421 ferences. In particular, while the distribution of binocularity indices (*Wiesel and Hubel, 1963*) and
422 disparities (*Sprague et al., 2015*) in healthy animals are relatively broad (*De Valois et al., 1982a;*
423 *Stevenson et al., 1992; Ringach et al., 1997*), we find more narrow ones in our model. These differ-
424 ences are likely due to a number of simplifications present in our model. In the brain, inputs from
425 both eyes into primary visual cortex are organized into ocular dominance bands such that individual
426 cortical neurons may receive input which is already biased towards one or the other eye (*Le Vay*
427 *et al., 1980; Crowley and Katz, 2000*). In contrast, in our model all neurons receive similar amounts
428 of input from both eyes and are therefore already predisposed for becoming binocular cells. This
429 might explain the model's narrower distribution of binocularity indices. Regarding the distribution
430 of preferred disparities, animals raised under natural conditions will experience a broad range of
431 disparities in different parts of the visual field, since objects in different locations will be at different
432 distances. In the model, the visual input is quite impoverished, as it is dominated by a single large
433 frontoparallel textured plane. Once this plane is accurately fixated, most parts of the visual field will
434 appear at close to zero disparity. This may explain the narrower distribution of preferred disparities
435 observed in the model.

436 Similarly, the distribution of preferred orientations in our model shows a very strong preference
437 for horizontal and vertical, that is accentuated relative to the normal oblique effect (*Li et al., 2003;*
438 *De Valois et al., 1982b*). Possible reasons for this include the discrete, rectangular pixel grid with
439 which visual inputs are sampled, the choice of our image data base (*Olmos and Kingdom, 2004a*),
440 which contains mostly man-made structures including buildings, etc., for which it is known that
441 they contain an abundance of horizontal and vertical edges (*Coppola et al., 1998*), and the model's

442 restriction to the central portion of the visual field, where the oblique effect is more pronounced
443 (*Rothkopf et al., 2009*).

444 Next to addressing the above limitations, an interesting topic for future work is to use the model
445 to study the development of amblyopia. For this, we have recently incorporated an interocular
446 suppression mechanism, since suppression is considered a central mechanism in the develop-
447 ment of amblyopia (*Eckmann et al., 2019*). Such models could be a useful tool for predicting the
448 effectiveness of novel treatment methods (*Papageorgiou et al., 2019; Gopal et al., 2019*).

449 In conclusion, we have presented a computational model that sheds new light on the central role
450 of behavior in the development of binocular vision. The model highlights how stimulus statistics,
451 sensory representation and behavior are all inter-dependent and influence one another and how
452 alternate rearing conditions affect every aspect of this system. The Active Efficient Coding approach
453 pursued here may be suitable for studying various other sensory modalities across species.

454 Methods

455 The following paragraphs will describe the different stages of the model, the experimental setup,
456 and the analysis. The implementation is publicly available at <https://github.com/Klimmasch/AEC/>.

457 Image processing

458 We use OpenEyeSim (*Priamikov and Triesch, 2014; Priamikov et al., 2016*) to render the left and
459 right eye image. It comprises a detailed biomechanical model of the human oculomotor system and
460 simulates a 3-dimensional environment. A rectangular plane is moved in front of the learning agent
461 (perpendicular to the gaze direction). On it we apply greyscale textures from the McGill Calibrated
462 Color Image Database (*Olmos and Kingdom, 2004b*) to simulate objects at different depths.

463 The two monocular images rendered by OpenEyeSim cover a horizontal field of view of 50° and
464 have $320 \text{ px} \times 240 \text{ px}$ (focal length $F = 257.34 \text{ px}$). We use Matlab to extract single patches in different
465 resolutions and combine corresponding patches from the left and right image. These binocular
466 patches will be jointly encoded by the sparse coder. The *coarse scale* corresponds to $128 \text{ px} \times 128 \text{ px}$
467 in the original image (corresponds to $26.6^\circ \times 26.6^\circ$) and is down-sampled by a factor of 4 to $32 \text{ px} \times 32 \text{ px}$.
468 The *fine scale* image corresponds to $40 \text{ px} \times 40 \text{ px}$ ($8.3^\circ \times 8.3^\circ$) and is not down-sampled. From coarse
469 and fine scale we extract $8 \text{ px} \times 8 \text{ px}$ patches with a stride of 4 px and combine corresponding left
470 and right patches to $16 \text{ px} \times 8 \text{ px}$ binocular patches (see Fig. 1). One patch in the coarse scale covers
471 a visual angle of 6.6° and in the fine scale one patch covers 1.6° . In total, we generate 81 fine scale
472 and 49 coarse scale patches that are subsequently normalized to have zero mean and unit norm.

473 Sparse coding

474 The patches from coarse and fine scale are used in the sparse coding step to construct a neural
475 representation of the visual input and to generate a reward signal that indicates the efficiency
476 of this encoding. Each scale $S \in \{c, f\}$ comprises a dictionary of binocular basis functions (BF)
477 $\phi_{S,i} \in \mathcal{B}_S$ with $|\mathcal{B}_S| = 400$. Each patch $p_{S,j}$ is reconstructed by a sparse linear combination of 10 BF:

$$\hat{p}_{S,j} = \sum_{i=1}^{|\mathcal{B}_S|} \kappa_{S,i}^j \phi_{S,i}, \quad (1)$$

478 where the vector of activations κ_S^j is allowed to have only 10 non-zero entries. The κ_S^j are chosen
479 by *matching pursuit* (*Mallat and Zhang, 1993*). This greedy algorithm selects the 10 BF from the
480 respective dictionary that yield the best approximation $\hat{p}_{S,j}$ of a patch.

481 The reconstruction error E_S is calculated as the sum over all squared differences between all
482 patches and their approximations, normalized by the total energy in the input patches:

$$E_S = \sum_{j=1}^{|p_S|} \frac{\|p_{S,j} - \hat{p}_{S,j}\|^2}{\|p_{S,j}\|^2}. \quad (2)$$

483 The total reconstruction error $E = E_c + E_f$ is used as the negative reward (see following section)
 484 while the errors for each scale are used to update the BF (*Olshausen et al., 1996*).

485 The state representation is given by a feature vector, where every entry describes the mean
 486 squared activation of one BF over the whole input image:

$$F_{S,i} = \sum_{j=1}^{|p_S|} \frac{(\kappa_{S,i}^j)^2}{|p_S|} . \quad (3)$$

487 Taken together, this feature vector F has $2|B_S|$ entries for both scales combined.

488 **Generation of motor commands**

489 The angular position of the eyes are controlled by two extra-ocular eye muscles responsible for
 490 rotations around the vertical axis. This *medial* and *lateral rectus* are simulated utilizing an elaborate
 491 muscle model (*Umberger et al., 2003*) inside OpenEyeSim (*Priamikov and Triesch, 2014; Priamikov*
 492 *et al., 2016*). Since we are interested in vergence movements only, we assume symmetrical eye
 493 movements so that the activities of the two muscles are mirrored for both eyes.

494 To generate those activations (between [0, 1] in arbitrary units) we use reinforcement learn-
 495 ing (*Sutton and Barto, 1998*). Specifically, the model employs the CACLA+VAR algorithm from
 496 *Van Hasselt and Wiering (2007)* that generates outputs in continuous action space. In short, it
 497 uses an actor-critic architecture (*Grondman et al., 2012*), where the actor and critic use neural
 498 networks as function approximators. These neural networks receive the state vector s_t that is the
 499 concatenation of the BF activations from both scales (see previous section) and the current muscle
 500 innervations. The entries in s_t are scaled by Welford's algorithm (*Welford, 1962*) to have zero mean
 501 and a fixed standard deviation (0.02 in our simulations).

502 The critic is a one-layer network that aims to learn the value of a state. From the state vector it
 503 approximates the discounted sum of all future rewards

$$V(s_t) = \sum_{i=0}^{\infty} \gamma^i r_{t+i} , \quad (4)$$

504 where r_t represents the reward achieved at time t and γ is the discount factor. To update this
 505 value network, we calculate the *Temporal Difference Error* (*Tesauro, 1995; Sutton and Barto, 1998*)
 506 as $\delta_t = r_t + \gamma V_t(s_{t+1}) - V_t(s_t)$. The parameters of the critic, θ^V , are updated by

$$\Delta \theta_{i,t}^V = \alpha \delta_t \frac{\partial V_t(s_t)}{\partial \theta_{i,t}^V} , \quad (5)$$

507 where α represents the learning rate for updating the critic.

The actor is a two layer artificial neural network with 50 hidden units (with \tanh activation functions) and a two-dimensional output that depicts changes in muscle innervation for the two relevant eye muscles (lateral and medial rectus). The generated motor outputs are random in the beginning and the network is updated whenever the given reward was higher than estimated by the critic:

$$\text{IF } \delta_t > 0 : \quad \Delta \theta_{i,t}^A = \beta (a_t - A_t(s_t)) \frac{\partial A_t(s_t)}{\partial \theta_{i,t}^A} \left[\frac{\delta_t}{\sqrt{\text{var}_t}} \right] , \quad (6)$$

508 where β is the actor's learning rate, $A_t(s_t)$ is the action selected by the actor at time t , and
 509 $a_t = A_t(s_t) + \mathcal{N}(0, \sigma^2)$ is the action that is actually executed. Adding Gaussian noise to the actor's
 510 output to discover more favourable actions is called *Gaussian exploration*. The last term scales the
 511 update depending on how much better the action was than expected with respect to its standard
 512 deviation.

513 **Simulation of alternate rearing conditions**

514 The deprivation of oriented edges is simulated by convolving the input images with elongated
515 Gaussian kernels defined by:

$$K_{\sigma_x, \sigma_y}(x, y) = \exp\left(-\left(\frac{x^2}{2\sigma_x^2} + \frac{y^2}{2\sigma_y^2}\right)\right), \quad (7)$$

516 where $\sigma_{x/y}$ represent the standard deviation in the horizontal/vertical direction.

517 Kernels with a large σ_x (σ_y) will blur out vertical (horizontal) edges. Specifically, to simulate the
518 deprivation of horizontal orientations, σ_x is set to 33 px (to cover one patch in the coarse scale
519 completely) and σ_y to a small value of 0.1 px. The numbers are reversed for the deprivation of
520 vertical orientations. In the case of orthogonal rearing, the left eye receives an image deprived of
521 horizontal orientations while the right eye receives one without vertical orientations. To make up
522 for the small standard deviation of 0.1 in the direction that should not be impaired, the images in
523 the *normal* case are convolved with a Gaussian kernel with $\sigma_x = \sigma_y = 0.1$ px.

524 To simulate monocular deprivation (MD) we set $\sigma_x = \sigma_y = 240$ px for the right input image only.
525 The small patches that we extract from this strongly blurred image contain almost no high spatial
526 frequencies.

527 A strabismus is artificially induced by rotating the right eye ball inwards as it is commonly done
528 in biological experiments by fixating a prism in front of the eye or by cutting the lateral rectus
529 muscle. In our Open-Eye-Simulator, however, we can rotate the eye by a specific angle. One input
530 patch in the coarse scale covers 6.6°. When we set the strabismic angle to 3° there is still an overlap
531 in the input images that will be reflected in the neural code. In contrast, when the strabismic angle
532 is set to 10°, the input patches become completely uncorrelated. Examples of the changes done to
533 the input images are displayed in Fig. 2.

534 **Analysis of receptive fields**

535 To determine the orientations of the basis functions (BFs) we use MATLAB's implementation of the
536 *trust region reflective algorithm* for non-linear curve fitting (*Coleman and Li (1996)*) to fit them to
537 two-dimensional Gabor functions as defined by:

$$G(x, y, \theta, f, \psi, \sigma, \xi) = \exp\left(-\frac{x'^2 + \xi^2 y'^2}{2\sigma^2}\right) \cos(2\pi f x' + \psi), \quad (8)$$

538 with $x' = x \cos(\theta) + y \sin(\theta)$ and $y' = -x \sin(\theta) + y \cos(\theta)$.

539 Here, f denotes the frequency, ψ the phase offset, σ the standard deviation of the Gaussian
540 envelope, ξ the spatial aspect ratio and θ the orientation, where $\theta = 0$ deg corresponds to a vertically
541 oriented Gabor function. We initialize the parameters randomly 150 times and fit the function
542 either to the left or right BFs (or to both, see below). To evaluate the quality of the fits, we record
543 the difference between the actual BFs and the Gabor fit. More specifically, the *residual* is defined as
544 the sum of the squared differences in single pixel values between BFs and the fit. To compare the
545 fits across the different experimental conditions, we only took those fits where this residual was
546 less than or equal to 0.2. This accounts for more than 96% of all BFs in the healthy case.

547 Another interpretation for these fits is a stimulus that maximally activates the particular neuron.
548 To investigate the binocularity of such a cell, we compare their monocular response to the left and
549 right Gabor fit. The eye with the greater response is the dominant eye for this neuron. Similar as in
550 *Hubel and Wiesel (1962)* we show the best stimulus (here the Gabor fit) to the dominant eye and
551 the same stimulus to the non-dominant eye and record the responses L and R . We then compare
552 these by

$$b = \frac{R - L}{R + L} \quad (9)$$

553 to get a binocularity index between -1 (monocular left) and +1 (monocular right), where 0 means
554 perfectly binocular.

555 When fitting this function to binocular BFs, we assume that all parameters are equal for the left
556 and right monocular sub-region of the BFs except for the phase offset ψ , that can be different for
557 left and right eye. Following the assumption that the disparity tuning in a binocular cell is encoded
558 by means of this phase shift, we can calculate the preferred disparity d of a neuron by:

$$d = \frac{\psi_L - \psi_R}{2\pi\theta f}. \quad (10)$$

559 The maximally detectable disparity is given by the RF size, that is, the visual angle one binocular
560 patch covers. BFs with a disparity preference bigger than that are excluded from the analysis.

561 **Laplacian disparity distribution**

562 The probability density distribution of a Laplacian distributed random variable X is defined as

$$p(x) = ce^{-s|x-\mu|}, \quad A < x < B, \quad (11)$$

563 To simulate the disparity distribution we set μ to the angle that is desired to fixate an object at a
564 certain distance d_o ,

$$\mu = 2 \arctan\left(\frac{d_I}{2d_o}\right), \quad (12)$$

565 where $d_I = 56$ cm is the interpupillary distance. The data shown in Fig. 8B are generated from a
566 model with only the fine scale, for simplicity.

567 **Acknowledgments**

568 This work was supported by the German Federal Ministry of Education and Research under Grants
569 01GQ1414 and 01EW1603A (within the frame of ERA-NET NEURON), the European Union's Horizon
570 2020 Grant 713010, the Hong Kong Research Grants Council under Grant 16244416, and the Quandt
571 Foundation.

572 **References**

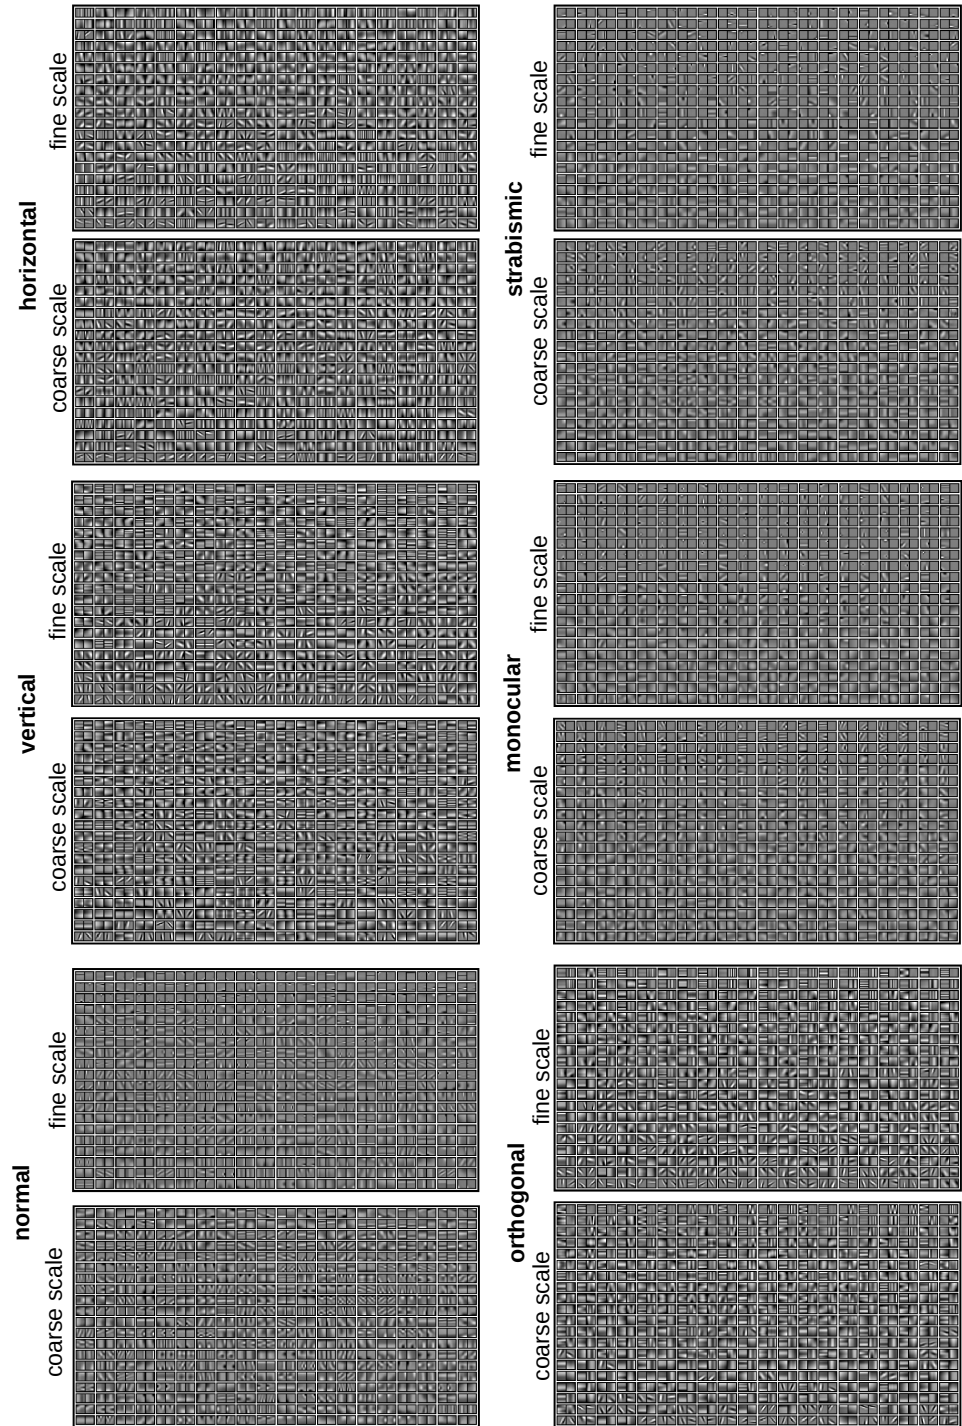
- 573 **Appelle S.** Perception and discrimination as a function of stimulus orientation: the " oblique effect" in man and
574 animals. *Psychological bulletin.* 1972; 78(4):266.
- 575 **Barlow HB.** Possible principles underlying the transformations of sensory messages. *Sensory Communication.*
576 1961; .
- 577 **Blake R, Wilson H.** Binocular vision. *Vision research.* 2011; 51(7):754-770.
- 578 **Chen Y, Qian N.** A coarse-to-fine disparity energy model with both phase-shift and position-shift receptive field
579 mechanisms. *Neural Computation.* 2004; 16(8):1545-1577.
- 580 **Coleman TF, Li Y.** An interior trust region approach for nonlinear minimization subject to bounds. *SIAM Journal*
581 *on optimization.* 1996; 6(2):418-445.
- 582 **Coppola DM, Purves HR, McCoy AN, Purves D.** The distribution of oriented contours in the real world. *Proceed-*
583 *ings of the National Academy of Sciences.* 1998; 95(7):4002-4006.
- 584 **Crowley JC, Katz LC.** Early development of ocular dominance columns. *Science.* 2000; 290(5495):1321-1324.
- 585 **Curcio CA, Sloan KR, Kalina RE, Hendrickson AE.** Human photoreceptor topography. *Journal of comparative*
586 *neurology.* 1990; 292(4):497-523.
- 587 **Dan Y, Atick JJ, Reid RC.** Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of
588 a computational theory. *Journal of Neuroscience.* 1996; 16(10):3351-3362.
- 589 **Daugman JG.** Uncertainty relation for resolution in space, spatial frequency, and orientation optimized by
590 two-dimensional visual cortical filters. *JOSA A.* 1985; 2(7):1160-1169.
- 591 **Daw NW.** Critical periods and amblyopia. *Archives of ophthalmology.* 1998; 116(4):502-505.

- 592 **De Valois RL**, Albrecht DG, Thorell LG. Spatial frequency selectivity of cells in macaque visual cortex. *Vision*
593 *research*. 1982; 22(5):545–559.
- 594 **De Valois RL**, Yund EW, Hepler N. The orientation and direction selectivity of cells in macaque visual cortex.
595 *Vision research*. 1982; 22(5):531–544.
- 596 **Ding J**, Levi DM. Recovery of stereopsis through perceptual learning in human adults with abnormal binocular
597 vision. *Proceedings of the National Academy of Sciences*. 2011; 108(37):E733–E741.
- 598 **Doi E**, Gauthier JL, Field GD, Shlens J, Sher A, Greschner M, Machado TA, Jepson LH, Mathieson K, Gunning
599 DE, et al. Efficient coding of spatial information in the primate retina. *Journal of Neuroscience*. 2012;
600 32(46):16256–16264.
- 601 **Eckmann S**, Klimmasch L, Shi BE, Triesch J. Active Efficient Coding Explains the Development of Binocular Vision
602 and its Failure in Amblyopia. *bioRxiv*. 2019; p. 571802.
- 603 **Fawcett SL**, Wang YZ, Birch EE. The critical period for susceptibility of human stereopsis. *Investigative*
604 *ophthalmology & visual science*. 2005; 46(2):521–525.
- 605 **Fleet DJ**, Wagner H, Heeger DJ. Neural encoding of binocular disparity: energy models, position shifts and phase
606 shifts. *Vision research*. 1996; 36(12):1839–1857.
- 607 **Fox R**, Aslin RN, Shea SL, Dumais ST. Stereopsis in human infants. *Science*. 1980; 207(4428):323–324.
- 608 **Franz A**, Triesch J. Emergence of disparity tuning during the development of vergence eye movements. In: 2007
609 *IEEE 6th International Conference on Development and Learning IEEE*; 2007. p. 31–36.
- 610 **Fregnac Y**, Trotter Y, Bienenstock E, Buisseret P, Gary-Bobo E, Imbert M. Effect of neonatal unilateral enucleation
611 on the development of orientation selectivity in the primary visual cortex of normally and dark-reared kittens.
612 *Experimental brain research*. 1981; 42(3-4):453–466.
- 613 **Furmanski CS**, Engel SA. An oblique effect in human primary visual cortex. *Nature neuroscience*. 2000; 3(6):535.
- 614 **Gopal SK**, Kelkar J, Kelkar A, Pandit A. Simplified updates on the pathophysiology and recent developments in
615 the treatment of amblyopia: A review. *Indian journal of ophthalmology*. 2019; 67(9):1392.
- 616 **Grondman I**, Busoniu L, Lopes GA, Babuska R. A survey of actor-critic reinforcement learning: Standard and
617 natural policy gradients. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and*
618 *Reviews)*. 2012; 42(6):1291–1307.
- 619 **Held R**, Birch E, Gwiazda J. Stereoacuity of human infants. *Proceedings of the National Academy of Sciences*.
620 1980; 77(9):5572–5574.
- 621 **Hirsch HV**, Spinelli D. Visual experience modifies distribution of horizontally and vertically oriented receptive
622 fields in cats. *Science*. 1970; 168(3933):869–871.
- 623 **Hubel DH**, Wiesel TN. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex.
624 *The Journal of physiology*. 1962; 160(1):106–154.
- 625 **Hubel DH**, Wiesel TN. Binocular interaction in striate cortex of kittens reared with artificial squint. *Journal of*
626 *neurophysiology*. 1965; 28(6):1041–1059.
- 627 **Hunt JJ**, Dayan P, Goodhill GJ. Sparse coding can predict primary visual cortex receptive field changes induced
628 by abnormal visual input. *PLoS computational biology*. 2013; 9(5):e1003005.
- 629 **Jordan MI**, Mitchell TM. Machine learning: Trends, perspectives, and prospects. *Science*. 2015; 349(6245):255–
630 260.
- 631 **Julesz B**. Foundations of cyclopean perception. . 1971; .
- 632 **Kandel ER**, Schwartz JH, Jessell TM, of Biochemistry D, Jessell MBT, Siegelbaum S, Hudspeth A. Principles of
633 neural science, vol. 4. McGraw-hill New York; 2000.
- 634 **Klimmasch L**, Lelais A, Lichtenstein A, Shi BE, Triesch J. Learning of Active Binocular Vision in a Biomechanical
635 Model of the Oculomotor System. *bioRxiv*. 2017; p. 160721.
- 636 **Kotsiantis SB**, Zaharakis I, Pintelas P. Supervised machine learning: A review of classification techniques.
637 *Emerging artificial intelligence applications in computer engineering*. 2007; 160:3–24.

- 638 **Le Vay S**, Wiesel TN, Hubel DH. The development of ocular dominance columns in normal and visually deprived
639 monkeys. *Journal of Comparative Neurology*. 1980; 191(1):1-51.
- 640 **Leventhal AG**, Hirsch HV. Cortical effect of early selective exposure to diagonal lines. *Science*. 1975;
641 190(4217):902-904.
- 642 **Li B**, Peterson MR, Freeman RD. Oblique effect: a neural basis in the visual cortex. *Journal of neurophysiology*.
643 2003; 90(1):204-217.
- 644 **Lonini L**, Forestier S, Teulière C, Zhao Y, Shi BE, Triesch J. Robust active binocular vision through intrinsically
645 motivated learning. *Frontiers in Neurorobotics*; 2013.
- 646 **Mallat SG**, Zhang Z. Matching pursuits with time-frequency dictionaries. *IEEE Transactions on Signal Processing*.
647 1993; 41(12):3397-3415.
- 648 **Okuda FC**, Apt L, Wanter BS. Evaluation of the TNO random-dot stereogram test. *American Orthoptic Journal*.
649 1977; 27(1):124-130.
- 650 **Olmos A**, Kingdom F, McGill calibrated colour image database; 2004.
- 651 **Olmos A**, Kingdom FA. A biologically inspired algorithm for the recovery of shading and reflectance images.
652 *Perception*. 2004; 33(12):1463-1473.
- 653 **Olshausen BA**, Field DJ. Sparse Coding with an Overcomplete Basis Set: A Strategy Employed by V1? *Vision*
654 *Research*. 1997; 37(23):3311-3325.
- 655 **Olshausen BA**, et al. Emergence of simple-cell receptive field properties by learning a sparse code for natural
656 images. *Nature*. 1996; 381(6583):607-609.
- 657 **Papageorgiou E**, Asproudis I, Maconachie G, Tsironi EE, Gottlob I. The treatment of amblyopia: current practice
658 and emerging trends. *Graefes Archive for Clinical and Experimental Ophthalmology*. 2019; p. 1-18.
- 659 **Priamikov A**, Fronius M, Shi B, Triesch J. OpenEyeSim: A biomechanical model for simulation of closed-loop
660 visual perception. *Journal of Vision*. 2016; 16(15):25-25.
- 661 **Priamikov A**, Triesch J. Openeyesim-a platform for biomechanical modeling of oculomotor control. In: *IEEE*
662 *International Conference on Development and Learning and on Epigenetic Robotics*; 2014. p. 394-395.
- 663 **Ringach DL**, Hawken MJ, Shapley R. Dynamics of orientation tuning in macaque primary visual cortex. *Nature*.
664 1997; 387(6630):281.
- 665 **Rothkopf CA**, Weisswange TH, Triesch J. Learning independent causes in natural images explains the space-
666 variant oblique effect. In: *2009 IEEE 8th International Conference on Development and Learning IEEE*; 2009. p.
667 1-6.
- 668 **Ruttum MS**. Visual screening with random dot stereograms. In: *Seminars in Ophthalmology*, vol. 3 Taylor &
669 Francis; 1988. p. 175-180.
- 670 **Shlaer R**. Shift in binocular disparity causes compensatory change in the cortical structure of kittens. *Science*.
671 1971; 173(3997):638-641.
- 672 **Simoncelli EP**. Vision and the statistics of the visual environment. *Current opinion in neurobiology*. 2003;
673 13(2):144-149.
- 674 **Smith EC**, Lewicki MS. Efficient auditory coding. *Nature*. 2006; 439(7079):978-982.
- 675 **Sprague WW**, Cooper EA, Tošić I, Banks MS. Stereopsis is adaptive for the natural environment. *Science*
676 *advances*. 2015; 1(4):e1400254.
- 677 **Stevenson SB**, Cormack LK, Schor CM, Tyler CW. Disparity tuning in mechanisms of human stereopsis. *Vision*
678 *research*. 1992; 32(9):1685-1694.
- 679 **Stidwill D**, Fletcher R. Normal binocular vision: Theory, investigation and practical aspects. John Wiley & Sons;
680 2017.
- 681 **Stryker MP**, Sherk H, Leventhal AG, Hirsch HV. Physiological consequences for the cat's visual cortex of
682 effectively restricting early visual experience with oriented contours. *Journal of Neurophysiology*. 1978;
683 41(4):896-909.

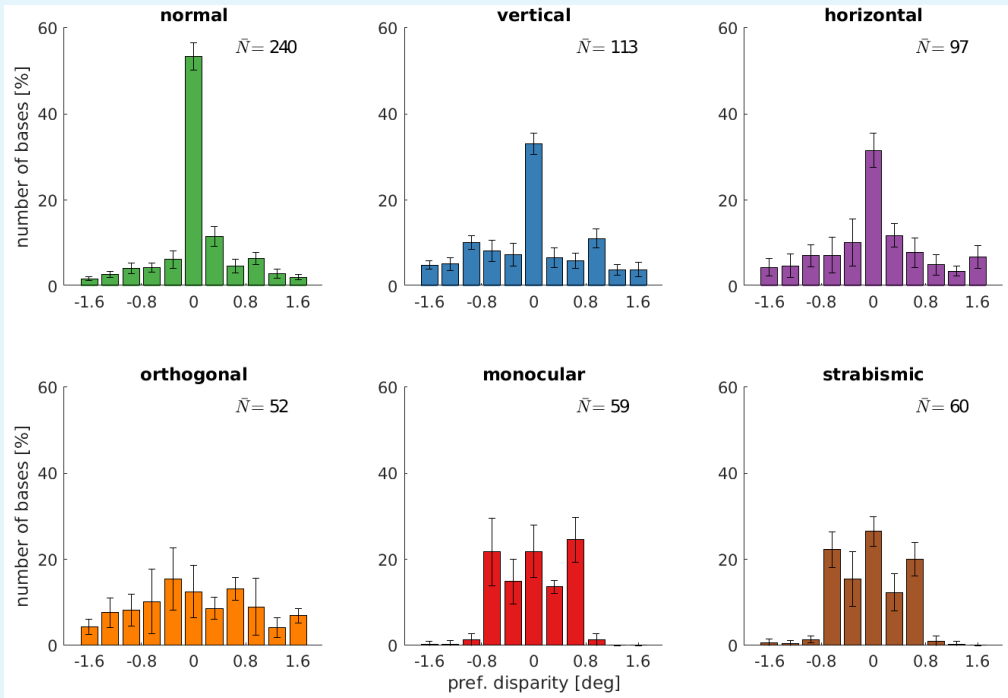
- 684 **Sun P**, Gardner JL, Costagli M, Ueno K, Waggoner RA, Tanaka K, Cheng K. Demonstration of tuning to stimulus
685 orientation in the human visual cortex: a high-resolution fMRI study with a novel continuous and periodic
686 stimulation paradigm. *Cerebral Cortex*. 2012; 23(7):1618–1629.
- 687 **Sutton RS**, Barto AG. Reinforcement Learning: An Introduction, vol. 1. MIT press Cambridge; 1998.
- 688 **Tanaka S**, Ribot J, Imamura K, Tani T. Orientation-restricted continuous visual exposure induces marked
689 reorganization of orientation maps in early life. *Neuroimage*. 2006; 30(2):462–477.
- 690 **Tesauro G**. Temporal difference learning and TD-Gammon. *Communications of the ACM*. 1995; 38(3):58–68.
- 691 **Umberger BR**, Gerritsen KG, Martin PE. A model of human muscle energy expenditure. *Computer Methods in*
692 *Biomechanics and Biomedical Engineering*. 2003; 6(2):99–111.
- 693 **Van Hasselt H**, Wiering MA. Reinforcement learning in continuous action spaces. In: *IEEE International*
694 *Symposium on Approximate Dynamic Programming and Reinforcement Learning*; 2007. p. 272–279.
- 695 **Vinje WE**, Gallant JL. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science*.
696 2000; 287(5456):1273–1276.
- 697 **Walraven J**. Amblyopia screening with random-dot stereograms. *American journal of ophthalmology*. 1975;
698 80(5):893–900.
- 699 **Welford B**. Note on a method for calculating corrected sums of squares and products. *Technometrics*. 1962;
700 4(3):419–420.
- 701 **Wiesel TN**, Hubel DH. Single-cell responses in striate cortex of kittens deprived of vision in one eye. *Journal of*
702 *neurophysiology*. 1963; 26(6):1003–1017.
- 703 **Yacoub E**, Harel N, Uğurbil K. High-field fMRI unveils orientation columns in humans. *Proceedings of the*
704 *National Academy of Sciences*. 2008; 105(30):10607–10612.
- 705 **Young IT**. Proof without prejudice: use of the Kolmogorov-Smirnov test for the analysis of histograms from flow
706 systems and other sources. *Journal of Histochemistry & Cytochemistry*. 1977; 25(7):935–941.
- 707 **Zhao Y**, Rothkopf CA, Triesch J, Shi BE. A unified model of the joint development of disparity selectivity and
708 vergence control. In: *IEEE International Conference on Development and Learning and Epigenetic Robotics (ICDL)*;
709 2012. p. 1–6.

710 **Appendix 1**



Appendix 1 Figure 1. Complete set of all BFs that are learned during training for all different rearing conditions.

715 **Appendix 2**



716

717

718

Appendix 2 Figure 1. Disparity tuning of the fine scale of models that were trained under different rearing conditions.

720 **Appendix 3**

