

Learning Divisive Normalization in Primary Visual Cortex

Max F. Günthner,^{1,2,*} Santiago A. Cadena,¹⁻³ George H. Denfield^{3,4}
Edgar Y. Walker,^{3,4} Andreas S. Tolias,^{2-5,§} Matthias Bethge^{1-3,§}
Alexander S. Ecker^{1-3,§,†}

¹ Inst. for Theoretical Physics and Centre for Integrative Neuroscience, University of Tübingen, Germany

² Bernstein Center for Computational Neuroscience, Tübingen, Germany

³ Center for Neuroscience and Artificial Intelligence, Baylor College of Medicine, Houston, TX, USA

⁴ Department of Neuroscience, Baylor College of Medicine, Houston, TX, USA

⁵ Department of Electrical and Computer Engineering, Rice University, Houston, TX, USA

* max.guenthner@bethgelab.org

§ These authors contributed equally

† Present address: Department of Computer Science, University of Göttingen, Germany

November 8, 2019

1 .
2 Deep convolutional neural networks (CNNs) have emerged as the state of the art
3 for predicting neural activity in visual cortex. While such models outperform classical
4 linear-nonlinear and wavelet-based representations, we currently do not know what
5 computations they approximate. Here, we tested divisive normalization (DN) for
6 its ability to predict spiking responses to natural images. We developed a model
7 that learns the pool of normalizing neurons and the magnitude of their contribution
8 end-to-end from data. In macaque primary visual cortex (V1), we found that
9 our interpretable model outperformed linear-nonlinear and wavelet-based feature
10 representations and almost closed the gap to high-performing black-box models.
11 Surprisingly, within the classical receptive field, oriented features were normalized
12 preferentially by features with similar orientations rather than non-specifically as
13 currently assumed. Our work provides a new, quantitatively interpretable and
14 high-performing model of V1 applicable to arbitrary images, refining our view on
15 gain control within the classical receptive field.

1 Introduction

16 1 Introduction

17 A crucial step towards understanding the visual system is to build models that predict neural
18 responses to arbitrary stimuli with high accuracy (Carandini et al., 2005). The classical standard
19 models of the primary visual cortex (V1) are based on linear-nonlinear models (Simoncelli
20 et al., 2004), energy models (Adelson and Bergen, 1985) and subunit (LN-LN) models (Rust
21 et al., 2005; Touryan et al., 2005; Willmore et al., 2008; Butts et al., 2011; McFarland et al.,
22 2013; Vintch et al., 2015). Fueled by advances in machine learning technology, recent studies
23 have shown that multi-layer convolutional neural networks (CNNs) can significantly improve
24 prediction of neural responses to complex images at several stages of the visual pathway
25 (Yamins et al., 2014; Khaligh-Razavi and Kriegeskorte, 2014; McIntosh et al., 2016; Zhang
26 et al., 2018; Cadena et al., 2019; Kindel et al., 2019), outperforming classical models. The
27 current state-of-the-art data-driven model of single-unit activity in monkey V1 is a three-layer
28 black-box CNN (Cadena et al., 2019). However, such models are difficult to interpret, limiting
29 our understanding of V1 function. In particular, we do not have first principles explaining what
30 kind of nonlinear mapping the black-box CNNs approximate.

31 A promising candidate to facilitate a more principled description of V1 neurons is to replace
32 the black-box computations by divisive normalization (Heeger, 1992), which has been proposed
33 to be a canonical neural computation throughout the visual pathway (Carandini and Heeger,
34 2012) because it explains a wide variety of neurophysiological phenomena (Carandini and
35 Heeger, 2012; Sawada and Petrov, 2017) and can be derived from first principles of redundancy
36 reduction (Schwartz and Simoncelli, 2001; Sinz and Bethge, 2008). A prominent example for
37 such normalization phenomena in V1 is cross-orientation inhibition. Here, the response of a
38 neuron to a driving grating stimulus in the receptive field (RF) is suppressed by superimposing
39 a second grating that would not elicit a response when presented alone: for instance, a grating
40 with orientation orthogonal to the neuron's preferred orientation (Bonds, 1989; Morrone et al.,
41 1982; DeAngelis et al., 1992; Heeger, 1992; Carandini et al., 1997; Busse et al., 2009).

42 The basic idea of divisive normalization (Fig. 1A) is that a neuron's driving input is normalized
43 divisively by a weighted sum over nearby neurons' responses (Heeger, 1992; Carandini and
44 Heeger, 2012). While the general idea is simple, elegant and powerful, our current knowledge of
45 DN is limited in two important ways: (1) DN has been studied mostly using simple stimuli
46 and we do not know whether incorporating DN into predictive models of neural responses
47 improves these models' performance on natural images, and (2) we currently do not know how
48 receptive field location and response properties determine whether a neuron contributes to the
49 normalization pool and, if so, with what normalization weight.

50 To explain normalization phenomena within the classical receptive field like cross-orientation
51 inhibition, current models of divisive normalization assume that all nearby neurons with diverse
52 orientation tuning preferences and with similar receptive field locations contribute equally to
53 the normalization pool (Heeger, 1992; Carandini et al., 1997; Busse et al., 2009). However,
54 some original experimental studies suggest that this assumption may not be correct for some
55 neurons (Bonds, 1989; DeAngelis et al., 1992), and normative models of normalization predict
56 that the magnitude with which a given neuron contributes to another neuron's normalization
57 depends on the relationship of their response properties (Schwartz and Simoncelli, 2001).

1 Introduction

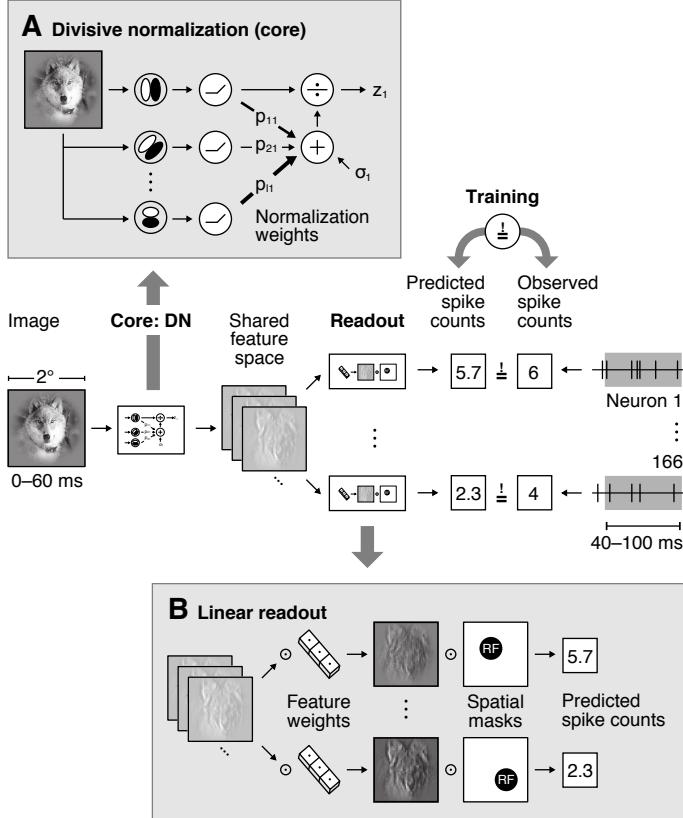


Figure 1: Overview of our divisive normalization (DN) model. The model takes as input an image and predicts neurons' spike counts in response to this image (details in Fig. 2). The model is split into two parts: a *core* that computes a shared nonlinear feature space and a *readout* that maps the shared feature space individually to each neuron's spike count. **A.** Divisive normalization mechanism (simplified). The visual input is convolved with 32 filters and then rectified to produce an excitatory output. The output of each filter is then divided by a weighted sum of the excitatory outputs of all filters with normalization weights p_{kl} and a semi-saturation constant σ_l . In our general formulation, all weights and constants are learned from the data. **B.** Linear readout that maps the shared feature space to each neuron's spike count through an individual weighted sum over the entire shared feature space. The readout weights are factorized into a feature vector – capturing the nonlinear feature(s) that a neuron computes – and a spatial mask – localizing each neuron's receptive field (RF).

58 In this paper, we address two main questions raised above: (1) can an interpretable model based
 59 on divisive normalization match the superior performance of black-box CNNs over simpler,
 60 interpretable subunit or energy models when predicting spiking responses to natural images
 61 and (2) how are V1 neurons normalized? We focus on responses to stimuli mostly restricted
 62 to the classical receptive field and on models that account only for normalization by neurons
 63 with overlapping receptive field locations. We developed an end-to-end trainable divisive
 64 normalization model to predict V1 spike counts from natural stimuli. Our model learns the
 65 filter coefficients of all neurons as well as their normalization weights directly from the data.
 66 We applied our model to natural image responses in monkey V1 and found that it outperforms
 67 linear-nonlinear and subunit models, and is competitive with that of state-of-the-art CNNs
 68 while requiring much fewer parameters and being directly interpretable. This result implies
 69 that divisive normalization is an important computation under stimulation with natural images.
 70 Importantly, we found that oriented features were normalized preferentially by features with
 71 similar orientation, in contrast to the current standard model of nonspecific normalization
 72 (Heeger, 1992; Busse et al., 2009). Our work thus advances our understanding of V1 function
 73 by establishing a new state-of-the-art interpretable model and predicting an orientation-specific
 74 divisive normalization mechanism under stimulation with natural stimuli.

2 Results

75 2 Results

76 2.1 Learning divisive normalization

The basic idea of divisive normalization (Fig. 1A) is that the response of neuron l

$$z_l(x) = \frac{y_l(x)}{\sigma_l + \sum_{k \in \mathcal{K}} p_{kl} \cdot y_k(x)} \quad (1)$$

77 is given by its driving input activity $y_l(x)$ divisively normalized by a weighted sum over nearby
78 neurons' responses $y_k(x)$ (Heeger, 1992; Carandini and Heeger, 2012), where x represents the
79 stimulus and σ_l is a semi-saturation constant. Here, the set of normalizing neurons \mathcal{K} and the
80 normalization weights p_{kl} define which neurons contribute to the normalization pool of neuron
81 l and with what strength, respectively.

82 While this formulation is straightforward to write down, it is challenging to build quantitative
83 models based on it that are applicable to arbitrary inputs. The denominator depends on a
84 potentially large population of neurons – which is unknown in general – and the structure of
85 the normalization weights has been studied only using very restricted sets of simple stimuli such
86 as oriented gratings and bars. Previous modeling work on divisive normalization has therefore
87 made specific assumptions about the filter properties of the underlying neuronal population
88 and either modeled only a closed set of stimuli such as gratings of different orientation (Heeger,
89 1992; Carandini et al., 1997; Freeman et al., 2002; Heuer and Britten, 2002; Busse et al., 2009)
90 or evaluated models only qualitatively (Schwartz and Simoncelli, 2001; Wainwright et al., 2002;
91 Froudarakis et al., 2014).

92 We developed a general, image-computable predictive model of divisive normalization following
93 Eq. (1), which is applicable to arbitrary images and whose parameters are learned by optimizing
94 the accuracy of the model in predicting the spiking activity of a large number of neurons in
95 response to natural images (see Fig. 1). Our model builds on a recent innovation in predictive
96 modeling (Antolik et al., 2016; Klindt et al., 2017; Batty et al., 2016; McIntosh et al., 2016;
97 Cadena et al., 2019), jointly modeling all recorded neurons instead of learning a predictive
98 model for each neuron individually. Because many neurons perform similar computations – up
99 to shifts in receptive field location – jointly modeling them makes more efficient use of the data
100 and we can learn more complex models. The basic idea is to split the model into two parts
101 (Fig. 1): (1) a *core* that transforms the input image into nonlinear features shared among all
102 neurons, and (2) a *readout* that linearly combines the features to produce a prediction of each
103 neuron's response.

104 We use a convolutional network for the core, whose architecture lends itself very well to model
105 divisive normalization. By construction, we have a model that contains all filters necessary
106 to account for the recorded neurons' responses. All of these filter responses are automatically
107 extracted at each location, providing a good approximation of the underlying population
108 of neurons in the brain although it is only sparsely sampled during the experiment. As a
109 consequence, we can optimize the pool of neurons providing normalizing inputs and their
110 corresponding weights p_{kl} (Eq. 1) to account for the neural responses.

2 Results

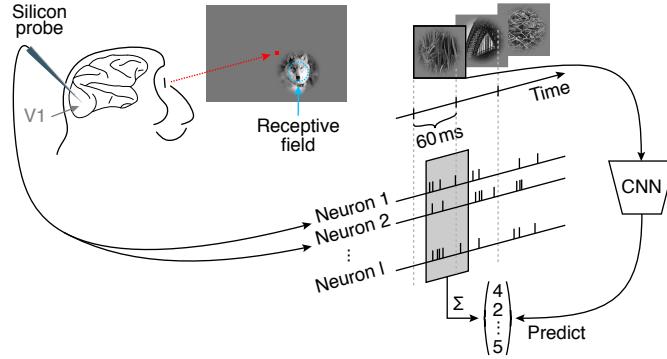


Figure 2: Experimental paradigm from Cadena et al. (2019). Natural images were flashed to a monkey covering 2° of their visual angle, and located at the center of the multi-unit receptive field. Multiple neurons were isolated from recordings with silicon probes inserted into V1 (Denfield et al., 2018). Natural images were shown in a fast sequence without blanks, each presented for 60 ms. Spike counts from all isolated neurons corresponding to each image were extracted from a window 40 ms after the image onset lasting 60 ms.

111 In summary, our model's core (Fig. 1A) consists of a set of convolutional filters (we use 32)
112 that provide the driving inputs, followed by a DN stage (Eq. 1). This core is shared among all
113 neurons and converts the image into a set of feature maps. These feature maps are converted
114 into response predictions by a linear readout step (Fig. 1B) that picks the relevant features
115 and spatial locations for each neuron. To ensure that the readout does not model any complex
116 computation, we constrain its weights to be non-negative. The non-negativity ensures that
117 activations can only add, preventing the readout stage from accounting for any suppressive
118 effects. The readout can, however, account for response invariances such as phase invariance of
119 complex cells; see Methods for an in-depth explanation. While our model reflects the general
120 formulation of divisive normalization, in this paper we mostly focus on normalization from the
121 vicinity of the receptive field.

122 2.2 DN model achieves competitive accuracy with fewer parameters

123 We fit the model described above to a dataset of 166 neurons recorded in V1 of two awake,
124 fixating monkeys (data from Cadena et al. 2019), who viewed a fast sequence of localized natural
125 images and textures (Fig. 2). The stimuli were centered on the neurons' receptive fields and
126 covered about twice the area of the classical receptive fields, mostly stimulating the near vicinity
127 of the RFs' center. Images were shown for 60 ms each, without blanks in between. Single unit
128 activity was recorded with laminar silicon probes sampling from all cortical layers. We fit the
129 model jointly to the responses of all neurons. As neurons were recorded in 17 recording sessions,
130 the dataset sampled a diverse range of preferred orientations. The objective function during
131 training was to minimize the difference between the model's prediction and the observed spike
132 counts of the neurons in a time window 40–100 ms after image onset (to account for response
133 latency).

2 Results

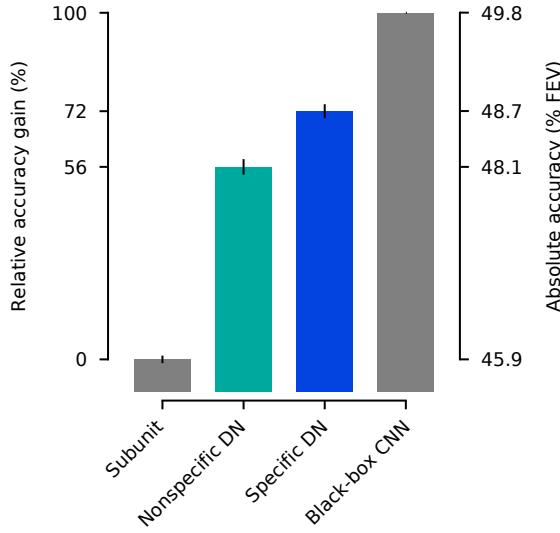


Figure 3: Performance comparison of our models fitted to the data from Cadena et al. (2019) relative to the gap between the best interpretable model – a subunit one layer convolutional neural network (CNN) – and the data-driven state-of-the-art three-layer CNN (Cadena et al., 2019) that offers little interpretability (black-box). Non-specific divisive normalization (DN) accounts for 56% of this gap, while specific DN improves it up to 72%. Absolute values in terms of percentage of explainable variance explained (FEV) on the right (mean over the ten best models selected in terms of validation set accuracy). Error bars show the corresponding standard error of the mean.

134 To evaluate model performance, we estimated the fraction of explainable variance explained
135 (FEV), which quantifies the fraction of the stimulus-driven response variance that is accounted
136 for by the model, and ignores unexplainable trial-to-trial variability in the response of the
137 neurons (see Methods). A perfect model would reach a FEV of 100%.

138 Subunit models are an established approach to model primary visual cortex responses (Rust
139 et al., 2005; Touryan et al., 2005; Willmore et al., 2008; Butts et al., 2011; McFarland et al., 2013;
140 Vintch et al., 2015). In addition to capturing a fair portion of the explainable variance, they
141 provide interpretability in the form of linear projections applied to the input images. Therefore,
142 we considered a convolutional subunit model – currently the best-performing interpretable
143 model of V1 (Cadena et al., 2019) – as a strong baseline for our model. It consists of a first
144 stage of rectified linear filtering followed by a static nonlinearity, and then a linear pooling
145 stage. Structurally, it is the same as our DN model, but without the normalization stage. This
146 subunit model accounted for 45.9% FEV. In comparison, a regularized linear nonlinear Poisson
147 model (LNP) only accounted for 16.3% FEV on the same dataset (Cadena et al., 2019) due to
148 its inability to model complex cells.

149 As recent developments in machine learning technology have allowed us to improve predictive
150 performance, we used the current best data-driven model as a gold standard. This model is a
151 black-box convolutional neural network with three convolutional layers and a linear-nonlinear
152 readout, reaching a performance of 49.8% FEV (Cadena et al., 2019). However, although this
153 model outperforms the simpler subunit model, we currently do not understand how it does

2 Results

Model	Number of parameters	
	Core	Readout per neuron
Subunit model	5 440	816
Nonspecific divisive normalization	5 536	816
Divisive normalization	6 528	816
Black-box CNN (Cadena et al., 2019)	23 936	867

Table 1: Number of parameters for different models.

154 SO.

155 To evaluate how well our DN model accounts for the data, we placed its performance on a scale
156 between 0% (baseline: subunit model) and 100% (gold standard: black-box CNN). On that scale,
157 our DN model achieved a score of 72% between the baseline and gold standard (48.7% FEV on
158 test set, mean over the ten best models selected in terms of validation set accuracy, Fig. 3), being
159 the new state-of-the-art interpretable model of primary visual cortex. Notably, we achieved
160 this performance gain by simply adding the trainable DN stage to the convolutional subunit
161 model, which shows that divisive normalization is an important computational mechanism in
162 V1 under stimulation with natural images.

163 While our DN model’s accuracy comes close to that of the state-of-the-art black-box CNN, it
164 requires substantially fewer parameters to achieve this performance (Table 1): The DN model’s
165 core – i. e. the shared computational backend before the linear readout – uses only 27.3% of
166 the parameters of the black-box CNN model’s core. This saving in parameters suggests that
167 the DN model captures important structure in the data, which we elaborate in the next section.
168 Compared to the number of parameters required by the subunit model’s core, the DN model
169 requires 20.0% more parameters (allocated to the divisive normalization module). The number
170 of readout parameters – i. e. the part that turns the shared nonlinear feature representation
171 into individual neurons’ responses – is very similar for all models.

172 2.3 Normalization is feature-specific

173 Having established that the DN model outperforms the current best interpretable model
174 and performs close to the black-box gold standard, we next investigated the structure of the
175 normalizing input, i. e. the sum in the denominator of Eq. (1) and how strongly different
176 features contribute to it. For this analysis, we focus on orientation-selective features. Visually
177 inspecting the strength of the normalizing inputs suggests that oriented features are normalized
178 preferentially by features with similar orientation preference (Fig. 4). In contrast, orthogonal
179 features seem to contribute less.

180 To quantify the difference with which the two groups contribute to normalization, we split
181 the sum in Eq. (1) into two parts and collect the contribution of normalizing features with
182 similar ($< 45^\circ$) orientation as the driving feature and that of features with dissimilar ($\geq 45^\circ$)
183 orientations. Analyzing the normalization of each oriented feature individually, we found that

2 Results

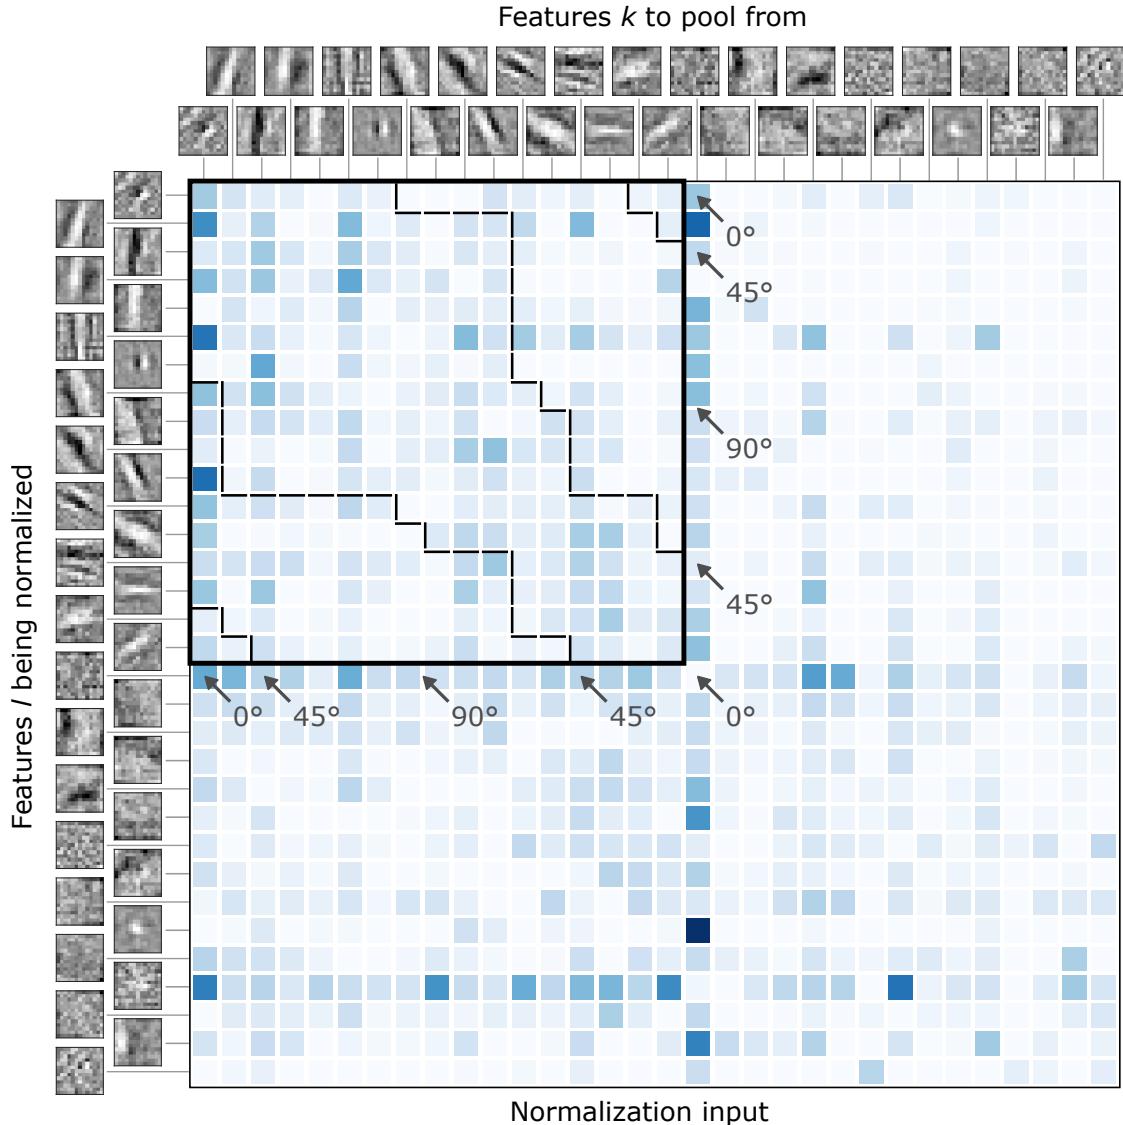


Figure 4: Structure of divisive normalization. The matrix shows the average strength (over images) of the normalizing inputs for each combination of filter response being normalized (rows) and filter response providing normalizing input (columns). Darker shades of blue indicate stronger normalization. Orientation-selective filters are grouped at the top, ordered by preferred orientation and marked by the black square. The dashed black lines within the square separate pairs of filters with similar ($< 45^\circ$) and dissimilar ($> 45^\circ$) orientations. Normalizing inputs are stronger for similarly tuned filters (see Fig. 6 for a quantification). Data of the model with highest accuracy on the validation set is shown.

2 Results

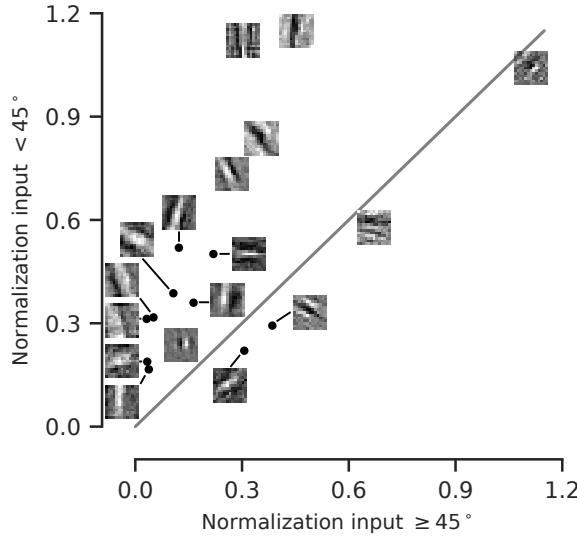


Figure 5: Normalization input from similar orientations ($< 45^\circ$) compared to the normalization input from dissimilar orientations ($\geq 45^\circ$) for each feature. Grey line: identity. Most features are normalized preferentially by the responses of filters with similar preferred orientations. Data of the model with highest accuracy on the validation set is shown.

most oriented features are more strongly normalized by features with similar orientations (Fig. 5). To assess whether our qualitative observation above is a general property of the data or a spurious characteristic of that one particular model we selected, we repeated this analysis for the top-10 models (assessed in terms of performance on the validation set) and observed similar behavior. Averaging over the features, we found that, for all of these models, similar orientations contributed more strongly than dissimilar orientations. Taking the data of all top-10 models into account, we found that, on average, similarly oriented features contributed 75% more normalizing input than dissimilar features (Wilcoxon signed rank test; $p < 0.006$, $N = 10$ models; Cohen's $d = 1.9$).

Having established that normalizing inputs are orientation-specific, we analyzed this specificity in more detail. Instead of using just two groups as before, we split up the normalizing inputs into nine bins of 10° width each and averaged those bins across the top-10 models. This analysis revealed that the strength of the normalizing inputs decreased as the difference in orientation increased (Fig. 6). Hence, the more similar a normalizing feature's orientation was to the feature to be normalized, the stronger was its contribution to normalization. In fact, features in the group most similar to the driving input contributed 133% more than those in the orthogonal group (Cohen's $d = 2.1$).

2.3.1 Control: Nonspecific divisive normalization reduces accuracy

To determine how important orientation-specific normalization is, we performed a control experiment: For each feature l being normalized, we constrained all of its incoming normalization

2 Results

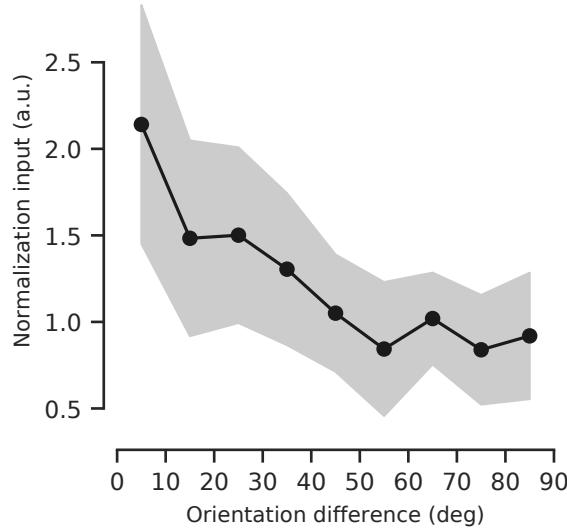


Figure 6: Normalization input, binned into orientation difference of 10° . Each bin was averaged over the top-10 models (assessed on the validation set). The shaded area depicts the standard deviation per bin.

weights p_{kl} to be identical. This constraint resembles non-specific normalization from all features, as assumed in previous models (Heeger, 1992; Carandini et al., 1997; Busse et al., 2009). This model achieved a performance of 56% between the baseline and gold standard (48.1% FEV). While it does not match the performance of our more general DN model, it does outperform the subunit baseline. Thus, orientation-specific normalization is necessary to achieve full performance.

2.3.2 Control: All channels contribute to our model's prediction

One potential caveat of our analyses so far is that we analyzed the orientation specificity of DN in terms of the convolutional feature maps in our model's core rather than the actual neurons we recorded. These features provide a much more compact view of the population of neurons, because they are invariant to the receptive field locations and the neural responses are simple linear combinations of those features. However, it is not clear a-priori whether all features are equally important for predicting the activity of the neurons in our population. Thus, considering convolutional features instead of actual neurons may lead to a skewed view of the population. To verify that this is not an issue, we quantified how much each feature contributed to the overall activity of all neurons by normalizing the feature readout weights across channels and averaging across neurons. The resulting distribution (Fig. 7) containing these averaged feature readout weights for the best ten models had a coefficient of variation of 0.2. We therefore concluded that all features were read out by roughly the same number of neurons and hence were similarly important to predict neural activity. Thus, our interpretation

2 Results

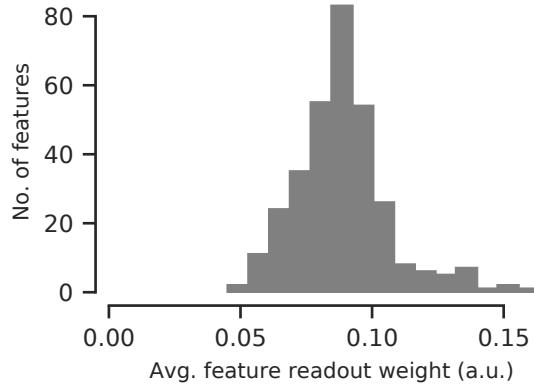


Figure 7: Histogram of feature readout weights of the ten best performing models in terms of validation set accuracy. For each model, feature weights are normalized across channels and averaged across individual neurons. All model's channels are used to predict neural activity.

224 of orientation-specific normalization is unlikely to be an artifact of analyzing the convolutional
225 features rather than the actual neurons.

226 2.3.3 Control: No surround influence in our results and dataset

227 We have observed orientation-specific divisive normalization in the classical receptive field.
228 Surround suppression is known to be orientation-specific (Blakemore and Tobin, 1972; DeAngelis
229 et al., 1994; Cavanaugh et al., 2002; Coen-Cagli et al., 2015), so a potential concern would be
230 that some of the extra-classical surround of a unit's RF contributed to the results presented
231 above. To rule out this possibility, we fit a more general DN model, where we additionally
232 learn the spatial structure of the normalization pool instead of just limiting it to neurons with
233 overlapping receptive fields (see Methods). This extended DN model included two normalization
234 pools that could have different patterns of weights along the feature dimension. It is therefore
235 general enough to account for the standard model of DN with a nonspecific center normalization
236 pool and orientation-specific surround suppression.

237 In contrast to what one may expect, spatially expanding the normalization pool to cover
238 larger surround areas did not increase our model's accuracy; in fact, for larger surrounds the
239 performance even decreased (Fig. 8). The best performance was achieved for models with a
240 normalization pool of approximately the size of the units' RF (approximately 0.5° diameter).
241 Since performance for larger normalization pools decreased, we used the model with the smallest
242 pool. The normalization weights of the extended spatial normalization pool showed no visible
243 separation into center and surround and exhibited no or only weak contributions from the
244 classical RF's surround (Fig. 9). From both the decrease in performance for larger models
245 and the spatial shape of the normalization pool, we concluded that our model does not learn
246 influence from the RF surround. The reason for this limitation is very likely that the surrounding
247 regions in our stimuli were masked out, so there is no surround information available to be

3 Discussion

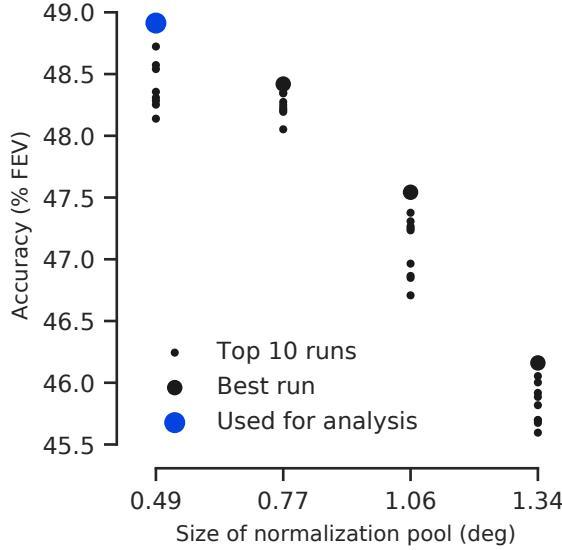


Figure 8: Validation set performance of our DN model for different normalization pool sizes (in space). The normalization pool has a square shape; x-axis denotes the edge length of the covered space that can contribute to normalization in units of visual angle in degrees.

248 learned. Consequently, our interpretation of orientation-specific normalization from nearby
249 units has no dependency on surrounding regions either.

250 3 Discussion

251 To improve our understanding of primary visual cortex, we asked what function state-of-the-art
252 black-box CNNs might implement for predicting V1 responses to localized natural stimuli.
253 To answer this question, we developed an end-to-end learnable divisive normalization model
254 and fit it to neural responses. Both the unspecific control model and the full model that
255 learned the normalization pool outperformed the current best-performing interpretable model
256 of V1, setting the new state-of-the-art. The full DN model improved performance even further,
257 reaching an accuracy competitive with the black-box CNN gold standard while having fewer
258 parameters. This result predicts that DN is a relevant mechanism to predict V1 responses to
259 natural images.

260 One may ask whether the difference between the non-specific DN model and the full model
261 learning orientation specific normalization weights is relevant, because the full model may simply
262 be able to better account for some insignificant biological heterogeneity due to its additional
263 parameters. Although it is possible, we believe that this explanation is unlikely, because oriented
264 features are preferentially normalized by channels with similar orientation. If the model was
265 simply picking up some biological imperfection, we would expect the normalization weights not
266 to depend systematically on preferred orientation.

3 Discussion

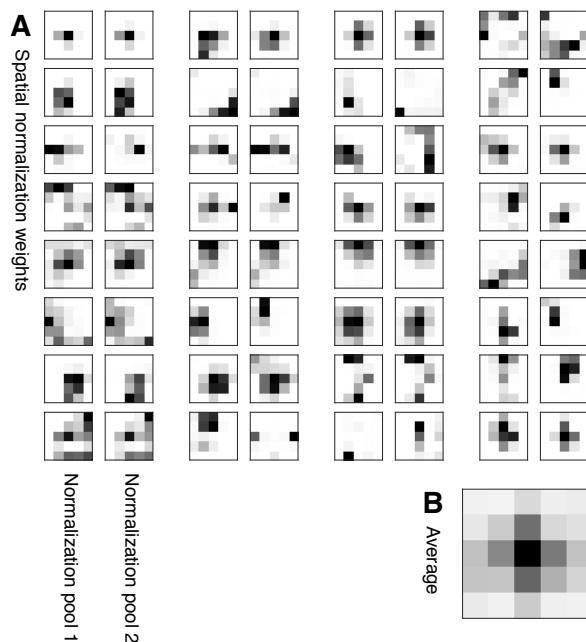


Figure 9: Weights of the spatial normalization pool for the best performing model (in terms of validation set accuracy) with an $5 \text{ px} \times 5 \text{ px}$ normalization pool (corresponding edge length in angle of the visual field: $1.06^\circ \times 1.06^\circ$). **A.** For each feature (rows), the two components of the in total 32 spatial normalization pools are shown. Darker color corresponds to higher weights. Both components are similar. **B.** Average across features and normalization pool components. The model learned normalization from the receptive field center (on average).

3 Discussion

267 Previous experimental work investigated suppressive phenomena within the receptive field only
268 with simple stimuli, mainly consisting of a combination of driving and mask gratings. Morrone
269 et al. (1982) find suppression at all orientations, but do not investigate orientations similar to
270 preferred orientation. Bonds (1989) report predominantly orientation-nonspecific suppression,
271 although three of fourteen cells exhibit stronger suppression with masks oriented similarly
272 to the neurons' preferred orientations, and a few other cells are suppressed most strongly by
273 mask orientations orthogonal to the preferred orientation. Similarly, DeAngelis et al. (1992)
274 find suppression to be predominantly independent of orientation, although for some cells an
275 increased suppression for a range of orientations near the optimal excitatory orientation is
276 apparent. Heeger (1992) explains those results by proposing an orientation-nonspecific divisive
277 normalization model. Carandini et al. (1997) consider the possibility of orientation-specific
278 normalization which provides a marginal improvement in the quality of their model fits to the
279 data. However, they conclude that their dataset was not specifically designed to provide a
280 strong test of this question and their results are inconclusive in this respect. Busse et al. (2009)
281 develop a quantitative model for the response of a population of neurons to a combination
282 of gratings. Assuming nonspecific normalization by overall contrast, their model predicts
283 the collective action of the whole neuron population better than linear and winner-take-all
284 baselines, but they do not test against an orientation-specific alternative model. To summarize,
285 these studies find phenomena that are predominantly explained by nonspecific normalization
286 (Heeger, 1992), some of them encountering only weak orientation-specific phenomena and only
287 in relatively few cells.

288 Thus, our findings are largely consistent with previous experimental results and quantitatively
289 refine them using a larger dataset, place them in the context of other models of V1 and show
290 that the same mechanisms observed with simple stimuli also apply under more natural stimulus
291 conditions. Interestingly, and somewhat unexpectedly based on earlier work, channels with
292 preferred orientations within 10° of the driving feature provided 133% stronger normalizing
293 input than those with orthogonal preferred orientations. The reason for this difference between
294 our findings and previous studies could be that we used natural stimuli, which have different
295 image statistics compared to simple stimuli used in earlier studies. Furthermore, most previous
296 studies of divisive normalization were performed in cats (Morrone et al., 1982; Bonds, 1989;
297 DeAngelis et al., 1992; Busse et al., 2009) and the results therein may not generalize to monkeys,
298 for which preceding studies are inconclusive regarding orientation specificity (Carandini et al.,
299 1997).

300 Recent work modeling a large set of classical psychophysical data also suggests an orientation-
301 specific divisive normalization: Schiütt and Wichmann (2017) developed an image-computable
302 model of early vision very similar in structure to ours, and found that in order to explain
303 classical data on contrast detection, contrast discrimination and oblique masking, their model
304 required divisive normalization to be orientation-specific. Similar results had been reported in
305 an earlier study (Itti et al., 2000).

306 Following a normative approach, Schwartz and Simoncelli (2001) derive an ecologically justified
307 divisive normalization model from the efficient coding hypothesis (Barlow, 1961) that is able to
308 describe the orientation masking data of Bonds (1989). Reducing the statistical redundancy of
309 responses to natural stimuli predicts that normalization should be stronger for neurons that

4 Methods

310 exhibit a higher dependency in their unnormalized responses. This theoretical result implies
311 that normalization weights should not be uniform, consistent with our empirical findings.

312 Is our discovery of divisive normalization by similar orientations actually implemented by the
313 connectivity of neurons in primary visual cortex? The answer to this question could be reflected
314 in the connectivity from inhibitory parvalbumin-expressing (PV) interneurons to pyramidal
315 cells and their relation to neurons' tuning properties. Hofer et al. (2011) find that, in the
316 mouse, pyramidal cells and PV cells are homogeneously connected. Although a weak bias
317 towards orientation tuning is apparent, they conclude that local inhibition in V1 is primarily
318 non-specific. However, despite the connection probability between PV and pyramidal cells
319 being homogeneous, it was found that connection strengths are quite heterogeneous: Individual
320 PV cells strongly inhibit those pyramidal cells that share their visual selectivity (Znamenskiy
321 et al., 2018). This result is in line with our finding of orientation-specific normalization.

322 A limitation of our study is that the stimuli in our dataset are spatially restricted to approx-
323 imately twice the size of the classical receptive field, which prevented us from learning the
324 influence of the surround on normalization. Moreover, we here focused on single images to
325 predict a spike count in a relatively short time window covering the transient response, and
326 ignored any temporal aspects or more sustained periods of the response. These limitations,
327 however, are imposed by the available data – the modeling approach generalizes very well to
328 cover both the surround and the temporal structure – and thus should be addressed in future
329 work.

330 In conclusion, we developed a model consisting of one layer of subunits followed by learned
331 orientation-specific divisive normalization, which accounted remarkably well for the V1 data.
332 We hope that this quantitative approach of evaluating theories of computation in the brain by
333 formalizing them as (components of) trainable predictive models will be used more widely in
334 the future, so the field will (slowly) converge to an accurate and interpretable general-purpose
335 model of the visual system applicable to natural inputs.

336 4 Methods

337 4.1 Experimental details

338 We used the dataset described in detail in Cadena et al. (2019) and provide a summary of the
339 most important characteristics here. Electrophysiological recordings from two healthy, male
340 rhesus macaque monkeys aged 12 and 9 years were performed with a 32-channel linear silicon
341 probe. The monkeys were head-fixed and placed in front of a screen. They were trained to
342 fixate on a target located at the center of the screen. The start of a trial was determined by
343 maintained fixation on the target for 300 ms. The fixation tolerance was set to 0.42° around
344 the center of the target. At the beginning of each recording session, population receptive
345 fields were mapped with a sparse random dot stimulus. Each dot was of size 0.12° of visual
346 angle and was presented over a uniform gray background, changing location and light intensity
347 (black or white) randomly every 30 ms. The receptive field profiles per electrode channel were
348 then obtained via reverse correlation (i. e. spike-triggered average). The center location of the

4 Methods

349 population receptive field was subsequently estimated by averaging over channels and fitting
350 a two-dimensional Gaussian to the reverse correlation profiles. Afterwards, this location was
351 used to place the images of the natural stimulus paradigm.

352 The dataset in Cadena et al. (2019) consists of 7250 distinct natural, greyscale images which
353 were presented two to four times each. A fifth of these images (1450) were taken from ImageNet
354 (Russakovsky et al., 2015). Four additional texturized images were synthesized from each
355 of them, preserving varying degrees of higher-order statistics. The images were cropped to
356 140 px \times 140 px covering two degrees of visual angle. Before displaying the images on the screen,
357 the images were normalized such that the central 1° (70 px) of each image had the same mean
358 and standard deviation. The mean was set to the screen's mean gray intensity (128) and the
359 standard deviation was set to the average standard deviation of the original images. Pixels
360 with an intensity that fell outside the display's range [0, 255] were clipped. Afterwards, all
361 images were overlaid with a circular mask with a soft cosine fade-out and an aperture with a
362 diameter of 1°.

363 Images were presented for 60 ms with no blanks in between. Neural responses were extracted
364 in time windows of 40–100 ms after image onset (Fig. 2), accounting for typical response
365 latencies in primary visual cortex. The image sequence was randomized with the restriction that
366 consecutive images do not belong to the same type (i. e. natural or one of the four texturized
367 versions).

We discarded a few isolated neurons if their stimulus driven variability was too low. The explainable variance in a dataset is smaller than the total variance because the observation noise prevents even a perfect model to account for all the variance in the data. Thus, targeting neurons that have sufficient explainable variance is necessary to train meaningful models of visually driven responses. For a neuron's spike count r , the explainable variance $\text{Var}_{\text{exp}}[r]$ is the difference between the total variance $\text{Var}[r]$ and the variance of the observational noise σ_{noise}^2 ,

$$\text{Var}_{\text{exp}}[r] = \text{Var}[r] - \sigma_{\text{noise}}^2. \quad (2)$$

We estimated the variance of the observational noise by computing the variance of a neuron's response r_t in multiple trials t in which we presented the same stimulus x_j and subsequently taking the expectation E_j over all images,

$$\sigma_{\text{noise}}^2 = E_j [\text{Var}_t [r_t | x_j]] . \quad (3)$$

368 We removed data of neurons if the ratio between the explainable to total variance was below 0.15.
369 The resulting dataset includes spike count data for 166 isolated neurons, with an average ratio
370 of explainable to total variance of 0.285. These neurons were recorded at 1° – 3° eccentricities
371 and had receptive field size diameters between 0.25° and 0.75°.

372 To keep our results of the full DN model without the extension to the surround consistent and
373 comparable to the gold standard baseline from Cadena et al. (2019), we down-sampled the
374 images by a factor of two to train our models. Likewise, images were cropped symmetrically,
375 keeping the 40 \times 40 central pixels. This size covers all of the recorded neurons' receptive fields,
376 with a slight variability in their spatial location. Furthermore, the stimuli light intensities
377 across all pixels and all images were centered around zero and normalized to have unit standard

4 Methods

378 deviation. Additionally, we used the same random dataset splits of Cadena et al. (2019) into
 379 training (64%), validation (16%) and testing (20%). We assessed our models' accuracy for a
 380 specific architecture or set of hyper-parameters in the validation set and we report performance
 381 on the test set. We consistently used the same split throughout our study.

382 4.2 Divisive normalization model

Our model consists of two parts, a nonlinear core and a linear readout (Section 2.1 and Fig. 1). The core (Fig. 1A) processes the input stimulus x by convolving it with 32 filters w_k of size 13 px \times 13 px without padding, defining a bank of features indexed by k . Subsequently, we apply batch normalization without re-scaling (BN*) leading to responses of unit variance (Ioffe and Szegedy, 2015), followed by a rectified linear unit (ReLU) nonlinearity

$$f(\cdot) = \max(0, \cdot) . \quad (4)$$

Hence, the resulting 32 feature maps of size 28 px \times 28 px for the excitatory drive are given by

$$y_k = f(\text{BN}^*(w_k * x)) . \quad (5)$$

383 Many neurons perform similar computations but respond at different localized areas of the
 384 visual field. Those receptive fields are represented by the kernels w_k , which we implemented
 385 convolutionally to exploit this knowledge. Furthermore, the ReLU nonlinearity (Eq. 4) ensures
 386 that all feature maps are positive, $y_k \geq 0$, which is coherent to the biological interpretation of
 387 an excitatory drive.

The feature maps y_k are then normalized divisively to produce 32 output feature maps

$$z_l = \frac{y_l^{n_l}}{\sigma_l^{n_l} + \sum_k p_{kl} \langle y_k^{n_k} \rangle} \quad (6)$$

388 shared by all neurons. Here, all operations are element-wise and the scalar semi-saturation
 389 constant $\sigma_l \geq 0$ is learned from the data. To include normalization by other channels k , we first
 390 exponentiate the excitatory feature maps y_k by the scalar $n_k \geq 0$ element-wise, which is learned
 391 from the data as well. Subsequently, low-pass filtering is performed through average pooling in
 392 space with pool-size 5 px \times 5 px, denoted by $\langle y_k^{n_k} \rangle$. We perform this pooling in order to achieve
 393 (approximate) phase invariance of the normalizing input without requiring a large number of
 394 filters with different phases. Subsequently, the results of the low-pass filtering are summed up,
 395 weighted by the normalization weights p_{kl} , and added into the denominator, resembling Eq. (1).
 396 Furthermore, the normalization weights are constrained to be non-negative, $p_{kl} \geq 0$. Together
 397 with $y_k \geq 0$ and $\sigma_l \geq 0$, this ensures that the denominator in Eq. (6) is non-negative, hence
 398 having a well-defined biological interpretation.

We converted the core's output feature maps z_l , shared by all neurons, to the activity of individual neurons via a linear readout for each of them (Fig. 1B). To do so, we factorized the readout into spatial readout weights $a_{uv,i} \geq 0$ and feature readout weights $b_{l,i} \geq 0$ that pick the relevant locations and features,

$$\hat{r}_i = (a_{uv,i} b_{l,i}) z_{uvl} . \quad (7)$$

4 Methods

399 Here, u, v index space and i indexes neurons. This factorization is beneficial because it reduces
 400 the number of parameters in the readout. Also, we wanted to ensure that the readout does
 401 not model any complex computations, which we achieved by this factorization and the non-
 402 negativity of the readout weights. Additionally, we limited complexity by imposing a sparseness
 403 prior on both weights, because each neuron should only respond to its receptive field which
 404 is represented by a sparse spatial readout weight and should not mix many different features
 405 which corresponds to a sparse feature readout weight. The readout can, however, model a
 406 complex cell (Hubel and Wiesel, 1962) by linearly combining multiple channels of the shared
 407 feature space.

To optimize our model's parameters, we *maximized* the log-likelihood of the model's predictions given the data. To do so, we assumed that neurons' spikes are produced by a Poisson process. Our model predicts the average spike count \hat{r} of a neuron, hence the probability of observing r spikes in the experiment is

$$P(r|\hat{r}) = \frac{\hat{r}^r}{r!} e^{-\hat{r}}. \quad (8)$$

From that follows the Poisson log-likelihood

$$\ln P(r|\hat{r}) = \sum_{i,j} (r_i(x_j) \ln \hat{r}_i(x_j) - \ln(r_i(x_j)!) - \hat{r}_i(x_j)) \quad (9)$$

for all neurons i and all stimuli x_j . A neuron's response $r_i \equiv r_i(x_j)$ depends on the stimulus x_j , which we suppress in our further notation for better readability. For implementation reasons, we wanted to *minimize* the Poisson loss function

$$\mathcal{L}_{\text{Poisson}} = \sum_{i,j} (\hat{r}_i - r_i \ln \hat{r}_i), \quad (10)$$

408 which is the *negative* of the Poisson log-likelihood (Eq. 9), where we omitted $\ln(r_i!)$ since this
 409 term does not depend on our model.

Furthermore, two terms regularizing the model's parameters were applied to the loss. We imposed a smoothness prior on the kernels w_k to ensure the spatial continuity of the predictors' receptive fields. The according penalty on the loss for not-smooth weights was determined with a Laplace filter L to be

$$\mathcal{L}_{\text{smooth}} = \sqrt{\sum_{u,v,k} (L * w_k)_{uv}^2}, \quad L = \begin{pmatrix} 0.25 & 0.5 & 0.25 \\ 0.5 & -3 & 0.5 \\ 0.25 & 0.5 & 0.25 \end{pmatrix}. \quad (11)$$

Due to their receptive fields, neurons only respond to a small, localized area of the visual field, which is why we imposed a sparsity regularizer on the spatial readout weights a_{uv} . Furthermore, neurons should only pool from a small set of feature maps to ensure that the readout does not perform complex computations. Thus we imposed a sparsity regularizer on the feature readout weights b_l as well. We achieved this by adding the L_1 -norm of both weights

$$\mathcal{L}_{\text{sparse}} = \sum_i \sum_{u,v,l} |a_{uv}| \cdot |b_l| \quad (12)$$

4 Methods

410 to the loss function.

The final loss function to minimize with respect to our model's parameters is

$$\mathcal{L} = \mathcal{L}_{\text{Poisson}} + \lambda_{\text{smooth}} \mathcal{L}_{\text{smooth}} + \lambda_{\text{sparse}} \mathcal{L}_{\text{sparse}} , \quad (13)$$

411 where λ_{smooth} and λ_{sparse} are hyper-parameters which set the strength of the smoothness and
412 the sparsity regularizer, respectively.

413 4.3 Divisive normalization model extended to normalization from surround

To extend our DN model to capture normalization from the spatial surround of a unit's classical RF, we replaced the weighted sum accounting for normalization (Eq. 6) by a convolution that also covers space, keeping the rest of the original DN model unchanged,

$$z_l = \frac{y_l^{n_l}}{\sigma_l^{n_l} + s_l} , \quad s_l = \sum_k p_{kl..} * \langle y_k^{n_k} \rangle . \quad (14)$$

414 The new shared feature space z_l consist of all element-wise operations where the normalization
415 feature maps s_l represent the strength by which the excitatory drive $y_l^{n_l}$ is normalized. The
416 normalization feature maps are the result of a convolution between $\langle y_l^{n_l} \rangle$ and normalization
417 pool kernels p_{kluv} . These kernels encode which features (indexed by k) are pooled over what
418 spatial extent (indexed by u, v). Note that for an $u \times v = 1 \times 1$ convolutional kernel p , this is
419 equal to the DN model without normalization from the surround (Eq. 6).

420 For a larger convolutional kernel p , the feature maps s have smaller spatial dimensions than the
421 excitatory feature maps y due to the valid convolution. To be able to perform the element-wise
422 division, we symmetrically cropped the excitatory feature maps y so that the resulting feature
423 maps had the same spatial dimensions as s .

424 Additionally, we wanted to keep the complexity (number of parameters) of the linear readout
425 constant for all the size choices of the normalization kernel p . To this end, we slightly modified the
426 image prepossessing: after down-sampling the full images by a factor of two, we symmetrically
427 cropped them to a size that corresponds – after a forward pass through our model – to a shared
428 feature space of spatial dimensions 34 px \times 34 px. In the particular case of a normalization
429 kernel p of size 7 px \times 7 px, the input images needed to be larger than the actual stimulus size
430 to fulfill that constraint. Thus, we removed any offset at the masked out edges of the images
431 by shifting their mean accordingly, and introduced the necessary zero padding. Overall, this
432 process enabled a fair comparison across all sizes of p .

To keep the kernel size of p computationally tractable, we used convolutions with a dilation
factor of five to be able to pool from a relatively large extra-classical RF while using few
parameters. If we would compute the convolution directly on the feature maps $y_k^{n_k}$, the dilation
would lead to a situation in which some elements in the feature map $y_k^{n_k}$ are not accounted
for by the convolution's inner product for one specific position of the convolutional kernel, i. e.
one specific element in the suppression feature maps s_l . To consider all those elements in the
inner product computation of the convolution, we introduced a preceding average pooling with

4 Methods

a $5 \text{ px} \times 5 \text{ px}$ pool size (same as the dilation factor) and stride one. Then, all the information is pooled over and weighted by exactly one weight of the convolutional kernel. In this view, the pools of neighbouring weights of the dilated kernel have coinciding boundaries. So in addition to implementing shift invariance (see Section 4.2), the average pooling makes sure that we do not loose information for the extended DN model. Due to this pooling, a normalization kernel p of spatial size $3 \text{ px} \times 3 \text{ px}$ would spatially cover a normalization pool of size $15 \text{ px} \times 15 \text{ px}$. We further reduced the number of parameters by a rank-two decomposition separating spatial integration c and the feature weighting d ,

$$p_{kluv} = \sum_{m=1}^2 c_{luv,m} \cdot d_{kl,m} . \quad (15)$$

433 Like before, u, v index space and k indexes the features to pool from. We constrained c and d
434 to be non-negative to make sure the denominator in Eq. (14) is strictly non-negative (recall
435 that in Eq. (14) $\langle y_k^{n_k} \rangle \geq 0$). Our motivation to use two normalization pools (indexed by m)
436 was to allow for both a localized feature-non-specific normalization pool and a feature-specific
437 surround normalization as suggested by the standard model of DN. We investigated models
438 with normalization kernel sizes of $1 \text{ px} \times 1 \text{ px}$, $3 \text{ px} \times 3 \text{ px}$, $5 \text{ px} \times 5 \text{ px}$ and $7 \text{ px} \times 7 \text{ px}$ which
439 spatially covered a five times larger normalization pool due to dilation. Those normalization
440 pools covered visual angles of 0.49° , 0.77° , 1.06° and 1.34° , respectively.

441 4.4 Baseline models

442 4.4.1 Black-box convolutional neural network

Since the divisive normalization computation in our model was completely learned from the data, we wanted to compare to a baseline model that is purely data-driven as well. For this, the current state-of-the-art model is a black-box convolutional neural network with three layers (Cadena et al., 2019). Its first convolutional layer consists of a kernel with spatial size of $13 \text{ px} \times 13 \text{ px}$ and for the second and third layer of size $3 \text{ px} \times 3 \text{ px}$ each. All layers use 32 channels, batch normalization (Ioffe and Szegedy, 2015) and ELU nonlinearity (Clevert et al., 2015)

$$\text{ELU}(h) = \begin{cases} h & \text{if } h \geq 0 , \\ \exp(h) - 1 & \text{else .} \end{cases} \quad (16)$$

443 Similar to our model's architecture, the core part of the CNN model results in a nonlinear
444 feature space shared by all neurons which is mapped to each neuron's activity with individual
445 readout weights factorized in spatial and feature weightings. Sparseness of both of them is
446 achieved by adding an L_1 -penalty to the according loss function. This readout differs from ours
447 in having no constraints on the weights and an additional sophisticated point-wise nonlinearity
448 requiring further parameters.

4 Methods

449 4.4.2 Convolutional subunit model

Our convolutional subunit baseline model is structurally a one-layer convolutional neural network with multiple filters followed by a readout. It is exactly the same as our divisive normalization model (Section 4.2) but with the normalization function (Eq. 6) replaced by the identity function

$$z_l = \text{id}(y_l) = y_l. \quad (17)$$

450 Hence, the only difference to our DN model is the lack of normalization. The shared feature
451 space z_l consists of rectified outputs of linear filters (Eq. 5) which approximate simple cells. The
452 subsequent linear readout can sum up those simple cell responses with additional weightings,
453 enabling the model to approximate complex cells (Hubel and Wiesel, 1962). We trained the
454 model with the same loss function (Eq. 13) as the divisive normalization model.

455 4.5 Number of learned parameters

456 In our model and the baseline models, parameters belong either to the core part that is shared
457 by all neurons (Table 1) or to the readout part in which parameters are specific for each
458 individual neuron.

459 Our DN model's first convolution consists of kernels w_k of spatial size $13 \text{ px} \times 13 \text{ px}$ for 32
460 output channels and batch-normalization without re-scaling, adding 32 bias weights. To learn
461 the normalization pool, 32 normalization weights p_{kl} were learned for each of the 32 output
462 channels l . Additionally, we learned 32 semi-saturation constants σ_l and 32 exponents n_l (one
463 for each channel l). Hence, we get $13 \cdot 13 \cdot 32 + 32 + 32 \cdot 32 + 32 + 32 = 6\,528$ weights for the
464 core. The resulting 32 feature maps are of spatial size of $28 \text{ px} \times 28 \text{ px}$ due to no padding in the
465 convolution. Our linear readout is factorized in spatial and feature weights, thus consisting
466 of $28 \cdot 28 + 32 = 816$ parameters per neuron. Note additionally that all weights except for
467 the convolution kernel and the bias are constrained to be non-negative, halving the according
468 weight-space.

469 In the nonspecific divisive normalization model (Section 2.3.1) the normalization weights are
470 constant for a given feature l , that is $p_{kl} = p_l$. Hence, it requires 32 instead of the $32 \cdot 32$
471 normalization weights of the full divisive normalization model. The other parts of the core
472 remain the same, leading to $6\,528 - 32 \cdot 32 + 32 = 5\,536$ weights for the core. The readout and
473 the number of 816 readout weights per neuron is the same for both models.

474 The spatially extended DN model covering the classical receptive field surround requires more
475 parameters. As before, we get from the first convolution, bias weights, exponents and semi-
476 saturation constants $13 \cdot 13 \cdot 32 + 32 + 32 + 32 = 5\,504$ weights. For each normalization pool
477 component (indexed by m), the factorized convolution learning the normalization contains
478 $32 \cdot 32 = 1\,024$ feature normalization weights $d_{kl,m}$ and $32 u v$ spatial normalization weights $c_{luv,m}$.
479 For the control experiments, we used 2 normalization pool components. So, the core consists in
480 total of $5\,504 + 2 \cdot 1\,024 + 2 \cdot 32 u v = 7\,552 + 64 u v$ weights. For the fitted spatial normalization
481 kernel sizes $u = v = 1, 3, 5, 7$ this results in 7 616, 8 128, 9 152 and 10 688 parameters for the core,

4 Methods

482 respectively. The shared feature space for this model is of larger spatial size of $34 \text{ px} \times 34 \text{ px}$.
483 Hence, the factorized readout consists of $34 \cdot 34 + 32 = 1188$ parameters per neuron.

484 The convolutional subunit model is the same as the divisive normalization model but with
485 the divisive normalization function (Eq. 6) replaced by the identity function (Eq. 17). Hence,
486 compared to the DN model, it saves $32 \cdot 32$ normalization weights p_{kl} , 32 semi-saturation
487 constants σ_l and 32 exponents n_l , leading to $6528 - (32 \cdot 32 + 32 + 32) = 5440$ parameters for
488 the core. The number of 816 readout parameters per neuron stays the same as for the divisive
489 normalization model.

490 The black-box CNN was trained on the same data with input stimuli of size $40 \text{ px} \times 40 \text{ px}$
491 (Cadena et al., 2019). We summarize the calculations in the following. The black-box CNN's
492 first convolution uses kernels with spatial size of $13 \text{ px} \times 13 \text{ px}$ and 32 channels as well as 32
493 biases due to batch normalization, leading to $13 \cdot 13 \cdot 32 + 32 = 5440$ parameters. The two
494 subsequent convolutions use kernels of spatial size $3 \text{ px} \times 3 \text{ px}$ with 32 input and output channels
495 as well as 32 biases each, $3 \cdot 3 \cdot 32 \cdot 32 + 32 = 9248$ parameters for each convolution. In total, the
496 core consists of $5440 + 2 \cdot 9248 = 23936$ parameters. To map to the neurons activities, a readout
497 is used that is factorized in space and features with one additional bias term. The utilized
498 nonlinearity is rather sophisticated, adding 50 parameters. In total, the readout consists of
499 $28 \cdot 28 + 32 + 1 + 50 = 867$ parameters per neuron. Since the CNN model uses smaller input
500 images than the DN model (Section 4.1), it requires less spatial readout weights.

501 4.6 Hyper-parameter optimization

502 Our model's accuracy depends on several hyper-parameters. We set the initial learning-rate to
503 10^{-3} and used an early stopping training scheme: We evaluated the Poisson loss (Eq. 10) every
504 100 training steps and after ten iterations of no improvement we decayed the learning-rate by a
505 factor of three. We repeated this four times to follow the same procedure used by Cadena et al.
506 (2019), because they find best validation set accuracy for this approach. For the filters w_k in
507 the first convolution, we found that a size of $13 \text{ px} \times 13 \text{ px}$ was optimal, the same is true for the
508 number of 32 channels.

509 The weight λ_{smooth} of the smoothness penalty (Eq. 11) and the weight λ_{sparse} of the readout
510 sparsity penalty (Eq. 12) in the full loss function (Eq. 13) were extensively cross-validated
511 using the validation set of our data (Section 4.1). After a first coarse grid search for the
512 divisive normalization model, we narrowed down the relevant parameter-space in which we
513 perform a fine-grained search for the divisive normalization, nonspecific divisive normalization
514 and convolutional subunit model. We randomly sampled the smooth-weight λ_{smooth} from a
515 logarithmic uniform distribution in the interval $[10^{-9}, 10^{-3}]$ for the subunit and nonspecific
516 DN model. For the full DN Model we sampled from a logarithmic uniform distribution in
517 $[10^{-9}, 10^{-4}]$. The readout sparse-weight λ_{sparse} was sampled from a logarithmic uniform
518 distribution in the interval $[10^{-9}, 10^{-4}]$ for the subunit and nonspecific DN model, for the full
519 DN model we used a smaller interval of $[10^{-9}, 10^{-5}]$. For all models, we sampled 1000 runs.

520 For the divisive normalization model (Section 4.2), we achieved the highest accuracy for
521 $\lambda_{\text{sparse}} = 2.25 \cdot 10^{-7}$ and $\lambda_{\text{smooth}} = 2.31 \cdot 10^{-9}$. For the convolutional subunit model, we found

4 Methods

522 the optimal parameters to be $\lambda_{\text{sparse}} = 2.59 \cdot 10^{-6}$ and $\lambda_{\text{smooth}} = 4.98 \cdot 10^{-5}$. The optimal
523 weights of the nonspecific DN model were $\lambda_{\text{sparse}} = 3.98 \cdot 10^{-7}$ and $\lambda_{\text{smooth}} = 1.11 \cdot 10^{-5}$.

524 4.7 Accuracy evaluation

525 4.7.1 Average correlation

526 For architecture search, hyper-parameter optimization and the selection of specific models for
527 analysis we evaluated models' accuracies on the validation set with the Pearson correlation
528 coefficient between the measured spike counts and our models' predictions, averaged over
529 neurons. If the prediction for one neuron is constant, the according standard deviation is zero.
530 Hence, the correlation coefficient was not computable due to division by zero. For those neurons,
531 we set the correlation coefficient to zero before averaging. This average correlation measure
532 does not consider observational noise (Eq. 3).

533 4.7.2 Fraction of explainable variance explained

For reporting accuracy values in this paper, we used the data's test set to compute the fraction of explainable variance explained (FEV)

$$\text{FEV} = 1 - \frac{\text{Var}_{\text{res}}[r]}{\text{Var}_{\text{exp}}[r]} \quad (18)$$

which utilizes the variance that is explainable in principle, $\text{Var}_{\text{exp}}[r]$ (Eq. 2), and the variance of the residuals corrected by the observation noise,

$$\text{Var}_{\text{res}}[r] = \sum_i^N (r_i - \hat{r}_i)^2 / N - \sigma_{\text{noise}}^2. \quad (19)$$

534 This measure corrects for observation noise, which variance σ_{noise}^2 we estimated with Eq. (3).

535 4.8 Evaluation of orientation-specific normalization

To analyse how the preferred orientation of the features being normalized depend on that of the features providing normalizing inputs (Fig. 4–6), we determined for each feature map whether it extracts oriented features and – if so – its preferred orientation. To do so, we windowed each convolutional kernel with a Gaussian window (SD: 3 px), normalized it and then computed its 2D power spectrum (using the discrete Fourier transform with 64×64 samples). We then quantified how power spectral density is distributed across orientations by computing a mean resultant vector m given by:

$$m = \sum_{u,v \in R} F_{uv} e^{2i\phi}, \quad (20)$$

4 Methods

536 where F_{uv} is the Fourier transformed kernel, $R = \{(u, v) : 0.3 < \sqrt{u^2 + v^2} < 0.7\}$ contains all
537 frequencies between 0.3 and 0.7 (with 1.0 being the Nyquist frequency), $\phi = \text{atan2}(v, u)$ is
538 the orientation, i the imaginary unit and the factor 2 in the complex exponential accounts for
539 the fact that we are interested in orientation, which is periodic in 180° or π . If all power in a
540 kernel is concentrated in one orientation, the mean resultant vector will be long, whereas an
541 unoriented kernel will have a mean resultant vector near zero. Based on visual inspection of the
542 kernels in one model fit, we found $m = 0.4$ to be a reasonable threshold for separating oriented
543 from unoriented features and used it as a heuristic for further analyses. We did not explore
544 other thresholds to avoid issues with multiple comparisons and post-hoc statistical testing.
545 To quantify how strong a feature l is normalized by other features k , we computed the average
546 normalizing input, which is given as the expected value (over images) of the product $p_{kl} \cdot y_{kuv}(x)$
547 in Eq. (1). Since this normalization input depends on the stimulus, we computed its expected
548 value of all images in the validation set. We removed the dependence on space by averaging
549 over all locations within the feature map.

550 4.9 Control: All channels contribute to our model's prediction

551 To verify that all features contribute to normalization, we analyzed the readout feature weights
552 for the best ten models (assessed in terms of performance on the validation set). However,
553 there are two issues that prevent a direct comparison across models and neurons of the feature
554 weightings. First, the factorization of the readout into spatial and feature weightings is not
555 unique: scaling the spatial weights (a in Eq. (7)) by a factor β whilst scaling the feature weights
556 (b in Eq. (7)) by $1/\beta$ yields the same output limiting comparisons across neurons. Second, a
557 similar exercise between the normalization weights p and the semi-saturation constant σ (Eq. 6)
558 impedes comparison across models. To solve these issues, we normalized the feature readout
559 weights across channels for this control analysis so that the resulting vectors for each neuron
560 and model convey how much a certain channel contributes to predict a neuron's response
561 compared to the other channels, making the feature readout weights comparable across neurons.
562 Next, we averaged these weights across neurons to assess the importance of the channels in a
563 model. Since these normalized feature readout weights were comparable across both neurons
564 and models, we calculated a collective distribution of the averaged feature readout weights
565 from the best ten models. To make sense of this distribution's absolute values, we evaluated its
566 width in terms of the coefficient of variation, which is the standard deviation in units of the
567 mean.

568 4.10 Implementation details

569 We used Tensorflow (Abadi et al., 2015) to implement models as well as Python, which we
570 additionally used for data analysis. We optimized models with the Adam optimizer (Kingma and
571 Ba, 2014) using mini-batches of size 256. In addition, we used the Python packages Numpy/Scipy
572 (Walt et al., 2011), Pandas (McKinney, 2010), Matplotlib (Hunter, 2007), Seaborn (Waskom
573 et al., 2017) and the tools Jupyter(Kluyver et al., 2016) and Docker (Merkel, 2014).

References

574 Acknowledgements

575 We thank Fabian H. Sinz, Felix A. Wichmann, David A. Klindt, Alexander Böttcher, Robert
576 Geirhos, and Claudio Michaelis for valuable discussions. M.F.G. thanks the International
577 Max Planck Research School for Intelligent Systems (IMPRS-IS). The research was supported
578 by the German Federal Ministry of Education and Research (BMBF) via the Competence
579 Center for Machine Learning (FKZ 01IS18039A); the German Research Foundation (DFG)
580 grant EC 479/1-1 (A.S.E.), the Collaborative Research Center (SFB 1233, Robust Vision)
581 and the Cluster of Excellence “Machine Learning – New Perspectives for Science” (EXC
582 2064/1, project number 390727645); the Bernstein Center for Computational Neuroscience
583 (FKZ 01GQ1002); the National Eye Institute of the National Institutes of Health under Award
584 Numbers R01EY026927 (A.S.T.), DP1 EY023176 (A.S.T.), and NIH-Pioneer award DP1-
585 OD008301 (A.S.T.). This research was also supported by NEI/NIH Core Grant for Vision
586 Research (EY-002520-37), NEI training grant T32EY00700140 (G.H.D) and F30EY025510
587 (E.Y.W.), DARPA grant N66001-17-C-4002, and the Intelligence Advanced Research Projects
588 Activity (IARPA) via Department of Interior/Interior Business Center (DoI/IBC) contract
589 number D16PC00003. The U.S. Government is authorized to reproduce and distribute reprints
590 for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The
591 views and conclusions contained herein are those of the authors and should not be interpreted
592 as necessarily representing the official policies or endorsements, either expressed or implied,
593 of IARPA, DoI/IBC, or the U.S. Government. The funders had no role in study design, data
594 collection and analysis, decision to publish, or preparation of the manuscript.

595 Competing interests

596 The authors declare no competing interests.

597 References

598 M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. Corrado, A. Davis,
599 J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia,
600 R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray,
601 C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke,
602 V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and
603 X. Zheng. Tensorflow: Large-scale machine learning on heterogeneous distributed systems,
604 2015.

605 E. H. Adelson and J. R. Bergen. Spatiotemporal energy models for the perception of motion.
606 *Journal of the Optical Society of America A*, 2: 284–299, 1985.

607 J. Antolík, S. B. Hofer, J. A. Bednar, and T. D. Mrsic-Flogel. Model constrained by visual
608 hierarchy improves prediction of neural responses to natural scenes. *PLoS computational
609 biology*, 12(6): e1004927, 2016.

References

610 H. B. Barlow. Possible principles underlying the transformations of sensory messages. In
611 W. A. Rosenblith, editor, *Sensory Communication*, pages 217–234. MIT Press, Cambridge,
612 Massachusetts, 1961.

613 E. Batty, J. Merel, N. Brackbill, A. Heitman, A. Sher, A. Litke, E. Chichilnisky, and L. Paninski.
614 Multilayer recurrent network models of primate retinal ganglion cell responses. 2016.

615 C. Blakemore and E. A. Tobin. Lateral inhibition between orientation detectors in the cat's
616 visual cortex. *Experimental Brain Research*, 15: 439–440, 1972.

617 A. B. Bonds. Role of inhibition in the specification of orientation selectivity of cells in the cat
618 striate cortex. *Visual Neuroscience*, 2: 41–55, 1989.

619 L. Busse, A. R. Wade, and M. Carandini. Representation of concurrent stimuli by population
620 activity in visual cortex. *Neuron*, 64: 931–942, 2009.

621 D. A. Butts, C. Weng, J. Jin, J.-M. Alonso, and L. Paninski. Temporal precision in the visual
622 pathway through the interplay of excitation and stimulus-driven suppression. *Journal of
623 Neuroscience*, 31(31): 11313–11327, 2011.

624 S. A. Cadena, G. H. Denfield, E. Y. Walker, L. A. Gatys, A. S. Tolias, M. Bethge, and A. S.
625 Ecker. Deep convolutional models improve predictions of macaque v1 responses to natural
626 images. *PLoS computational biology*, 15(4): e1006897, 2019.

627 M. Carandini and D. J. Heeger. Normalization as a canonical neural computation. *Nature
628 Reviews Neuroscience*, 13: 51–62, 2012.

629 M. Carandini, D. J. Heeger, and J. A. Movshon. Linearity and normalization in simple cells of
630 the macaque primary visual cortex. *Journal of Neuroscience*, 17: 8621–8644, 1997.

631 M. Carandini, J. B. Demb, V. Mante, D. J. Tolhurst, Y. Dan, B. A. Olshausen, J. L. Gallant,
632 and N. C. Rust. Do we know what the early visual system does? *Journal of Neuroscience*,
633 25: 10577–10597, 2005.

634 J. R. Cavanaugh, W. Bair, and J. A. Movshon. Nature and Interaction of Signals From the
635 Receptive Field Center and Surround in Macaque V1 Neurons. *Journal of Neurophysiology*,
636 88: 2530–2546, 2002.

637 D.-A. Clevert, T. Unterthiner, and S. Hochreiter. Fast and accurate deep network learning by
638 exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.

639 R. Coen-Cagli, A. Kohn, and O. Schwartz. Flexible gating of contextual influences in natural
640 vision. *Nature Neuroscience*, 18: 1648–1655, 2015.

641 G. C. DeAngelis, J. G. Robson, I. Ohzawa, and R. D. Freeman. Organization of suppression
642 in receptive fields of neurons in cat visual cortex. *Journal of Neurophysiology*, 68: 144–163,
643 1992.

644 G. C. DeAngelis, R. D. Freeman, and I. Ohzawa. Length and width tuning of neurons in the
645 cat's primary visual cortex. *Journal of Neurophysiology*, 71: 347–374, 1994.

References

646 G. H. Denfield, A. S. Ecker, T. J. Shinn, M. Bethge, and A. S. Tolias. Attentional fluctuations
647 induce shared variability in macaque primary visual cortex. *Nature communications*, 9(1):
648 2654, 2018.

649 T. C. Freeman, S. Durand, D. C. Kiper, and M. Carandini. Suppression without Inhibition in
650 Visual Cortex. *Neuron*, 35: 759–771, 2002.

651 E. Froudarakis, P. Berens, A. S. Ecker, R. J. Cotton, F. H. Sinz, D. Yatsenko, P. Saggau,
652 M. Bethge, and A. S. Tolias. Population code in mouse v1 facilitates readout of natural
653 scenes through increased sparseness. *Nature neuroscience*, 17(6): 851, 2014.

654 D. J. Heeger. Normalization of cell responses in cat striate cortex. *Visual Neuroscience*, 9:
655 181–197, 1992.

656 H. W. Heuer and K. H. Britten. Contrast dependence of response normalization in area mt of
657 the rhesus macaque. *Journal of neurophysiology*, 88(6): 3398–3408, 2002.

658 S. B. Hofer, H. Ko, B. Pichler, J. Vogelstein, H. Ros, H. Zeng, E. Lein, N. A. Lesica, and T. D.
659 Mrsic-Flogel. Differential connectivity and response dynamics of excitatory and inhibitory
660 neurons in visual cortex. *Nature Neuroscience*, 14: 1045–1052, 2011.

661 D. H. Hubel and T. N. Wiesel. Receptive fields, binocular interaction and functional architecture
662 in the cat's visual cortex. *The Journal of physiology*, 160: 106–154, 1962.

663 J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in science & engineering*, 9
664 (3): 90–95, 2007.

665 S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing
666 internal covariate shift. In *International conference on machine learning*, pages 448–456,
667 2015.

668 L. Itti, C. Koch, and J. Braun. Revisiting spatial vision: Toward a unifying model. *Journal of
669 the Optical Society of America A*, 17(11): 1899–1917, 2000.

670 S.-M. Khaligh-Razavi and N. Kriegeskorte. Deep supervised, but not unsupervised, models
671 may explain it cortical representation. *PLoS Computational Biology*, 10: e1003915, 2014.

672 W. F. Kindel, E. D. Christensen, and J. Zylberberg. Using deep learning to probe the neural
673 code for images in primary visual cortex. *Journal of vision*, 19(4): 29–29, 2019.

674 D. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint
675 arXiv:1412.6980*, 2014.

676 D. Klindt, A. S. Ecker, T. Euler, and M. Bethge. Neural system identification for large
677 populations separating “what” and “where”. In I. Guyon, U. V. Luxburg, S. Bengio,
678 H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural
679 Information Processing Systems 30*, pages 3506–3516. Curran Associates, Inc., 2017.

References

680 T. Kluyver, B. Ragan-Kelley, F. Pérez, B. Granger, M. Bussonnier, J. Frederic, K. Kelley,
681 J. Hamrick, J. Grout, S. Corlay, P. Ivanov, D. Avila, S. Abdalla, and C. Willing. Jupyter
682 notebooks – a publishing format for reproducible computational workflows. In F. Loizides
683 and B. Schmidt, editors, *Positioning and Power in Academic Publishing: Players, Agents
684 and Agendas*, pages 87 – 90. IOS Press, 2016.

685 J. M. McFarland, Y. Cui, and D. A. Butts. Inferring nonlinear neuronal computation based on
686 physiologically plausible inputs. *PLoS computational biology*, 9(7): e1003143, 2013.

687 L. McIntosh, N. Maheswaranathan, A. Nayebi, S. Ganguli, and S. Baccus. Deep learning models
688 of the retinal response to natural scenes. In *Advances in Neural Information Processing
689 Systems*, pages 1369–1377, 2016.

690 W. McKinney. Data structures for statistical computing in python. In S. van der Walt and
691 J. Millman, editors, *Proceedings of the 9th Python in Science Conference*, pages 51 – 56, 2010.

692 D. Merkel. Docker: Lightweight linux containers for consistent development and deployment.
693 *Linux J.*, 2014(239), Mar. 2014.

694 M. C. Morrone, D. C. Burr, and L. Maffei. Functional implications of cross-orientation inhibition
695 of cortical visual cells. I. Neurophysiological evidence. *Proceedings of the Royal Society of
696 London. Series B. Biological Sciences*, 216: 335–354, 1982.

697 O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy,
698 A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. Imagenet large scale visual recognition
699 challenge. *International Journal of Computer Vision*, 115: 211–252, 2015.

700 N. C. Rust, O. Schwartz, J. A. Movshon, and E. P. Simoncelli. Spatiotemporal Elements of
701 Macaque V1 Receptive Fields. *Neuron*, 46: 945–956, 2005.

702 T. Sawada and A. A. Petrov. The divisive normalization model of V1 neurons: A comprehensive
703 comparison of physiological data and model predictions. *Journal of Neurophysiology*, 118:
704 3051–3091, 2017.

705 H. H. Schütt and F. A. Wichmann. An image-computable psychophysical spatial vision model.
706 *Journal of Vision*, 17: 12–12, 2017.

707 O. Schwartz and E. P. Simoncelli. Natural signal statistics and sensory gain control. *Nature
708 Neuroscience*, 4: 819–825, 2001.

709 E. P. Simoncelli, L. Paninski, J. Pillow, and O. Schwartz. Characterization of neural responses
710 with stochastic stimuli. *The cognitive neurosciences*, 3: 327–338, 2004.

711 F. Sinz and M. Bethge. The conjoint effect of divisive normalization and orientation selectivity
712 on redundancy reduction. In *Advances in Neural Information Processing Systems 21*, pages
713 1521–1528, 2008.

714 J. Touryan, G. Felsen, and Y. Dan. Spatial structure of complex cell receptive fields measured
715 with natural images. *Neuron*, 45(5): 781–791, 2005.

References

716 B. Vintch, J. A. Movshon, and E. P. Simoncelli. A Convolutional Subunit Model for Neuronal
717 Responses in Macaque V1. *Journal of Neuroscience*, 35: 14829–14841, 2015.

718 M. J. Wainwright, O. Schwartz, and E. P. Simoncelli. Natural image statistics and divisive
719 normalization: Modeling nonlinearity and adaptation in cortical neurons. *Probabilistic Models
720 of the Brain: {Perception} and Neural Function*, 2002.

721 S. v. d. Walt, S. C. Colbert, and G. Varoquaux. The numpy array: a structure for efficient
722 numerical computation. *Computing in Science & Engineering*, 13(2): 22–30, 2011.

723 M. Waskom, O. Botvinnik, D. O’Kane, P. Hobson, S. Lukauskas, D. C. Gemperline,
724 T. Augspurger, Y. Halchenko, J. B. Cole, J. Warmenhoven, J. de Ruiter, C. Pye, S. Hoyer,
725 J. Vanderplas, S. Villalba, G. Kunter, E. Quintero, P. Bachant, M. Martin, K. Meyer, A. Miles,
726 Y. Ram, T. Yarkoni, M. L. Williams, C. Evans, C. Fitzgerald, Brian, C. Fonnesbeck, A. Lee,
727 and A. Qalieh. mwaskom/seaborn: v0.8.1 (september 2017), Sept. 2017.

728 B. Willmore, R. J. Prenger, M. C.-K. Wu, and J. L. Gallant. The berkeley wavelet transform:
729 a biologically inspired orthogonal wavelet transform. *Neural computation*, 20(6): 1537–1564,
730 2008.

731 D. L. K. Yamins, H. Hong, C. F. Cadieu, E. A. Solomon, D. Seibert, and J. J. DiCarlo.
732 Performance-optimized hierarchical models predict neural responses in higher visual cortex.
733 *Proceedings of the National Academy of Sciences*, 111: 8619–8624, 2014.

734 Y. Zhang, T. S. Lee, M. Li, F. Liu, and S. Tang. Convolutional neural network models of V1
735 responses to complex patterns. *Journal of computational neuroscience*, pages 1–22, 2018.

736 P. Znamenskiy, M.-H. Kim, D. R. Muir, M. F. Iacaruso, S. B. Hofer, and T. D. Mrsic-Flogel.
737 Functional selectivity and specific connectivity of inhibitory neurons in primary visual cortex.
738 *bioRxiv*, page 294835, 2018.