A Developmental Model of Hemispheric Asymmetries of Spatial Frequencies

Ben Cipollini (bcipolli@ucsd.edu)
Department of Cognitive Science, 9500 Gilman Dr., MC 0515, La Jolla, CA 92093

Garrison W. Cottrell (gary@ucsd.edu)
Computer Science and Engineering, 9500 Gilman Dr., MC 0404, La Jolla, CA 92093

Abstract

Lateralization touches virtually every function we think makes us human and interacts fundamentally with development. Here we connect lateralized function to anatomical asymmetries, and connect those anatomical asymmetries to temporal asymmetries in development.

Our differential encoding (DE) model (Cipollini, Hsiao, & Cottrell, 2012; Hsiao, Cipollini, & Cottrell, 2013; Hsiao, Shahbazi, & Cottrell, 2008) shows that lateralization in visual processing of spatial frequencies can be explained by a postulated asymmetry in the spatial spread of connections within retinotopic visual cortex. Here, we present three new modeling results supporting our previous conclusions. First, we show that our model results persist when trained on natural images, warped to match physical distortions of V1, showing that greater biological realism does not diminish our results. Second, we show that the hypothesized anatomical asymmetry can emerge from normal development, due to 1) the known temporal asymmetry in developmental pruning, coupled with 2) known acuity changes. This results in the two hemispheres being trained with images of different spatial frequency content. Third, results from this developmental model suggest that the LH is not specialized for HSF processing; rather, the RH is specialized for LSF processing to the detriment of HSF processing.

Keywords: Lateralization, local/global, spatial frequency, development, double filtering by frequency, differential encoding, visual processing, hemispheric asymmetry

Introduction

Lateralization is an essential part of virtually every function that we believe makes us human. Speaking, fine motor skills, spatial reasoning, emotion, reading, and face perception are all functions with an uneven representation across most individual's cortical hemispheres, but with a consistent hemispheric distribution across the human population.

Lateralization of visual processing, in particular, has long been established (see Ivry and Robertson (1998) for a review). Subjects tend to respond more quickly and accurately to task-relevant low spatial frequency (LSF) information when it is presented to the left visual field (LVF, which the right hemisphere (RH) has preferred access to) vs. the right visual field (RVF, which the left hemisphere (LH) has preferred access to). The opposite holds for task-relevant high spatial frequency (HSF) information. These results fit nicely with LH lateralization for word reading (a HSF task) and RH lateralization for face perception (a task using configural information found in LSFs). The more general inference is generally that the RH is specialized for LSF processing, while the LH is specialized for HSF processing. We believe that understanding mechanisms behind lateralization of spatial frequency (SF) processing may give insight into word reading, face perception, and general mechanisms that may lead to other lateralized functions.

Like lateralization, development is also key to understanding human cognition. Human development differs from that of any other primate (Martin, 1983; D. Geschwind & Rakic, 2013), including extinct homo species such as Neanderthal (Gunz, Neubauer, Maureille, & Hublin, 2010). Developmental disorders come with a wide variety of cognitive impairments, including many involving atypical patterns of lateralization and inter-hemispheric transfer.

How do development and learning interact with hemispheric lateralization of visual processing? Several hypotheses exist. A few are based on data showing that the right hemisphere develops earlier than the left (N. Geschwind & Galaburda, 1985; Hellige, 1993). As Hellige (1993) noted, during that time, the retina is also developing, during which acuity changes from predominantly LSF ranges to adult-like levels. Howard and Reggia (2007) theorized that during this period, magnocellular afferents to visual cortex enervate V2 in the RH, while later-developing parvocellular afferents innervate V2 in the LH to a greater extent, leading to lateralization of spatial frequency processing. Other approaches exist; Plaut and Behrmann (2011) showed that anatomical constraints on wiring length, the differential projection onto the retina of words (central) and faces (peripheral), and the left lateralization of language could lead to lateralization of faces to the RH (Fusiform Face Area) and words to the LH (Visual Word Form Area).

In this paper, we show that the hypothesized asymmetry that leads to lateralization can emerge from a plausible interaction between the asymmetric timing of connection pruning and visual acuity changes. We show this in a biologically plausible model under "natural image" experience and with cortical distortions known to exist in retinotopic visual areas. Then, in order to compare the results of our developmental model to our previous work, we also implemented a more biologically plausible version of our model, also using "natural image" experience and the same cortical distortions.

The Differential Encoding (DE) Model

Our approach and model of lateralization of visual processing was initially a response to the Double Filtering by Frequency (DFF) model by Ivry and Robertson (1998). Following the lead of Sergent (1982), they argued that the hemispheres are generically lateralized for SF processing across modalities, and proposed that lateralization of spatial frequency processing plays a causal role in the local/global effects in hierarchical letter stimuli and in other tasks with information at multiple spatial scales. However, their connectionist implementation of their model simply assumed a spatial frequency bias existed between the hemispheres, without any indication how such frequency biasing could be neurally implemented and exist in each relevant modality.

Inspired by the finding that long-range lateral connections differed in their spatial spread between left and right BA22 (Wernicke's area and its RH homologue) (Galuske, Schlote, Bratzke, & Singer, 2000), we hypothesized that the same asymmetry exists in visual cortex. We then showed in a simple connectionist model how frequency filtering could arise from such a connectivity asymmetry (Cipollini et al., 2012; Hsiao et al., 2013), and could lead to lateralization in classical behavioral tasks (Cipollini et al., 2012; Hsiao et al., 2008, 2013). We also argued that, due to the dependence of lateralization on both task and stimulus features, that long-range lateral connections are most likely involved, as they are hypothesized to enhance stimuli via top-down attention (Li, Piech, & Gilbert, 2008; Piech, Li, Reeke, & Gilbert, 2013) as well as bottom-up processing (Swadlow & Alonso, 2009).

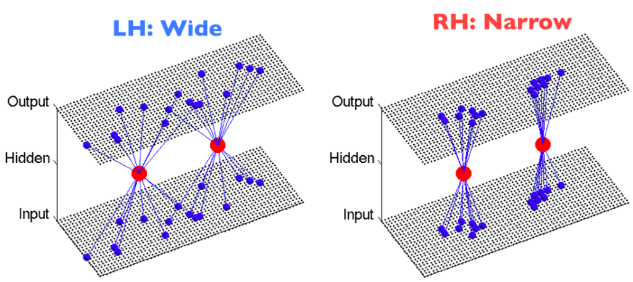

Figure 1: Two (of 850) hidden units for each hemispheric model, each with 8 connections. In our simulations below, each hidden unit has 15 connections.

The differential encoding model is a three layer feed-forward autoencoder model with sparse connectivity between the hidden layer and the input and output layers (Figure 1), where inputs and outputs are images. Each hidden unit has a 2D position in the input/output space and a small, fixed number of connections. Connections for each unit are sampled from a Gaussian distribution centered at the hidden unit's input/output location. The only difference between the LH and RH models is the standard deviation (σ) parameter of the Gaussian distribution: $\sigma_{LH} > \sigma_{RH}$, such that the spatial spread of connections is greater in the LH vs. the RH model. Note that this Gaussian PDF is used to create *connections* between layers and thus is different from the Gaussian receptive field functions used in some models of lateralization (e.g., Ivry and Robertson (1998); Monaghan and Shillcock (2004)). In fact, the difference in connection spread in our model hemispheres (more spread LH connections) is the *opposite* of theirs (e.g. more spread RH connections).

The model is trained using backpropagation of error (see Cipollini et al. (2012) for detailed methods and training parameters). The training task is to reproduce the output image from the input image through the sparse connectivity matrix described above. This forces the images to be recoded in a manner dependent on the sparse connectivity matrix. The hidden unit encoding represents the lateral interaction between nearby retinotopic locations in cortex.

For LH and RH analysis, many networks instances are generated and trained, with their results compiled and analyzed by hemisphere. After training, differences in spectral content of the input and output images indicate lateralized differences in SF encoding abilities (see Hsiao et al. (2013) for detailed methods). Hidden unit encodings are computed for images related to a human behavioral task, and are then used as inputs to independent RH and LH classification networks. These classification networks are trained (using the backpropagation of error algorithm; see Cipollini et al. (2012) for detailed methods) on the same classification task as in the human behavioral task. After training, network performance is summarized over all LH and RH network instances and is then compared to the summary statistics for the human data.

The Developmental DE Model

A primary finding of our previous work is an association between connection spread and spatial frequency processing, where a more spatially constrained connection spread is biased for lower spatial frequency processing. We discovered this by querying what image information is best learned when the connection distribution is varied. Here, we explore the complementary approach: we query what connection distributions are preferred when the spatial frequency content of training images is varied.

Human visual development is an example of this complementary approach. This is due to an interaction among the following three factors:

- Visual acuity / contrast sensitivity is initially poor and improves as the retina develops (see Wang and Cottrell (2012) for a summary).
- Long-range lateral connections are profuse at birth, with die-off of presumably unused connections and strengthening of the remaining connections, occurring during early visual experience (Katz & Callaway, 1992).
- The RH begins maturing earlier than the LH (for reviews, see N. Geschwind and Galaburda (1985); Hellige (1993, 2006)).

Because the RH begins maturing earlier, RH connections are pruned more during blurrier, lower-frequency visual experience, while the LH connections are pruned more when visual acuity is better. This is just the complementary mechanism we described above.

Methods

Here, we construct LH and RH autoencoder models similarly to our previous work. Input images are 34x25 pixel images.

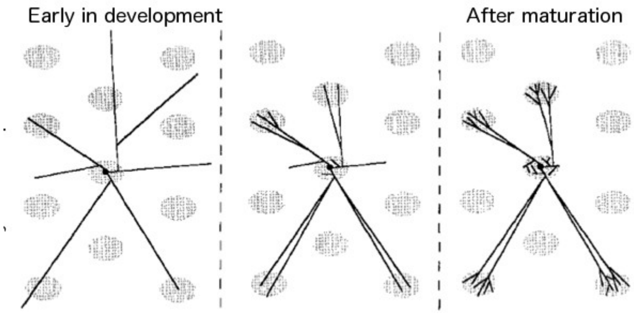

Figure 2: Maturation of long-range lateral connections between "patches" in the developing cat visual cortex. Through visual experience, connections are pruned and elaborated, while synapses are strengthened. Adapted from Katz and Callaway (1992), without permission.

Each model has 850 hidden units distributed across the input/output space, with connections sampled from a Gaussian distribution ($\sigma = 10$ pixels; see top row of Figure 3). Unlike in previous work, connections for LH and RH hidden units are selected from the same Gaussian distribution, simulating initial symmetry between the hemispheres.

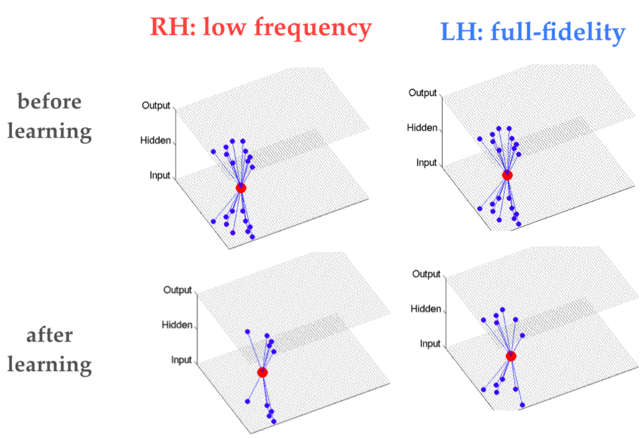

Figure 3: Pruning results differ in the LH and RH models, despite the original connection patterns being identical. This is due to differences in connection removal, induced by different spatial frequency content in the training images.

There are four major differences in the training methods from our previous work[1]:

- Rather than training only on task-specific images (such as hierarchical letter stimuli), we train on 250 natural image patches sampled randomly from the Van Hateren database (Hateren & Schaaf, 1998)[2]. This simulates more accurately the visual experienced gained during development.

- Each hidden unit has 30 connections to start–twice as many as previous models had–and will eliminate synapses until each hidden unit has, on average, 15 input/output connections (see Figure 3). This simulates initial connection proliferation before maturation, followed by elimination during visual experience.

- LH and RH networks differ only in the spectral content of the images they're trained on. Both networks are trained on low-pass images where the image quality improves over time (i.e. the cutoff frequency increases over time), but on average the image quality is higher for the LH network than the RH network (i.e. on average, the cutoff frequency is at a higher frequency for the LH network). The different schedules of training inputs are detailed in Figure 4. Note that iteration 1 of the LH training coincides with iteration 3 of the RH training. This simulates the interaction between changes in visual acuity and hemispheric development.

- In order to simulate the cortical expansion of the fovea, we trained on log-polar version of our original images. The log-polar transform is thought to closely represent retinotopic visual cortex that we aim to simulate (Schwartz, 1985).

After both networks are trained, we compile the empirical connectivity distribution of the unpruned connections across all hidden units within LH and RH models. We compare each distribution with the original connection distribution (before pruning) to see how training on different SF content affected pruning.

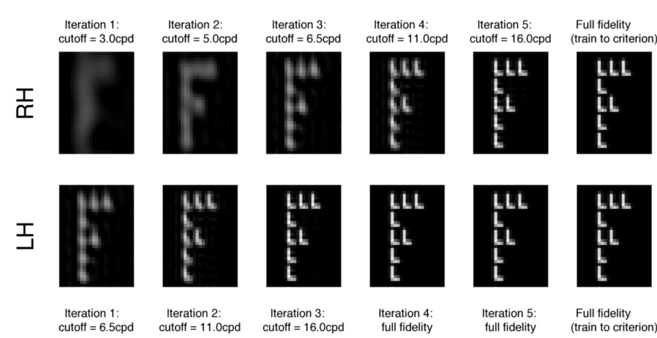

Figure 4: Low-pass filtering schedule of image training. During each iteration, the model was trained on all 250 natural images for 7 epochs. Before moving on to the next iteration, connections containing the smallest $(0.5)^{\frac{1}{6}}\%$ weight values were pruned, such that after the 6 iterations, 50% total connections were pruned. After these 6 training / pruning iterations, both models were trained on full-fidelity natural images until reaching an equal error criterion (summed over all input images and pixels), simulating equal visual experience. Note that hierarchical letter stimuli are pictured here as they show variations in spatial frequency content better than the natural images that were actually used throughout the simulations here.

[1]In addition, weight decay was set to $\lambda = 0.05$, to accentuate differences between used and unused weights.

[2]Greater numbers of image patches were tried and made no qualitative difference in the results, but did increase training time.

In order to compare our developmental model to our previous work, we trained our previous model with the same 250 natural image patches and with 15 connections per hidden unit, just like the developmental model after pruning occurs.[3]. We verified that the models show qualitatively similar results in both frequency processing and behavioral modeling as previously reported (Cipollini et al., 2012; Hsiao et al., 2013), and therefore are appropriate for direct comparison to the developmental model.

Results

Summary

For our network following our previous work, but trained on natural images, we found the same spatial frequency differences as previously reported. We also tested the same network (without retraining) on target detection of letters within hierarchical letter stimuli (Sergent, 1982). These networks showed the same hemisphere × target level interaction as previously found (see Figure 5)[4].

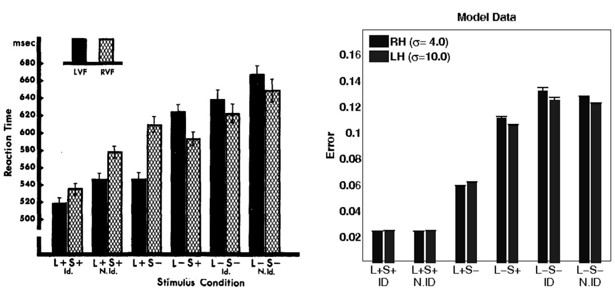

Figure 5: Behavioral results for our previous model, but with autoencoders trained on natural images rather than hierarchical letter stimuli. These results are more consistent with the overall pattern of behavioral results found in Sergent (1982). They are also more consistent across the 6 groupings of $[H, L, T, F]$ into groupings of 2 targets and 2 distractors. Note that we did not test our developmental model on this behavioral task.

For our developmental networks, while we used a complementary approach to the problem, we found the same association between spatial spread of connections and spatial frequency processing: networks trained and pruned under low-frequency images kept connections with a relatively smaller spatial spread than networks trained and pruned on full-fidelity images.

Connection Distributions

In these developmental networks, connection distributions can only differ from variations in visual experience that lead to variations in what connections are pruned. These networks show a difference pattern very similar to our previous model, which had LH and RH connections sampled from Gaussians with different standard deviations. This shows that spatial frequency input differences can drive connectivity differences qualitatively similar to those we had previously postulated, and suggests that these connectivity differences can arise through typical human visual development.

Despite the similar appearance of these connection distributions, the size of the connectivity spread was overall smaller in our developmental model (see Figure 7). In our previous work, LH connections were 30% farther from their nearest connecting neighbor than RH connections on average; here, this number dropped to 5%. We note that re-running the developmental model with a greater difference in the spatial frequency content of the input images can drive connection distance differences equal or greater to the 30% postulated in our previous study.

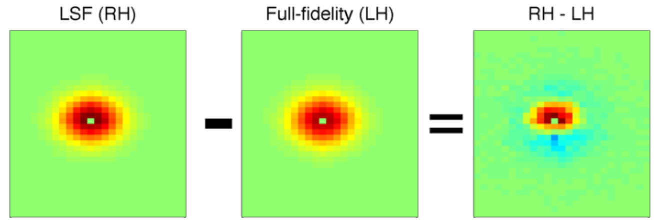

Figure 6: RH and LH connection distributions from the developmental model, and their difference. Here, warm colors are positive, cold colors are negative, and green is zero. To compile the RH and LH distributions, all hidden units were placed at the center of the figure, and a histogram of connections was created. Note in the difference plot the central positive values indicating more short connections in the RH model, and the surrounding blue ring indicating more spread connections in the LH model.

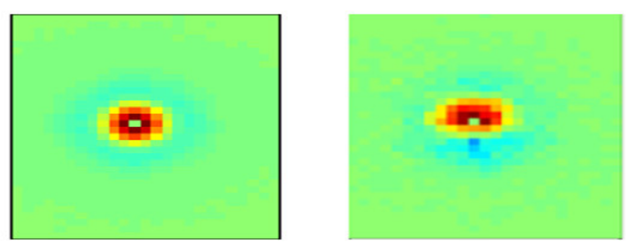

Figure 7: $RH - LH$ connection distribution differences between our previous model trained on natural images (left) and our developmental model (right). Warm colors show connections with greater representations in the RH, cool in the LH.

Spatial Frequency Content

The developmental model also showed spatial frequency differences similar in shape, but attenuated, as compared to those found in our previous work (see Figure 8). We found

[3] $\sigma_{RH} = 4$ pixels, $\sigma_{LH} = 10$ pixels, weight decay $\lambda = 0.025$

[4] In fact, results from this network were more robust to which letters were chosen as targets than in previous work, likely due to a reduction in overfitting of the network due to having a larger training set and more robust regularization procedures.

that this was related to the smaller average connection spread reported above; when the developmental model was re-run on a greater difference in frequency content, the spatial frequency differences met or exceeded those reported in our previous work.

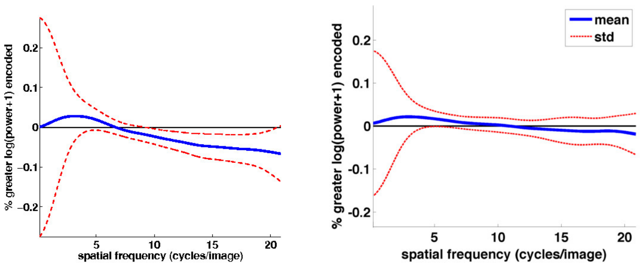

Figure 8: Difference in log-power for 1D spatial frequencies (*RH - LH*); our previous model trained on natural images is on the left, our developmental model is on the right. Note the similar character of both, but attenuated in this developmental scenario. Note also that all spatial frequency differences, besides those very close to the x-axis crossing point, are statistically significant.

Connection Changes

In the literature we've reviewed, it has been consistently suggested that the RH is specialized for low spatial frequency (LSF) processing, and the LH for high spatial frequency (HSF) processing. However, the performance of each hemisphere is measured relative to the other. We don't have a baseline to compare each hemisphere's abilities to, which would be necessary to determine whether both hemispheres are biased, or whether one hemisphere is biased and the other is not.

We can examine this directly issue directly in our developmental model. In Figure 9, the RH (top row) and LH (bottom row) changes over training are shown. We can see that the RH and LH changed similarly, but that the LH network is simply less changed from the original distribution than the right.

This suggests the novel hypothesis that, in fact, the RH is biased towards LSF information at the cost of HSF information, while the LH is essentially similarly, but less biased. Under this hypothesis, the LH only looks specialized for HSF information because it is being compared to the RH, which has sacrificed HSF processing more than the LH has (for the benefit of better LSF processing). We are currently developing a model to examine this hypothesis in greater detail.

Discussion & Conclusions

Here, we described a developmental model of lateralization in visual processing, where improvements in visual acuity interacts the differential timing of connection pruning in left and right hemispheres. In this developmental model, we fixed spatial frequency content and allowing connections to vary via connection pruning during learning. This led to an association between a smaller connection spread, enhanced low

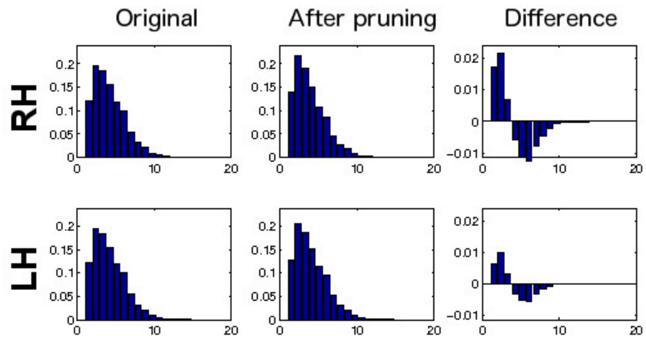

Figure 9: These are histograms of the distance from each connection to the hidden unit location. RH and LH networks begin with the same distribution. Each model hemisphere changes its connection distribution via connection pruning during its (differing) visual experience. The difference between beginning and ending distributions is pictured on the right. Note the similar character of the differences, with the LH network essentially being an attenuated version of the RH network.

spatial frequency processing, and attenuated high spatial frequency processing.

The results of our developmental model are consistent with those found from our previous (adult) models, where connection spread was fixed and spatial frequency processing measured. This consistency in association between spatial frequency processing and connection spread suggests that the assumptions of our adult model could plausibly arise during normal human visual development.

In addition to these findings, we also saw the first evidence that the RH could be specialized for LSF processing at the detriment of HSF processing, while the LH is similarly, but less strongly, biased in how it processes and represents spatial frequency content. In the context of modeling long-range lateral connections, this might suggest that their effect is overall detrimental to LSF processing–consistent with evidence of their important role in contour processing (e.g. (Li et al., 2008)), present in LSFs. Thus, rather than lateralization being about HSF and LSF per se, it may related to computational trade-offs focused in contour processing.

In the future, we plan to follow up on two issues here, and extend this work to central vision:

- We did not test the encodings from our developmental model in any behavioral paradigm. Our first order of business is to verify that the developmental model also shows the behavioral lateralization seen in humans and replicated by other versions of our model.

- We plan to implement a new model to systematically explore how spatial frequency processing relates to spatial spread of connections. This would be a simple 2-layer receptive field model–one output neuron with a sparse set of input connections. We will use this model to map out how spatial spread affects frequency tuning preferences of the output neuron.

- We also hope to explore how interhemispheric connectivity affects the development of lateralization and the interaction between task, stimuli, and measures of functional lateralization. Specifically, we're interested in embedding these connectivity differences in a model with inter-hemispheric interactions, so that we could try and model data for central fixation in the Navon paradigm (Sergent, 1982).

Acknowledgments

This work was partly funded by a Center for Academic Research and Training in Anthropogeny (CARTA) fellowship and by NSF grant SMA 1041755 to the Temporal Dynamics of Learning Center, an NSF Science of Learning Center.

References

Cipollini, B., Hsiao, J., & Cottrell, G. (2012). Connectivity asymmetry can explain visual hemispheric asymmetries in Local/Global, face, and spatial frequency processing. In *Proceedings of the 34th annual meeting of the cognitive science society.* Austin, TX: Cognitive Science Society.

Galuske, R. A., Schlote, W., Bratzke, H., & Singer, W. (2000). Interhemispheric asymmetries of the modular structure in human temporal cortex. *Science (New York, N.Y.)*, *289*(5486), 1946–1949.

Geschwind, D., & Rakic, P. (2013). Cortical evolution: Judge the brain by its cover. *Neuron*, *80*(3), 633–647.

Geschwind, N., & Galaburda, A. M. (1985). Cerebral lateralization. biological mechanisms, associations, and pathology: I. a hypothesis and a program for research. *Archives of neurology*, *42*(5), 428–459.

Gunz, P., Neubauer, S., Maureille, B., & Hublin, J.-J. (2010). Brain development after birth differs between neanderthals and modern humans. *Current Biology*, *20*(21), R921–R922.

Hateren, J. H. van, & Schaaf, A. van der. (1998). Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings of the Royal Society B: Biological Sciences*, *265*(1394), 359–366.

Hellige, J. B. (1993). *Hemispheric asymmetry: What's right and what's left*. Harvard University Press.

Hellige, J. B. (2006). Evolution of brain lateralization in humans. *Cognitie, Creier, Comportament/Cognition, Brain, Behavior*, *10*(2), 211–234.

Howard, M. F., & Reggia, J. A. (2007). A theory of the visual system biology underlying development of spatial frequency lateralization. *Brain and Cognition*, *64*(2), 111–123.

Hsiao, J. H., Cipollini, B., & Cottrell, G. W. (2013). Hemispheric asymmetry in perception: A differential encoding account. *Journal of Cognitive Neuroscience*, 1–10.

Hsiao, J. H., Shahbazi, R., & Cottrell, G. (2008). Hemispheric asymmetry in visual perception arises from differential encoding beyond the sensory level. In *Proceedings of the 30th annual meeting of the cognitive science society.* Austin, TX: Cognitive Science Society.

Ivry, R. B., & Robertson, L. C. (1998). *The two sides of perception*. The MIT Press.

Katz, L. C., & Callaway, E. M. (1992). Development of local circuits in mammalian visual cortex. *Annual Review of Neuroscience*, *15*(1), 31–56.

Li, W., Piech, V., & Gilbert, C. D. (2008). Learning to link visual contours. *Neuron*, *57*(3), 442–451.

Martin, R. D. (1983). *Human brain evolution in an ecological context*. American Museum of Natural History.

Monaghan, P., & Shillcock, R. (2004). Hemispheric asymmetries in cognitive modeling: connectionist modeling of unilateral visual neglect. *Psychological Review*, *111*, 283-308.

Piech, V., Li, W., Reeke, G. N., & Gilbert, C. D. (2013). Network model of top-down influences on local gain and contextual interactions in visual cortex. *Proceedings of the National Academy of Sciences*, *110*(43), E4108–E4117.

Plaut, D. C., & Behrmann, M. (2011). Complementary neural representations for faces and words: A computational exploration. *Cognitive Neuropsychology*, *28*, 251275.

Schwartz, E. (1985, March). On the mathematical structure of the visuotopic mapping of macaque striate cortex. *Science*, *227*(4690), 1065–1066.

Sergent, J. (1982). The cerebral balance of power: confrontation or cooperation? *Journal of Experimental Psychology. Human Perception and Performance*, *8*(2), 253–72.

Swadlow, H. A., & Alonso, J.-M. (2009). Spikes are making waves in the visual cortex. *Nat Neurosci*, *12*(1), 10–11.

Wang, P., & Cottrell, G. W. (2012). A computational model of the development of hemispheric asymmetry of face processing. In *Proceedings of the 35th annual conference of the cognitive science society.* Austin, TX: Cognitive Science Society.