A SPARSE CODING MODEL OF V1 PRODUCES SURROUND SUPPRESSION EFFECTS IN RESPONSE TO NATURAL SCENES

A Thesis

Presented to

The Academic Faculty

by

Allie Del Giorno

In Partial Fulfillment

of the Requirements for the Degree

Bachelor of Science of Electrical Engineering in the

School of Electrical and Computer Engineering

Georgia Institute of Technology

May 2013

A SPARSE CODING MODEL OF V1 PRODUCES SURROUND SUPPRESSION EFFECTS IN RESPONSE TO NATURAL SCENES

Approved by:

Dr. Christopher J. Rozell, Advisor School of Electrical and Computer Engineering *Georgia Institute of Technology*

Dr. Robert Butera School of Electrical and Computer Engineering *Georgia Institute of Technology*

Dr. Bonnie Ferri School of Electrical and Computer Engineering *Georgia Institute of Technology*

Date Approved: _____

To my parents, Mark and Sherry Del Giorno, who continue to instill values that have developed me as a student, researcher, and person.

Acknowledgments

I cannot imagine being more blessed as an undergraduate, and I cannot help but acknowledge all those to whom my success should be attributed.

I wish to first thank the Georgia Tech community. My friends, especially those in the President's Scholarship Program, raised the bar far higher than one that I could have imagined for myself. My teachers have made me realize how exciting the field engineering is and I am now convinced I will never want to leave it. The graduate students here have treated me as an equal and have taught me most of what I know about research. I'd like to especially thank mentors and friends Natalia Toporikova, Mengchen Zhu, Risa Lin, and Yogi Patel for demonstrating that excellence in research and pursuit of outside passions are not mutually exclusive.

I would like to thank faculty that have acted as advisors and mentors, from whom I've learned more at this school than anyone else. They have given me many more opportunities than I could hope to deserve. I would especially like to thank Dr. Chris Rozell and Dr. Rob Butera for dedicating time to training me as a researcher and for so willingly acting as advocates on my behalf. Their ability to balance roles as parents, spouses, researchers, and teachers is an art I hope to master in my future career. Most importantly, I wish to thank my family. My sisters, Jackie and Laura, have been outstanding role models, and I continue to look up to them throughout all stages of life. I cannot thank my parents, Mark and Sherry, enough for putting up with me all these years. They have instilled the discipline and time management skills that have gotten me here and have supported my education without asking for anything in return. I love you all and hope to never forget how much of my success is attributed to you.

Table of contents

Acknowledgments	iv
Introduction vii	
Literature review	1
Human vision and neuroscience	1
Classical receptive fields	2
Surround suppression	3
Systems involved	4
Biological mechanisms	5
Perceptual advantages	6
Information theoretic advantages	7
Efficient coding models and modern signal processing	8
Methods 8	
Stimulus	9
Model	9
Locally Competitive Algorithm (LCA)	9
Receptive fields: Learned dictionary	11
Simulation and output	12
Results 14	
Suppression in time series responses	14
Tuning curves	15
Lifetime sparsity	16
Discussion 17	

Future directions

References 19

Introduction

Recent electrophysiology research has made significant advancements toward revealing the neural basis of early visual processing. The brain is optimized to draw conclusions from natural scenes, and models of the human visual system may uncover principles by which to develop better automated vision systems. In turn, the neuroscience community would benefit from deeper understanding of human vision through the implementation and testing of models of this neural system.

While many neural coding models have been proposed for the primary visual cortex (V1), it remains an open question as to which model best describes the diversity of observed response properties. For instance, the canonical linear-nonlinear model (LN) partially explains some fundamental mechanistic and phenomenological properties of V1, but is unable to explain many nonlinear response properties that are likely associated with the keys to efficient and robust human vision.

Surround suppression is one such nonlinear response property in which visual stimuli extending beyond the classical receptive field (CRF) selectively diminish neural responses. This property has been studied through electrophysiology experiments with synthetic stimuli (e.g., gratings). Surprisingly, high level sparse coding models implemented in a biologically plausible dynamical system have been shown to produce surround suppression effects that match individual and population observed responses. More recently, surround suppression has been investigated experimentally using natural stimuli, and these experiments have shown an increase in the sparsity of measured responses. Despite these findings, it remains unclear whether a functional sparse coding model is sufficient to produce the types of surround suppression observed with natural

vii

stimuli. This thesis demonstrates that the surround suppression effects recently observed with natural stimuli are also emergent properties of a sparse coding model. First, relevant literature in human vision and signal processing will be reviewed. The methods for implementing the model and the results from simulations will then be presented followed by discussion of implications of these results and future work.

Literature review

Human vision and neuroscience

A full understanding of the computations underlying human vision has yet to be established. Studies of human vision have accelerated in recent years due to advancements in hardware and software used to acquire neural signals during electrophysiology experiments. In addition, the computational modeling community has grown in the attempt to develop hypotheses that explain the neural activity collected from electrophysiology and psychophysics studies. The first part of this literature review will summarize findings relevant to research on the primary visual cortex (V1).

Classical receptive fields



that is smaller than stimulation of the classical receptive field alone. Stimulation of the surround only has no effect. **B.** Firing rate decreases with increasing stimulus size beyond the classical receptive field.

The classical linear-nonlinear (LN) model of early vision presumes that each neuron responds independently to a specific stimulus in a localized area of the visual field. The stimulus (image) and corresponding location in the visual field are collectively called the classical receptive field, which resembles a Gabor filter in its canonical form. The traditional model of V1 treats each neuron as a filter where the representation scheme used by the network is a linear combination of the filters weighted by each neuron's response level (firing rate).

Surround suppression

Not all effects can be explained by this idea of independent filters. Surround suppression is one of several effects in which neurons do not behave according to the classical model. These effects are collectively called non-classical receptive field effects and have recently become a substantial research area in vision and other sensory systems. During surround suppression, neurons are inhibited by a stimulus outside their classical receptive field, in an area deemed the 'surround' (Figure 1). The characteristics, mechanisms, and perceptual consequences of this phenomenon are of potential interest to many communities including neurobiology, computational neuroscience, psychology, and computer vision.

Electrophysiology studies have been used to characterize the surround suppression effect. Vision researchers that record neural activity in V1 have seen that spike rates, or neural responses, can be suppressed in as many as 90% of neurons [1,2] by stimuli outside of their surround. In these cells, the spike rates are reduced by as much as 70% [3].

3

Stimulus dependence

The suppressive effect is often dependent on the contrast, orientation, and direction of motion of the stimulus stimulating the surround. These properties are highly dependent on the brain area and the individual neuron being studied. In MT, for instance, cells can be sensitive to the direction and velocity of stimuli up to 50 to 100 times the area of their classical receptive fields [4]. The statistical properties of the stimuli used to probe these neurons affect the properties of the surround as well. Because these areas are so highly interconnected, stimulation of one cell can affect the response properties of other cells, and therefore researchers have become increasingly aware of the choice of stimuli they use in these experiments. In addition to studies with simple stimuli (dots, bars, sinusoidal gratings) [3,5,6], recent studies have used more realistic stimuli (natural scenes) to study these effects [7]. Stimuli that better represent natural scenes tend to induce higher levels of suppression, indicating this effect is tied closely to the properties of natural scenes such as textures and local context.

Systems involved

Surround suppression was formally discovered in the visual pathway, and noticed first by Hubel and Wiesel [5] while mapping receptive fields. Parts of the visual pathway, including V1 and MT, are among the most well-studied.

Surround suppression has also been seen in sensory systems other than vision. One example in somatosensation is surround suppression in the barrel cortex of mice, in which bending one whisker can suppress the response of a neuron responding to a whisker nearby [10]. It has even been seen in the frequency response properties of electoreception in electric fish [11].

Biological mechanisms

The biological mechanisms behind surround suppression have remained a point of contention [10].

Several theories have been proposed for the biological basis of this effect. Based on the diversity of the stimulus characteristics that cause this effect and the variety of responses that are generated, it seems that many mechanisms may be at play (Figure 2).



Lateral connections are connections between neurons in the same layer. There are many of these connections in all areas of the visual system, which means that a neuron representing one piece of the visual field can influence a neuron representing another piece. Even within lateral connections, there are potentially different mechanisms at play. Monocular mechanisms, requiring stimulation in only one eye, may drive this effect with stimuli with high spatial frequency. When the stimulus frequency is lowered, however, binocular mechanisms come into play, where neurons from different eyes may suppress each other [12].

It has been posited that lateral connections are too slow and cover too little of the visual field to fully explain surround suppression [13]. Feedback from higher areas may explain the discrepancies seen in mechanism for surround suppression based purely on lateral connections. There is evidence that inactivation of higher order areas results in reduced strength of surround suppression [13]. At least one model of excitatory connections from higher levels has been formed in the effort to more fully explain surround suppression [14]. However, recurrent feedback is difficult to determine using electrophysiology, and the potential mechanisms at play are not as well studied as feedforward or lateral connections.

Perceptual advantages

Surround suppression likely participates in context-dependent perceptual tasks. Some specific tasks in which surround suppression may aid include: (1) Motion [3] and velocity [15] detection: In areas such as MT and even V1, the selectivity of neurons to the motion of contrasts may play a potential role in representing the structure of moving objects. (2) Contour integration[16]: Detecting continuity of curved and/or 'broken' edges. (2) Texture segregation [17] (3) Perceptual constancies [2]: Recognizing continuity in objects despite changes in lighting, color, or size. (4) Figure-ground segmentation [14]: In this process, local contrast must be used to identify and assign borders.(5) Depth perception (through motion parallax) [2].

These tasks require the use of inputs over wide regions of visual space, meaning that independent responses to small parts of the visual field (a classical linear model of V1) would not be able to produce these effects. There is evidence that surround suppression participates in these tasks by either adjusting the representation of the classical receptive field or representing entirely different features that include both the classical receptive field and the surround. Direct comparison between physiology and psychophysical experiments have been done on several perceptual effects. These include: (1) the reduced apparent contrast of a grating texture embedded in a surrounding grating, (2) target identification when flanked by other features, (3) saliency of broken contours surrounded by edge segments of different orientations, and (4) orientation discrimination when surrounded by features of different orientations and spatial frequencies [19].

Information theoretic advantages

It has recently been shown that stimulation of the surround may support the efficient coding hypothesis proposed by Horace Barlow in 1961 [20]. This hypothesis suggests that the goal of the sensory system is to create an efficient representation of the stimulus. Recently, this has intersected with the idea of a 'sparse' code, one that is represented using the fewest units possible. It has been shown that surround suppression increases the efficiency of transmitting visual information, and may form a sparse code [21]. If many cells respond to parts of the same stimulus, for instance, a lot of redundant information is encoded [22]. The cell needs metabolic energy for each action potential it produces. Therefore, surround suppression likely helps to produce a neural code that is more metabolically efficient. There are additional theoretical advantages, including the removal of statistical redundancy inherent in natural scene statistics, as well as decorrelation of neural responses [7] which means less information to process later in the pathway.

7

Efficient coding models and modern signal processing

Recent developments in modern signal processing share the goal of uncovering efficient methods for encoding signals. Many of these methods formulate efficiency by imposing a sparsity constraint on the representation scheme. Sparsity is the idea that as few coefficients as possible should be used to represent a given signal.

A landmark paper by Olshausen and Field revealed that if a dictionary is learned by imposing a sparsity constraint on their responses to natural scenes, the filters that emerge resemble the receptive fields measured in V1 [23]. This could imply that V1 is using a sparse representation scheme for encoding images. The Locally Competitive Algorithm (LCA) was developed as an efficient way to solve for the optimal representation given a sparsity constraint [24]. Surprisingly, high level sparse coding models implemented in a biologically plausible dynamical system have been shown to produce surround suppression effects that match individual and population observed responses with bars and gratings [25]. Despite these findings, it has remained unclear whether a functional sparse coding model is sufficient to produce the types of surround suppression observed with natural stimuli.

Methods

Simulated electrophysiology experiments were carried out using MATLAB stimulations to investigate whether the LCA model can explain surround suppression effects with natural scenes input.

Stimulus

Natural scenes were obtained from a natural movie stimulus set used in electrophysiology experiments by Charles Cadieu at Berkeley. 10,000 successive images were cropped and then blurred and subsampled to form 32x32 pixel image patches the size of V1 receptive fields. A mask was then applied to each stimulus set, revealing inside circles of radii 0.5 through 3 times the classical receptive field (CRF) size. Each mask was centered on the CRF of the filter, or neuron, being tested (Figure 3).



Model

Locally Competitive Algorithm (LCA)

The Locally Competitive Algorithm (LCA) was implemented in MATLAB. The LCA is one method for generating sparse representations. This algorithm uses a network of linear filters where each filter, ϕ_i , is weighted by a coefficient, a_i (Figure 4). Through

converging network dynamics, the system minimizes the following cost function [24]:

$$\sum_{x,y} \left(I(x,y) - \sum_{i} a_{i}\phi_{i}(x,y) \right)^{2} + \lambda \sum_{i} |a_{i}|$$

Mean-Squared Error (MSE) Sparsity

The mean-squared error term ensures accurate representation and the sparsity term penalizes the amount of active coefficients. Only the filters that best represent the image retain non-zero coefficients. The LCA is an ideal implementation because (1) it is tunable by a single parameter λ , the weight of the sparsity term, and (2) recent work that

shows it can be implemented in real-time hardware [26].



Euler's method was used to find an iterative solution to this optimization using the following update equations:

$$\dot{u}(t) = f(u(t)) = \frac{1}{\tau} \left[b(t) - u(t) - \left(\Phi^t \Phi - I \right) a(t) \right],$$
$$a(t) = T_{\lambda}(u(t)).$$

The variable u(t) is an internal variable representing the excitability of the neuron. The thresholding function T_{λ} is a linear thresholding function in which everything below λ is set to 0 and everything above it is set to u(t)- λ . This not only enforces the sparsity constraint but also ensures that all neurons in the population have nonnegative firing rates a(t). The simulation parameters Δ (time step) and n (number of iterations) were chosen so that $1/(\Delta \cdot n)$, the number of frames per second, matched the physiologically-relevant 30 frames per second of the retina. The LCA converges appropriately when $\Delta/\tau \approx 0.1$, so the parameters chosen for the simulation were $\tau=60$, $\Delta=1$, n=35. The threshold, λ , was set to 0.07 based on empirical testing for discernible image representation and reasonable population sparsity in comparison with electrophysiology data.

Receptive fields: Learned dictionary

The dictionary Φ of filters, or receptive fields, is a collection of 1024 filters of 1024 (32.32) pixels each (Figure 5). This dictionary was learned using the method described in Olshausen and Field [23]. It was optimized for a separate natural scene training set and was provided courtesy of Charles Cadieu.



Simulation and output

28 cells, or dictionary elements, were chosen for simulated stimulation. They were chosen because their receptive fields are localized within the center of the patch so that edge effects of expanding the stimulus would not be a confounding factor. Each cell was fitted using an automated Gabor fit routine, describing its center and radius (Figure 6).



The entire neural population 'viewed' the stimulus centered on the cell's CRF, and the response of the selected neuron was stored. These simulations were repeated for all 8 stimulus sizes (0.5 through 3x CRF size) for each of the 28 cells. The output of the simulation for each test neuron were 8 time series responses, a(t), of 10,000 firing rates or coefficients (one for each frame), from that neuron (Figure 7).



scenes, the cells in the network then compete to represent the image by forming a solution to an optimization function, and a time series of responses is formed for each cell based on its response to the input image sequence.

Results



Suppression in time series responses

Figure 8. A. Neural responses are suppressed as stimulus size increases in a V1 neuron in response to natural scene stimuli [21]. **B**. A sample neuron from the LCA model shows response suppression as stimulus size increases in response to natural scenes. **C.** The same neuron does not show surround suppression in response to white noise stimuli.

Tuning curves

Tuning curves were plotted for each neuron and averaged over all sampled neurons. These tuning curves plot the average response for each size stimulus. Each tuning curve is normalized (maximum value set to 1) for better comparison and averaging. The LCA model shows a decreasing trend beyond the stimulus size similar to that seen in electrophysiology (Figure 9).



Figure 9. Population averages of mean spike rate with increasing stimulus size. A. Mean spike rate decreases as stimulus size increases beyond the CRF in V1 neurons. Mean spike rate across the sampled neurons decreased by 59% from 1x to 3xCRF in this study [21]. **B**. Mean spike rate similarly decreases in the LCA model in response to natural scene input. The radius of the CRF (1xCRF) is the optimal stimulus size. The suppression from 1x to 3xCRF was a 20% reduction in mean response level, smaller than that seen in electrophysiology. **C**. Mean spike rate does not show the same suppression beyond the CRF.

Lifetime sparsity

Lifetime sparsity was calculated for each time series of neural responses, a_i of length n. This metric is based on a metric used to determine the selectivity of neurons given [27]:

$$S = \frac{1 - \frac{\left(\sum_{i=1}^{n} \frac{a_{i}}{n}\right)^{2}}{\sum_{i=1}^{n} \frac{a_{i}^{2}}{n}}}{1 - \frac{1}{n}}$$

S can take values between 0 and 1, where 0 represents low lifetime sparsity (constant firing rate) and 1 represents maximal lifetime sparsity (non-zero for only one frame of the stimulus). With natural scene input, 23 of 28 cells showed increased lifetime sparsity as the stimulus size tripled. With white noise input, however, less than half of the cells (13 of 28) showed the same trend (Figure 10).



Figure 10. A. Lifetime sparsity S increases when the stimulus is expanded to 4x the CRF (CRF + nCRF) [7]. **B**. Lifetime sparsity similarly increases as stimulus size is expanded to 3x the CRF with natural scene input. **C**. Lifetime sparsity shows no definitive trend as stimulus size increases with white noise input (less than half of the points lie above the unity line).

Discussion

It has been shown that the LCA model is able to reproduce several recent results in natural scene surround suppression studies of primate V1 neurons. The similar effects seen in time series suppression, tuning curves, and lifetime sparsity support the hypothesis that a high-level objective of V1 may be to encode visual information in a sparse way.

An important result from these studies is that the natural scene statistics of the input images were crucial to seeing the same trends in surround suppression effects. By using the same model with both white noise and natural scene stimuli, it is made clear that a representation scheme optimized for natural stimuli will likely not produce the same effects when given artificial stimuli. This result supports recent efforts in neuroscience to have subjects perform tasks with natural scenes rather than bars, gratings, or white noise. While the LCA model qualitatively shows the same effects, these effects are not quantitatively as drastic compared to those reported in electrophysiology. There are several potential explanations for this. First, the threshold was empirically chosen to ensure the representation was visually recognizable; if this constraint had been loosened, a larger threshold would have been used that would have created more drastic effects. It is likely that there is more redundancy in V1 filters than was in our dictionary, especially since the dictionary was complete rather than overcomplete (the number of neurons equaled the number of dimensions of the input image). Second, in electrophysiology there is an inherent selection bias toward neurons that respond selectively to stimuli, as their receptive field must be mapped in the visual space before performing the surround suppression experiments. This selection process likely increases the perceived average

17

surround suppression effects. The model has no such selection bias based on selectivity of neural responses because the receptive fields of the neurons are already known.

Future directions

There are several unanswered questions left to be explored. A deeper, low-level understanding of the root cause of these effects in LCA would help determine what mechanisms may be at play within V1. At a higher level, establishing links between surround suppression and perception could reveal what the purpose of this effect is as well as whether reproducing this effect in machine vision systems would lead to benefits in perception. The hope is that if there are significant perceptual advantages to this type of system, a hardware implementation of the LCA could lead to improved performance and efficiency in encoding and interpreting high-dimensional visual data. Finally, other non-classical effects beyond surround suppression could be explored with the LCA model in response to natural scenes.

References

[1] M. J. Cavanaugh JR, Bair W, "Selectivity and spatial distribution of signals from the receptive field surround in macaque V1 neurons.," *J. Neurophys*, vol. 88, no. 5, pp. 2547–56, 2002..

 [2] J. Allman, F. Miezin, and E. McGuinness, "Stimulus specific responses from beyond the classical receptive field: Neurophysiological Mechanisms for Local-Global Comparisons in Visual Neurons," *Ann. Rev. Neurosci.*, vol. 8, pp. 407–30, 1985.

[3] H. E. Jones, K. L. Grieve, W. Wang, and A. M. Sillito, "Surround suppression in primate V1," *J Neurophysiol*, vol. 86, pp. 2011–2028, 2001.

[4] M. E. Allman J, Miezin F, "Direction- and velocity-specific responses from beyond the classical receptive field in the middle temporal visual area (MT).," *Perception*, vol. 14, no. 2, pp. 105–26, 1985.

[5] D. H. Hubel and T. N. Wiesel, "Receptive fields and functional architecture in two nonstriate visual areas (18 and 19) of the cat," *J. Neurophys*, vol. 28, pp. 229–89, Mar. 1965.

[6] F. D. Ozeki H, Finn IM, Schaffer ES, Miller KD, "Inhibitory stabilization of the cortical network underlies visual surround suppression.," *Neuron*, vol. 62, no. 4, pp. 478–92, 2009.

[7] B. Haider, M. R. Krause, A. Duque, Y. Yu, J. Touryan, J. a Mazer, and D. a
McCormick, "Synaptic and network mechanisms of sparse and reliable visual cortical activity during nonclassical receptive field stimulation.," *Neuron*, vol. 65, no. 1, pp. 107–21, Jan. 2010.

[8] R. J. Sundberg KA, Mitchell JF, "Spatial attention modulates center-surround interactions in macaque visual area v4.," *Neuron*, vol. 61, no. 6, pp. 952–63, 2009.
[9] J. N. Hunter and R. T. Born, "Stimulus-dependent modulation of suppressive influences in MT.," *J Neurosci : the official journal of the Society for Neuroscience*, vol. 31, no. 2, pp. 678–86, Jan. 2011.

[10] R. N. S. Sachdev, M. R. Krause, and J. A. Mazer, "Surround suppression and sparse coding in visual and barrel cortices.," *Frontiers in Neural Circuits*, vol. 6, no. 00043, p. 43, Jan. 2012.

[11] M. J. Chacron, B. Doiron, L. Maler, A. Longtin, and J. Bastian, "Non-classical receptive field mediates switch in a sensory neuron's frequency tuning.," *Nature*, vol. 423, no. 6935, pp. 77–81, May 2003.

[12] B. S. Webb, N. T. Dhruv, S. G. Solomon, C. Tailby, and P. Lennie, "Early and late mechanisms of surround suppression in striate cortex of macaque.," *J Neurosci : the official journal of the Society for Neuroscience*, vol. 25, no. 50, pp. 11666–75, Dec. 2005.
[13] A. Angelucci and J. Bullier, "Reaching beyond the classical receptive field of V1 neurons: horizontal or feedback axons?," *J Phys, Paris*, vol. 97, no. 2–3, pp. 141–54, 2003.

[14] T. J. Sullivan and V. R. de Sa, "A model of surround suppression through cortical feedback.," *Neural networks*, vol. 19, no. 5, pp. 564–72, Jun. 2006.

[15] J. A. Perrone, "A neural-based code for computing image velocity from small sets of middle temporal (MT/V5) neuron inputs.," *J Vision*, vol. 12, no. 8, Jan. 2012.

[16] R. Hess and D. Field, "Integration of contours: new insights.," *Trends in Cog Sci*, vol. 3, no. 12, pp. 480–486, Dec. 1999.

[17] D. J. Field, A. Hayes, and R. F. Hess, "Contour integration by the human visual system: evidence for a local 'association field'.," *Vision Research*, vol. 33, no. 2, pp. 173–93, Jan. 1993.

[18] K. M. Supèr H, Romeo A, "Feed-forward segmentation of figure-ground and assignment of border-ownership.," *PloS one*, vol. 5, no. 5, 2010.

[19] P. Seriès, J. Lorenceau, and Y. Frégnac, "The 'silent' surround of V1 receptive fields: theory and experiments.," *J Phys*, vol. 97, no. 4–6, pp. 453–74, 2004.

[20] H. Barlow, "Possible principles underlying the transformation of sensory messages," *Sensory Communication*, pp. 217–234, 1961.

[21] W. E. Vinje and J. L. Gallant, "Natural stimulation of the nonclassical receptive field increases information transmission efficiency in V1.," *J Neurosci*, vol. 22, no. 7, pp. 2904–15, Apr. 2002.

[22] O. Schwartz and E. Simoncelli, "Natural signal statistics and sensory gain control," *Nature Neuroscience*, no. August, pp. 819–825, 2001.

[23] B. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, pp. 607–609, 1996.

[24] C. J. Rozell, D. H. Johnson, R. G. Baraniuk, and B. a Olshausen, "Sparse coding via thresholding and local competition in neural circuits.," *Neural Computation*, vol. 20, no. 10, pp. 2526–63, Oct. 2008.

[25] Zhu M, Rozell C: Biologically realistic excitatory and inhibitory cell properties emerge from a sparse coding network [abstract, CNS 2012]. *BMC Neuroscience* 2012, 13(Suppl 1):P55. [26] Shapero, S., Rozell, C., Balavoine, A., & Hasler, P. (2011). A scalable implementation of sparse approximation on a field programmable analog array. *IEEE BioCAS* (141–144).

[27] W. E. Vinje and J. L. Gallant, "Sparse coding and decorrelation in primary visual cortex during natural vision.," *Science*, vol. 287, no. 5456, pp. 1273–6, Mar. 2000.