

Sparse Coding on the Spot: Spontaneous Retinal Waves Suffice for Orientation Selectivity

Jonathan J. Hunt

jjh@42quarks.com

*Queensland Brain Institute, University of Queensland, St Lucia,
Queensland 4072, Australia*

Michael Ibbotson

mibbotson@nrvri.org.au

*National Vision Research Institute, Australian College of Optometry,
Carlton, Victoria 3053, Australia*

Geoffrey J. Goodhill

g.goodhill@uq.edu.au

*Queensland Brain Institute and School of Mathematics and Physics,
University of Queensland, St. Lucia, Queensland 4072, Australia*

Ohshiro, Hussain, and Weliky (2011) recently showed that ferrets reared with exposure to flickering spot stimuli, in the absence of oriented visual experience, develop oriented receptive fields. They interpreted this as refutation of efficient coding models, which require oriented input in order to develop oriented receptive fields. Here we show that these data are compatible with the efficient coding hypothesis if the influence of spontaneous retinal waves is considered. We demonstrate that independent component analysis learns predominantly oriented receptive fields when trained on a mixture of spot stimuli and spontaneous retinal waves. Further, we show that the efficient coding hypothesis provides a compelling explanation for the contrast between the lack of receptive field changes seen in animals reared with spot stimuli and the significant cortical reorganisation observed in stripe-reared animals.

1 Introduction ---

The study of receptive field development in the mammalian primary visual cortex has been an important source of insight into developmental plasticity. There is strong evidence for environmental influence during critical periods (Stryker, Sherk, Leventhal, & Hirsch, 1978; Sengpiel, Stawinski, & Bonhoeffer, 1999; Tanaka, Ribot, Imamura, & Tani, 2006; Hirsch & Spinelli, 1970; Tani & Tanaka, 2008; Hubel & Wiesel, 1965, 1970; Wiesel & Hubel, 1965a, 1965b; Shatz & Stryker, 1978; Stryker, 1978; Shatz, Lindstrom, & Wiesel, 1977;

Kind et al., 2002; Schwarzkopf, Vorobyov, Mitchell, & Sengpiel, 2007; Vorobyov, Schwarzkopf, Mitchell, & Sengpiel, 2007; Mitchell, Kennie, Schwarzkopf, & Sengpiel, 2009; Li, Fitzpatrick, & White, 2006), but receptive field structure develops prior to eye opening (Hubel & Wiesel, 1963; Crair, Gillespie, & Stryker, 1998) and spontaneous retinal waves are essential for normal development (McLaughlin, Torborg, Feller, & O'Leary, 2003; Cang et al., 2005; Huberman, Speer, & Chapman, 2006; Gjorgjieva & Eglen, 2011). The relative contributions of these various influences remain an area of active interest.

One theoretical approach to understanding receptive field development, expounded by Barlow (1961), hypothesizes that the role of early sensory cortex is to encode input in efficient representations for comprehension by higher brain areas. Sparse coding has been a particularly successful implementation of this hypothesis. Olshausen and Field (1996) demonstrated that a sparse coding model was able to learn realistic receptive fields from natural images. More recently, independent component analysis (ICA) has been used (Bell & Sejnowski, 1997; Van Hateren & Van der Schaaf, 1998), which also learns sparse codes when trained on natural scenes (Hyvärinen, Hurri, & Hoyer, 2009). Sparse coding models have been successful at explaining many aspects of receptive field development, including temporal changes (Van Hateren & Ruderman, 1998) and color (Hoyer & Hyvärinen, 2000; Caywood, Willmore, & Tolhurst, 2004). Of importance to the work presented here, Albert, Schnabel, and Field (2008) demonstrated that ICA trained on a simple model of spontaneous retinal waves also learned oriented receptive fields.

Recently, Ohshiro, Hussain, and Weliky (2011) demonstrated an intriguing experimental result that appears to refute the efficient coding hypothesis. They reared ferrets with visual experience consisting of flashing spot stimuli. Despite receiving no oriented external visual experience, these animals developed near-normal levels of oriented receptive fields. Ohshiro and colleagues interpreted this result as a failure of the efficient coding hypothesis and proposed an alternative correlation-based model of receptive field development. Here we demonstrate that the influence of spontaneous retinal waves during development provides an alternative explanation for their findings. We show that an ICA model of receptive field development trained with a mixture of spot stimuli and spontaneous retinal waves continues to learn oriented receptive fields even when retinal waves constitute only a small proportion of the training input.

We also examine why spontaneous retinal waves are able to exert this disproportionate influence on receptive field formation and demonstrate that the sparsity of the spot stimuli is an important consideration in interpreting the Ohshiro et al. result. In particular, we show that another well-known visual manipulation, stripe rearing, results in changes in receptive fields that are more compatible with encoding retinal waves. This provides a coherent explanation for the lack of influence that spot stimuli had on receptive

field development, while stripe rearing results in significant cortical reorganization (Blakemore & Cooper, 1970; Sengpiel et al., 1999; Tanaka et al., 2006).

2 Methods

FastICA (Hyvärinen, 1999) was used for unsupervised learning of static receptive fields. The model was trained on 16×16 pixel image patches. As is standard practice (Hyvärinen et al., 2009), dimensionality was reduced to 100 dimensions using principal component analysis prior to ICA training. The resulting projection matrix was taken to be the receptive fields.

Patches were acquired from 13 images of the natural world (same images as in Hyvärinen & Hoyer, 2001). Retinal waves were simulated using a simple percolation model (Albert et al., 2008). In this model, a fraction p sites on a square lattice were initially marked as potentially active. A starting point was chosen at random, and all potentially active points within a radius r were marked as active. The wave was propagated by iteratively activating potential sites with at least t active points within a distance r until no further sites could be activated. After wave propagation was completed, the result was low-pass-filtered to fill in small holes in the wave (see Albert et al., 2008, for full details of the model). This model was used with parameters $p = 0.48$, $r = 3$, $t = 5$. As in Albert et al., the percolation images were downsampled to 128×128 pixels before patches were extracted. Stripe-rearing input was simulated by filtering the natural scenes with an oriented gaussian filter similar to Hsu and Dayan (2007). Spot patches were created by randomly positioning 6 gaussian spots ($\sigma^2 = 1$ pixel) within each patch. The sign of the spot was chosen at random, and 100,000 patches were used for training each condition.

Receptive field orientation selectivity index and bandwidth were calculated as in Ohshiro et al. (2011) using methods from Wörgötter, Muehle, and Eysel (1991) and Ringach, Shapley, and Hawken (2002). Model training and characterization were implemented in Python using the modular tool kit for data processing (Zito, Wilbert, Wiskott, & Berkes, 2008).

3 Results

Four types of training input were used in this study: natural scene input (see Figure 1A); percolation patterns (see Figure 1B), which are a model of spontaneous retinal waves, spot patches (see Figure 1C); and stripe-filtered natural scenes (see Figure 1D). Spot patches, similar to the training input in Ohshiro et al. (2011), were the only unoriented stimuli.

The ICA model was trained with different combinations of these input types. As expected, receptive fields in the natural case and the percolation case are primarily edge-like (see Figures 1E, and 1F). In agreement with Ohshiro and colleagues, and intuition, the spot stimuli do not lead to

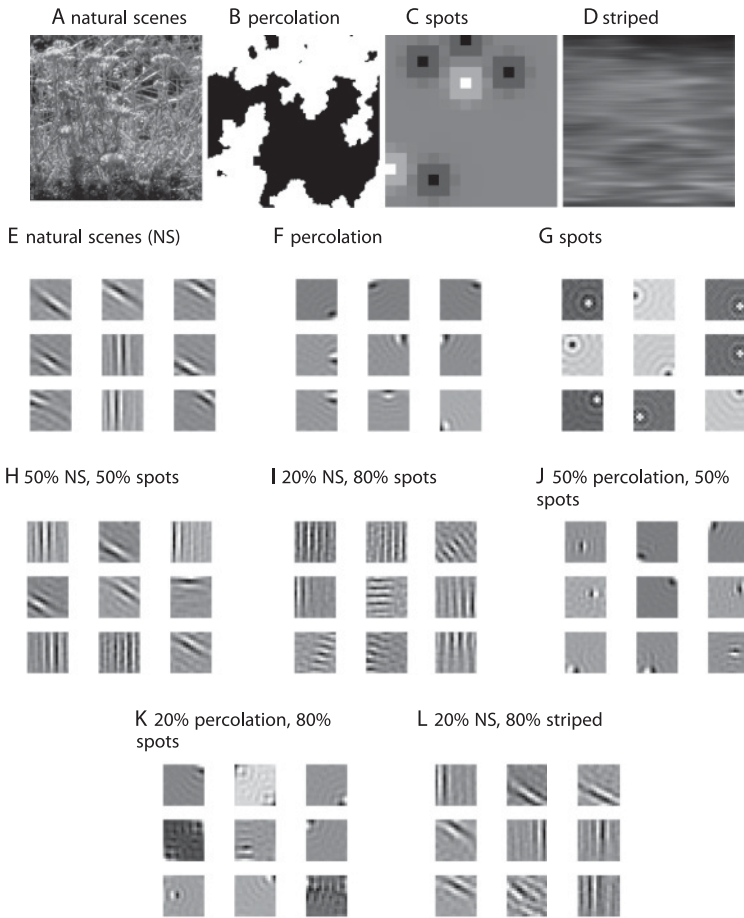


Figure 1: Example inputs and receptive fields. Representative examples of the four input types used for training the ICA model: (A) natural scenes, (B) percolation patterns, (C) spots, and (D) stripe-filtered scenes. (E) Natural scene input leads to strongly oriented receptive fields, as do the (F) percolation patterns, which simulate spontaneous retinal waves. (G) Spot stimuli do not result in oriented receptive fields; however, mixtures of natural scenes and spot stimuli (H, I) or percolation patterns and spot stimuli (J, K) all lead to significant increases in oriented receptive fields (quantified in Figure 2). Stripe-reared mixtures (L) also result in strongly oriented receptive fields.

edge-like receptive fields (see Figure 1G). However, when the ICA model is trained on a mixture of spot stimuli and edge stimuli, significant edge-detecting structure develops (see Figures 1H–1K). This occurs whether the edge stimuli are natural scenes or percolation patches.

This result is quantified in Figure 2, which shows the orientation selectivity and bandwidth for the different conditions. The spot condition leads to an almost complete lack of orientation selectivity and broad receptive fields. However, when even small amounts (20%) of edge-like stimuli are included in the training, there is a dramatic recovery in edge-like structure. This mixture may represent a better approximation to conditions experienced by the animals in Ohshiro et al. (2011), which spent the majority of their life in the dark. The influence of edge-like receptive fields in these mixtures is titrated in Figure 3, which demonstrates that even a small fraction of edge-like input influences receptive field development disproportionately, and 40% edge-like input leads to orientation selectivity levels equal to receptive fields learned from pure edge-like input. This demonstrates that the influence of spontaneous retinal waves may explain the oriented receptive fields found in the spot-reared animals.

The mixtures containing spot stimuli also lead to a broadening of the receptive fields. This broadening does not occur when receptive fields are learned using infomax ICA (data not shown), although in that case, the receptive fields are generally broader in the normal case. Regardless of ICA algorithm, oriented receptive fields develop with only a small amount of edge-like input in the training data.

In contrast to the lack of change observed in spot rearing, several groups have previously demonstrated that animals reared with visual experience limited to a particular orientation develop significant overrepresentation of the exposed orientation (Blakemore & Cooper, 1970; Sengpiel et al., 1999; Tanaka et al., 2006). We examined whether these different outcomes could be understood by the efficient coding hypothesis. We simulated stripe rearing using a mixture that contained 20% unfiltered natural scenes (see Figure 1L), which resulted in a two times overrepresentation of horizontal edges. We examined whether different types of receptive fields are more or less compatible with representing retinal waves. We found that edge-like receptive fields resulting from training with natural scenes or stripes provided a much sparser representation of percolation patches than the unoriented receptive fields learned from the spot stimuli (see Figure 4). This was true even in the stripe-rearing case when there was significant horizontal overrepresentation. If receptive fields are optimized during development toward sparsity, while receiving a combination of visual input and spontaneous retinal waves, this result may explain why some experimental manipulations, such as stripe rearing, result in dramatic receptive field changes while other, less compatible manipulations result in a compromise that retains significant edge structure.

4 Discussion

We have demonstrated that an efficient coding framework can explain the development of oriented receptive fields even in animals reared without




Figure 2: Orientation selectivity and bandwidth. The (A) orientation selectivity and (B) orientation bandwidth of the receptive fields for the different rearing conditions (NS stands for natural scene input, perc for percolation). The spot condition, trained without oriented input, has little orientation selectivity and broad orientation tuning. However, adding even a small proportion of edged input (20% percolation, 80% spots) results in a significant increase in orientation selectivity ($p < 0.0001$) and a significant reduction in orientation bandwidth ($p < 0.0001$) to near-normal levels ($p = 0.07$ for the significance of bandwidth differences between natural scenes condition and 20% percolation, 80% spots). When natural scenes rather than percolation mixtures are used, the results show an even sharper return to normal levels of orientation selectivity. Error bars show SEM. p -values are calculated with a two-sided unpaired t -test.

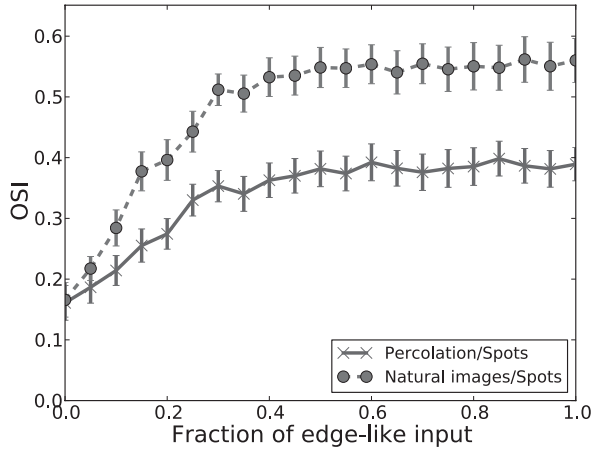


Figure 3: Titration of orientation selectivity with input mixtures. The orientation selectivity of the learned receptive fields increases rapidly with a small fraction of edge-like training input and asymptotes at $\sim 40\%$ edge-like input. This occurs for both percolation mixtures and natural scenes, although natural scene input has a stronger effect.

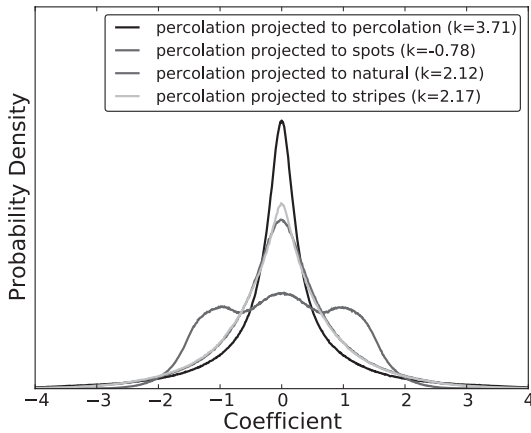


Figure 4: Sparsity of percolation representations. The probability distribution function of coefficients for percolation patches when represented by different sets of receptive fields. The kurtosis value (k) is given in the legend. Spot-like receptive fields are extremely poor at representing edge-like stimuli such as percolation patches sparsely (negative kurtosis). Other edge-like receptive fields are able to represent other types of edged input relatively sparsely (large positive kurtosis values).

exposure to oriented input, provided spontaneous retinal waves are included. The input mixture used here is not identical to the experience of the animals, and any neural implementation of sparse coding may differ markedly from ICA; however, these details are unlikely to be crucial to the results. Our key finding is that sparse coding with a mixture of edge types may lead to important and nonintuitive changes in receptive field structure.

Additionally, we have provided a possible explanation for why animals reared with spot-like stimuli show few changes in receptive field structure, while stripe rearing leads to large cortical reorganization. We demonstrated that receptive fields learned from striped stimuli are better at representing edged input, such as retinal waves, sparsely than the unoriented receptive fields learned from spot stimuli.

As in previous work (Frégnac & Imbert, 1978), Ohshiro et al. (2011) found that animals with no light experience had reduced orientation selectivity, which would seem to argue against an instructive role for retinal waves. However, there is evidence that light exposure may be necessary to retain pre-eye-opening receptive field structure (Frégnac & Imbert, 1978; Crair et al., 1998), which may explain the lack of oriented receptive fields in this case. We have not attempted to include this mechanism in our model. The case for an instructive influence of retinal waves is strengthened by the finding that time spent in the dark has a protective effect against monocular deprivation (Tanaka et al., 2006).

Ohshiro et al. (2011) proposed an alternative model of receptive field development based on input correlations. Our findings here do not exclude this model; however, we have demonstrated that their experimental findings remain compatible with the efficient coding hypothesis, which has had significant success explaining other aspects of receptive field development, including the changes in receptive fields observed in stripe rearing (Hsu & Dayan, 2007).

As in the models tested by Ohshiro et al. (2011), we reduced complexity by examining only static stimuli, so we could not directly examine the development of direction selectivity. Van Hateren and Ruderman (1998) demonstrated that ICA can learn direction selectivity when trained with natural movies, though recent findings in ferret indicate that direction selectivity may have differing developmental mechanisms from orientation (Li, Hooser, Mazurek, White, & Fitzpatrick, 2008).

Endogenous activity occurs at multiple levels of the visual system (Huberman, Feller, & Chapman, 2008). Here, we have focused on retinal waves as the best-characterized of this intrinsic activity and because it has been demonstrated to provide an intrinsic source of oriented input (Albert et al., 2008). However, intrinsic activity also occurs in the lateral geniculate nucleus (Ohshiro & Weliky, 2006) and long-range correlated activity in the cortex (Chiu & Weliky, 2001). Although not captured by our model, it is possible that these sources also provide a scaffold for receptive field structure

before eye opening and contribute a normalizing effect in animals reared with deprived visual input.

Kaschube et al. (2010) recently demonstrated that many aspects of orientation preference map layouts may be well explained by intrinsic self-organization rules. Such intrinsic rules may provide an explanation for the resilience of cortical development in the face of abnormal visual input. However, as highlighted in section 1, there is ample evidence for plasticity of receptive fields in response to visual input. As we have demonstrated here, the efficient coding hypothesis provides a compelling explanation for the differences in receptive field plasticity to two types of visual input manipulation during development: stripe rearing and spot rearing. This hypothesis has previously been demonstrated to provide a good description of receptive field development (Van Hateren & Van der Schaaf, 1998; Hsu & Dayan, 2007). One possibility is that a self-organizing map seeds receptive fields, which are then refined by intrinsic activity and external visual input.

Overall, these findings highlight the complexity of visual system development. Given the myriad of genetic, environment, and retinal wave influences during development and the significant changes in plasticity at differing time points, it is unlikely that any one simple model provides a complete description of receptive field development. However, here we have demonstrated that the efficient coding hypothesis, coupled with a model of retinal waves, is able to explain many recent experimental findings.

Acknowledgments

We thank Mark V. Albert for helpful discussions regarding retinal wave patterns. This research was supported by the Australian National Health and Medical Research Council (Project grant 525459).

References

- Albert, M. V., Schnabel, A., Field, D. J. (2008). Innate visual learning through spontaneous activity patterns. *PLoS Computational Biology*, *4*, e1000137.
- Barlow, H. B. (1961). Possible principles underlying the transformations of sensory messages. In W. A. Rosenblith (Ed.), *Sensory communication* (pp. 217–234). Cambridge, MA: MIT Press.
- Bell, A. J., & Sejnowski, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vision Res.*, *37*, 3327–3338.
- Blakemore, C., & Cooper, G. F. (1970). Development of the brain depends on the visual environment. *Nature*, *228*, 477–478.
- Cang, J., Renteria, R., Kaneko, M., Liu, X., Copenhagen, D., & Stryker, M. (2005). Development of precise maps in visual cortex requires patterned spontaneous activity in the retina. *Neuron*, *48*, 797–809.

- Caywood, M. S., Willmore, B., & Tolhurst, D. J. (2004). Independent components of color natural scenes resemble V1 neurons in their spatial and color tuning. *J. Neurophysiol.*, *91*, 2859–2873.
- Chiu, C., & Weliky, M. (2001). Spontaneous activity in developing ferret visual cortex in vivo. *Journal of Neuroscience*, *21*, 8906–8914.
- Crair, M. C., Gillespie, D. C., & Stryker, M. P. (1998). The role of visual experience in the development of columns in cat visual cortex. *Science*, *279*, 566–570.
- Frégnac, Y., & Imbert, M. (1978). Early development of visual cortical cells in normal and dark-reared kittens: Relationship between orientation selectivity and ocular dominance. *J. Physiol.*, *278*, 27–44.
- Gjorgjieva, J., & Eglen, S. J. (2011). Modeling developmental patterns of spontaneous activity. *Current Opinion in Neurobiology*, *21*, 1–6.
- Hirsch, H. V. B., & Spinelli, D. N. (1970). Visual experience modifies distribution of horizontally and vertically oriented receptive fields in cats. *Science*, *168*, 869–871.
- Hoyer, P., & Hyvärinen, A. (2000). Independent component analysis applied to feature extraction from colour and stereo images. *Network: Computation in Neural Systems*, *11*, 191–210.
- Hsu, A. S., & Dayan, P. (2007). An unsupervised learning model of neural plasticity: Orientation selectivity in goggle-reared kittens. *Vision Research*, *47*, 2868–2877.
- Hubel, D., & Wiesel, T. (1963). Receptive fields of cells in striate cortex of very young, visual inexperienced kittens. *J. Neurophysiol.*, *26*, 994–1002.
- Hubel, D. H., & Wiesel, T. N. (1965). Binocular interaction in striate cortex of kittens reared with artificial squint. *Journal of Neurophysiology*, *28*, 1041–1059.
- Hubel, D., & Wiesel, T. (1970). The period of susceptibility to the physiological effects of unilateral eye closure in kittens. *Journal of Physiology*, *206*, 419–436.
- Huberman, A., Feller, M., & Chapman, B. (2008). Mechanisms underlying development of visual maps and receptive fields. *Annual Review of Neuroscience*, *31*, 479–509.
- Huberman, A. D., Speer, C. M., & Chapman, B. (2006). Spontaneous retinal activity mediates development of ocular dominance columns and binocular receptive fields in V1. *Neuron*, *52*, 247–254.
- Hyvärinen, A. (1999). Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks*, *10*, 626–634.
- Hyvärinen, A., & Hoyer, P. (2001). A two-layer sparse coding model learns simple and complex cell receptive fields and topography from natural images. *Vision Research*, *41*, 2413–2423.
- Hyvärinen, A., Hurri, J., & Hoyer, P. O. (2009). *Natural image statistics: A probabilistic approach to early computational vision*. New York: Springer.
- Kaschube, M., Schnabel, M., Löwel, S., Coppola, D. M., White, L. E., & Wolf, F. (2010). Universality in the evolution of orientation columns in the visual cortex. *Science*, *330*, 1113–1116.
- Kind, P. C., Mitchell, D. E., Ahmed, B., Blakemore, C., Bonhoeffer, T., & Sengpiel, F. (2002). Correlated binocular activity guides recovery from monocular deprivation. *Nature*, *416*, 430–433.
- Li, Y., Fitzpatrick, D., & White, L. E. (2006). The development of direction selectivity in ferret visual cortex requires early visual experience. *Nat. Neuroscience*, *9*, 676–681.

- Li, Y., Hooser, S. V., Mazurek, M., White, L., & Fitzpatrick, D. (2008). Experience with moving visual stimuli drives the early development of cortical direction selectivity. *Nature*, *456*, 952–956.
- McLaughlin, T., Torborg, C. L., Feller, M. B., & O'Leary, D.D.M. (2003). Retinotopic map refinement requires spontaneous retinal waves during a brief critical period of development. *Neuron*, *40*, 1147–1160.
- Mitchell, D. E., Kennie, J., Schwarzkopf, D. S., & Sengpiel, F. (2009). Daily mixed visual experience that prevents amblyopia in cats does not always allow the development of good binocular depth perception. *Journal of Vision*, *9*, 7–22.
- Ohshiro, T., Hussain, S., & Weliky, M. (2011). Development of cortical orientation selectivity in the absence of visual experience with contour. *Journal of Neurophysiology*, *106*, 1923–1932.
- Ohshiro, T., & Weliky, M. (2006). Simple fall-off pattern of correlated neural activity in the developing lateral geniculate nucleus. *Nat. Neurosci.*, *9*, 1541–1548.
- Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, *381*, 607–609.
- Ringach, D., Shapley, R., & Hawken, M. (2002). Orientation selectivity in macaque V1: Diversity and laminar dependence. *Journal of Neuroscience*, *22*, 5639–5651.
- Schwarzkopf, D. S., Vorobyov, V., Mitchell, D. E., & Sengpiel, F. (2007). Brief daily binocular vision prevents monocular deprivation effects in visual cortex. *Eur. J. Neurosci.*, *25*, 270–280.
- Sengpiel, F., Stawinski, P., & Bonhoeffer, T. (1999). Influence of experience on orientation maps in cat visual cortex. *Nat. Neurosci.*, *2*, 727–732.
- Schatz, C., Lindstrom, S., & Wiesel, T. (1977). The distribution of afferents representing the right and left eyes in the cat's visual cortex. *Brain Res.*, *131*, 103–116.
- Schatz, C. J., & Stryker, M. P. (1978). Ocular dominance in layer IV of the cat's visual cortex and the effects of monocular deprivation. *J. Physiol.*, *281*, 267–283.
- Stryker, M. P. (1978). Postnatal development of ocular dominance columns in layer IV of the cat's visual cortex and the effects of monocular deprivation. *Archives italiennes de biologie*, *116*, 420–426.
- Stryker, M. P., Sherk, H., Leventhal, A. G., & Hirsch, H. V. (1978). Physiological consequences for the cat's visual cortex of effectively restricting early visual experience with oriented contours. *Journal of Neurophysiology*, *41*, 896–909.
- Tanaka, S., Ribot, J., Imamura, K., & Tani, T. (2006). Orientation-restricted continuous visual exposure induces marked reorganization of orientation maps in early life. *Neuroimage*, *30*, 462–477.
- Tani, T., & Tanaka, S. (2008). *Interocular imbalance in visual experience affects orientation map formation in cat visual cortex*, program no. 769.8/kk14. Abstract. Society for Neuroscience Annual Meeting, Washington, D.C.
- Van Hateren, J. H., & Ruderman, D. L. (1998). Independent component analysis of natural image sequences yields spatio-temporal filters similar to simple cells in primary visual cortex. *Proc. Roy. Soc. B*, *265*, 2315–2320.
- Van Hateren, J. H., & Van der Schaaf, A. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc. Roy. Soc. Ba*, *265*, 359–366.

- Vorobyov, V., Schwarzkopf, D. S., Mitchell, D. E., & Sengpiel, F. (2007). Monocular deprivation reduces reliability of visual cortical responses to binocular disparity stimuli. *Eur. J. Neurosci.*, *26*, 3553–3563.
- Wiesel, T. N., & Hubel, D. H. (1965a). Comparison of the effects of unilateral and bilateral eye closure on cortical unit responses in kittens. *Journal of Neurophysiology*, *28*, 1029–1040.
- Wiesel, T., & Hubel, D. (1965b). Extent of recovery from the effects of visual deprivation in kittens. *J. Neurophysiol.*, *28*, 1060–1072.
- Wörgötter, F., Mücke, T., & Eysel, U. (1991). Correlations between directional and orientational tuning of cells in cat striate cortex. *Experimental Brain Research*, *83*, 665–669.
- Zito, T., Wilbert, N., Wiskott, L., & Berkes, P. (2008). Modular toolkit for data processing (MDP): A Python data processing framework. *Frontiers in Neuroinformatics*, *2*, 1–7.

Received December 21, 2011; accepted April 12, 2012.