An account of the genesis of the parvo- and magnocellular division based on early visual experience

Lukas Vogelsang* (lvogelsa@mit.edu)
Department of Brain and Cognitive Sciences, MIT
Cambridge, MA 02139

Marin Vogelsang* (ozaki@mit.edu)
Department of Brain and Cognitive Sciences, MIT
Cambridge, MA 02139

Gordon Pipa (gpipa@uos.de)
Institute of Cognitive Science, Osnabrueck University
Osnabrueck, Germany

Sidney Diamond (spdiamon@mit.edu)
Department of Brain and Cognitive Sciences, MIT
Cambridge, MA 02139

Pawan Sinha (psinha@mit.edu)
Department of Brain and Cognitive Sciences, MIT
Cambridge, MA 02139

$*$ shared first authors

Abstract

The division of the early visual pathway into parvo- and magnocellular systems with distinct response properties is broadly acknowledged to be a prominent organizing principle in the mammalian visual system. Characteristic of this division are differential sensitivities to high spatial frequencies and colors, with parvo cells exhibiting high sensitivity and magno cells exhibiting low sensitivity to both. While this distinction is generally accepted, its genesis remains unclear. Here, we provide a potential account based on trajectories of early sensory development. Specifically, we hypothesize that the temporal confluence of constraints on spatial frequency and chromatic sensitivities during development may play a critical role in shaping neuronal response properties characteristic of this division. Receptive field analyses of deep neural networks trained on developmentally inspired 'biomimetic' protocols strongly support this hypothesis. Further, biomimetic training induces a more human-like classification bias towards global shape rather than local texture information, driven by magnocellular-like units. These results provide a potential account for the emergence of a key aspect of visual pathway organization and have implications for designing improved training procedures for deep networks.

Keywords: parvocellular; magnocellular; deep neural networks; sensory development; neurophysiology

Introduction

Cells in the mammalian visual pathway can be broadly segregated into magnocellular and parvocellular groups (Livingstone & Hubel, 1988). Two key characteristics of this division are color (Wiesel & Hubel, 1966) and spatial frequency sensitivities (Derrington & Lennie, 1984), with the magnocellular group exhibiting low spatial frequency and low chromatic sensitivity, and the parvocellular group exhibiting both high spatial frequency and high chromatic tuning. While the division between the two pathways is widely accepted in neurophysiology, details of its emergence from the molecular and cellular processes of fetal development into its functional use in experiential vision are not yet clearly established. Here, we propose an account of the origin of the magno-parvo distinction based on early developmental trajectories of sensory experience and report computational tests of its predictions.

Akin to other aspects of early visual development, visual acuity (Dobson & Teller, 1978) and color sensitivity (Adams & Courage, 2002) undergo a systematic developmental progression from limited to proficient over the months following birth. We propose that the joint progression of these two dimensions of perceptual function could causally influence how they are encoded in the underlying neural system. Specifically, as the start of visual experience is characterized by both low acuity and poor color sensitivity, the cell response properties emerging early in development could come to jointly encode these two attributes. Later in development, higher acuity and richer color information become available and may be conjoined in the underlying neuronal response properties. Thus, the joint coding of low spatial frequency and low color information in magnocellular units, and high spatial frequency and high color sensitivity in parvocellular units, could be an outcome of the co-occurrence of these features at different time points during visual development. We report tests of this account through simulations with deep neural networks as computational models, which provide a useful methodology for directly probing the consequences of deliberate manipulations of sensory experience and allow for analyses at both the representational and behavioral level.

Methods

We utilized the AlexNet architecture (Krizhevsky, Sutskever, & Hinton, 2017), which was minimally adjusted to feature fewer but larger receptive fields (rfs) in the first convolutional layer (including a total of 48 instead of 96 rfs, each increased from 11x11 to 22x22 pixels). This network was trained on the ImageNet object database (Deng et al., 2009) while undergoing two separate temporal progressions of stimulus properties:

1. In a 'standard' regimen, as a non-developmental control, we trained our network on high-resolution, full-color images for the entire training duration comprising 200 epochs.
2. As part of a developmentally-inspired 'biomimetic' regimen, we trained the network on blurry, achromatic images for the first 100 epochs, and on high-resolution, full-color images for the subsequent 100 epochs.

Following training, we examined the first-layer receptive fields of the two differently trained networks in terms of color and spatial frequency metrics (with high values indicating strong tuning to colors and high spatial frequencies, respectively). Subsequently, we probed the two networks' behavior in terms of their tendency to classify images based on global shape vs. local texture features.

Results

Analysis of receptive field structures

Figure 1 depicts a quantification of individual rfs in terms of their joint frequency and color coding. In the 'standard' network, no clear relationship between the two attributes is evident (if any, an anti-correlation between the two properties becomes apparent). By contrast, in the biomimetic model, a cluster of magno-like rfs is evident (indicated by the red ellipse), featuring cells coding for low spatial frequencies while having low color sensitivity.

It is worth noting that the rfs of the biomimetic model tuned to high spatial frequency or high chromatic content are more heterogeneous than the magno-like rfs. This is in agreement with physiological reports revealing two subpopulations of parvocellular units: interblobs, exhibiting high spatial frequency selectivity and strong orientation tuning, and blobs exhibiting lower spatial frequency selectivity and low orientation tuning

(Hubel & Livingstone, 1990). A closer examination of the biomimetic model in Figure 1, in which colors code for a given rf's orientation selectivity, suggests that the units outside of the magnocellular cluster that are tuned to higher frequencies have a slight tendency to indeed be more orientation-selective than those with lower spatial frequency tuning.

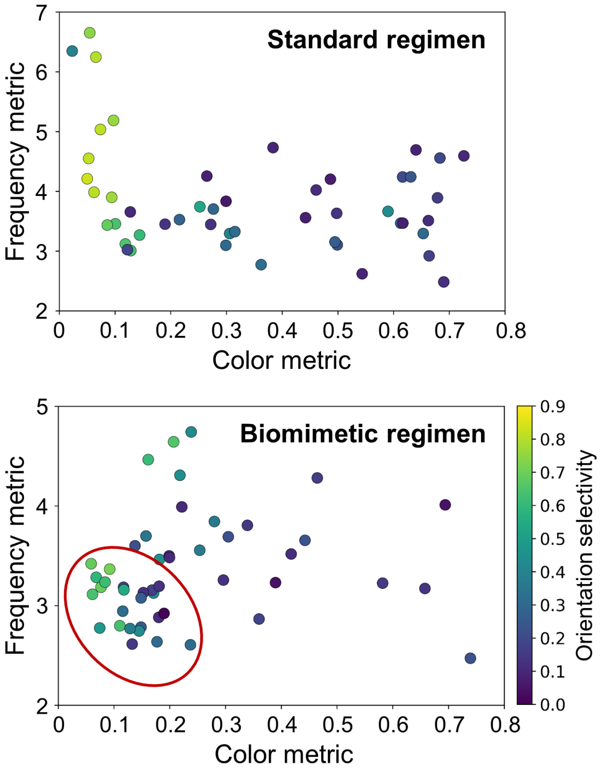

Figure 1: Results of receptive field analysis. Characterization of rf properties in a standard network trained on full-color, full-resolution inputs as well as a network trained on a developmentally-inspired biomimetic progression of inputs. Scatter plots depict the joint frequency and color coding of individual receptive fields, along with their orientation tuning (color-coded). The red ellipse marks a group of consistent magno-like rfs emerging from biomimetic training.

Analysis of classification decisions based on global shape vs. local texture information

From the functional perspective, the magno system is believed to be involved in coarse-grained processing, and the parvo system in fine-grained spatial analysis. Pursuing this idea, we also examined the relationship between receptive fields and network behavior. An important dimension in this regard is that of local texture vs. global shape encoding. We tested whether classification decisions are biased toward texture or shape, using the texture-shape conflict methodology detailed in Geirhos et al. (2018). As is evident in Table 1, and as expected from Geirhos et al. (2018), the standard network does not exhibit a bias to classify images based on shape rather than texture. Notably, biomimetic training leads to a strong and more human-like shape bias. Further, this shape bias appears driven by magno-like rfs, as ablation of as little as the 10 least colorful rfs (thus, the most magno-like units) eliminates the shape bias almost entirely. Ablating the 10 most colorful rfs (thus, the least magno-like units) does not have a comparably strong effect.

Table 1: Results of performance-based analysis. Percentage of total classifications that are correct in terms of texture and, separately, correct in terms of shape. The models 'Biomimetic (magno-abl.)' and 'Biomimetic (parvo-abl.)' thereby refer to the biomimetic model when ablating the 10 least colorful (magno) and 10 most colorful (parvo) rfs, respectively.

Network	Shape correct	Texture correct
Standard	22.4%	23.4%
Biomimetic	31.1%	10.0%
Biomimetic (magno-abl.)	7.8%	4.3%
Biomimetic (parvo-abl.)	23.4%	10.2%

Conclusion

We have presented an account of the potential emergence of the parvo- and magnocellular pathway division based on early developmental trajectories of sensory experience. Our computational results provide evidence supporting this possibility based on the analysis of individual receptive fields. Further, in accordance with the expected psychophysical correlates of this division, we found that magno-like rfs emerging from biomimetic training induced more human-like classification decisions based on global shape rather than local texture information.

More broadly, the finding that training with a developmentally-inspired progression of inputs yields more robust and more human-like representations and performance profiles is well-aligned with the notion that initial degradations during perceptual development may be adaptive and provide a scaffold rather than act as hurdles for the acquisition of later perceptual skills (Turkewitz & Kenny, 1982; Elman, 1993; Newport, 1988; Dominguez & Jacobs, 2003; Vogelsang et al., 2018; Vogelsang, Vogelsang, Pipa, Diamond, & Sinha, 2024).

In conclusion, the work presented here provides a potential account for the genesis of the parvo/magno distinction and thus also presents a teleological perspective on why normal development progresses in the way that it does. It further demonstrates how findings from biological development may help inspire training procedures of computational model systems, as also reviewed in Zaadnoordijk, Besold, and Cusack (2022).

Acknowledgments

This work was supported by grant R01EY020517 from the National Eye Institute (NIH) to Pawan Sinha, and by the Office of the Director of National Intelligence, Intelligence Advanced Research Projects Activity, via [2022-21102100009]. Lukas Vogelsang is supported by a grant from the Simons Foundation International to the Simons Center for the Social Brain at MIT, and Marin Vogelsang is supported by the Japan Society for the Promotion of Science (JSPS), Overseas Research Fellowship.

References

Adams, R. J., & Courage, M. L. (2002). A psychophysical test of the early maturation of infants' mid-and long-wavelength retinal cones. *Infant Behavior and Development*, *25*(2), 247–254.

Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., & Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *2009 ieee conference on computer vision and pattern recognition* (pp. 248–255).

Derrington, A., & Lennie, P. (1984). Spatial and temporal contrast sensitivities of neurones in lateral geniculate nucleus of macaque. *The Journal of physiology*, *357*(1), 219–240.

Dobson, V., & Teller, D. Y. (1978). Visual acuity in human infants: a review and comparison of behavioral and electrophysiological studies. *Vision research*, *18*(11), 1469–1483.

Dominguez, M., & Jacobs, R. A. (2003). Developmental constraints aid the acquisition of binocular disparity sensitivities. *Neural Computation*, *15*(1), 161–182.

Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition*, *48*(1), 71–99.

Geirhos, R., Rubisch, P., Michaelis, C., Bethge, M., Wichmann, F. A., & Brendel, W. (2018). Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*.

Hubel, D. H., & Livingstone, M. S. (1990). Color and contrast sensitivity in the lateral geniculate body and primary visual cortex of the macaque monkey. *Journal of neuroscience*, *10*(7), 2223–2237.

Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2017). Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, *60*(6), 84–90.

Livingstone, M., & Hubel, D. (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science*, *240*(4853), 740–749.

Newport, E. L. (1988). Constraints on learning and their role in language acquisition: Studies of the acquisition of american sign language. *Language sciences*, *10*(1), 147–172.

Turkewitz, G., & Kenny, P. A. (1982). Limitations on input as a basis for neural organization and perceptual development: A preliminary theoretical statement. *Developmental Psychobiology: The Journal of the International Society for Developmental Psychobiology*, *15*(4), 357–368.

Vogelsang, L., Gilad-Gutnick, S., Ehrenberg, E., Yonas, A., Diamond, S., Held, R., & Sinha, P. (2018). Potential downside of high initial visual acuity. *Proceedings of the National Academy of Sciences*, *115*(44), 11333–11338.

Vogelsang, L., Vogelsang, M., Pipa, G., Diamond, S., & Sinha, P. (2024). Butterfly effects in perceptual development: A review of the 'adaptive initial degradation'hypothesis. *Developmental Review*, *71*, 101117.

Wiesel, T. N., & Hubel, D. H. (1966). Spatial and chromatic interactions in the lateral geniculate body of the rhesus monkey. *Journal of neurophysiology*, *29*(6), 1115–1156.

Zaadnoordijk, L., Besold, T. R., & Cusack, R. (2022). Lessons from infant learning for unsupervised machine learning. *Nature Machine Intelligence*, *4*(6), 510–520.