



Biologically plausible learning in neural networks with modulatory feedback



W. Shane Grant^{a,*}, James Tanner^a, Laurent Itti^{a,b}

^a Department of Computer Science, University of Southern California, Los Angeles, CA, 90089, USA

^b Department of Psychology and Neuroscience Graduate Program, University of Southern California, Los Angeles, CA, 90089, USA

ARTICLE INFO

Article history:

Received 17 June 2016

Received in revised form 6 December 2016

Accepted 17 January 2017

Available online 28 January 2017

Keywords:

Border ownership
Computational modeling
Feedback
Modulatory
Plasticity
Self-organization

ABSTRACT

Although Hebbian learning has long been a key component in understanding neural plasticity, it has not yet been successful in modeling modulatory feedback connections, which make up a significant portion of connections in the brain. We develop a new learning rule designed around the complications of learning modulatory feedback and composed of three simple concepts grounded in physiologically plausible evidence. Using border ownership as a prototypical example, we show that a Hebbian learning rule fails to properly learn modulatory connections, while our proposed rule correctly learns a stimulus-driven model. To the authors' knowledge, this is the first time a border ownership network has been learned. Additionally, we show that the rule can be used as a drop-in replacement for a Hebbian learning rule to learn a biologically consistent model of orientation selectivity, a network which lacks any modulatory connections. Our results predict that the mechanisms we use are integral for learning modulatory connections in the brain and furthermore that modulatory connections have a strong dependence on inhibition.

© 2017 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

The brain has a remarkable ability to learn to process complicated input through self-organization, and since the studies of Hubel and Wiesel (1963) it has been known that the development of early visual processes is dependent on experience. In the decades since, models of visual development have focused on feedforward pathways, with little attention given to the learning of modulatory connections. Modulatory connections, which adjust existing neuron activations instead of directly driving them, dominate feedback pathways, which themselves constitute a majority of the connections in the brain (Markov et al., 2014). Hebbian-based models have come a long way in explaining potential mechanisms of learning (Clopath, Büsing, Vasilaki & Gerstner, 2010; Hebb, 1949; Widloski & Fiete, 2014), especially in feedforward models of V1 (Stevens, Law, Antolik & Bednar, 2013), but an increasing amount of literature suggests that more comprehensively explaining plasticity requires novel approaches (Lim et al., 2015; Zenke, Agnes & Gerstner, 2015). We will argue that the principles of Hebbian learning, known colloquially as fire together, wire together, cannot be used alone to learn correctly or maintain stability in the context of modulatory connections.

The primary contributions of this work are twofold: the development of a new learning rule that handles modulatory connections, and showing that a stimulus driven feedback model of border ownership can be learned in a biologically plausible way as a result of the new learning rule. The new learning rule, which we call conflict learning, is composed of three conceptually simple, physiologically plausible mechanisms: adjusting plasticity based on the activation of strongly learned connections, using inhibition as an error signal to explicitly unlearn connections, and exploiting several timescales. With border ownership as our prototypical example, we show that a Hebbian learning rule fails to properly learn modulatory connections, while the components of our proposed rule enable it to learn the required connections. Border ownership, which involves the assignment of edges to owning objects, is perhaps one of the earliest and simplest visual processes dependent upon modulatory feedback (Kogo & van Ee, 2014), appearing in V1, V2, and V4 (Zhou, Friedman & Von Der Heydt, 2000). Although many models of its function exist (e.g., lateral models: Sakai and Nishimura (2006); Zhaoping (2005), feedforward: Supér, Romeo and Keil (2010), and feedback: Craft, Schütze, Niebur and Von Der Heydt (2007)) those incorporating feedback are especially promising, integrating well with models of attention (Mihalas, Dong, von der Heydt & Niebur, 2011; Qiu, Sugihara & von der Heydt, 2007) and concepts of grouping (Martin & von der Heydt, 2015). However, until now, all of these models have used fixed, hand-crafted weights, with no demonstration of how the connection patterns for border ownership might be learned.

* Corresponding author.

E-mail addresses: wgrant@usc.edu (W.S. Grant), jetanner@usc.edu (J. Tanner), itti@pollux.usc.edu (L. Itti).

With our new learning rule, we demonstrate that inhibitory modulation of plasticity, in conjunction with competition, is likely a crucial mechanism for learning modulatory connections. Additionally, we show that the rule can be used as a drop-in replacement for a Hebbian learning rule even in networks lacking any modulatory connections, such as an orientation selective model of primary visual cortex. Conflict learning is compared against a recent Hebbian learning based rule (GCAL; Stevens et al. (2013)), which is a good baseline rule for comparison because its weight updates are governed purely by Hebbian logic and it operates at a level of abstraction that captures important physiological behaviors while still being usable in large scale neural network models (e.g., orientation selectivity) and being adaptable for use in new network architectures (e.g., border ownership). We demonstrate that conflict learning, like a Hebbian rule such as GCAL, can be used to learn a biologically consistent model of orientation selectivity. Our results further suggest that networks learned with conflict learning have improved noise and stability responses.

Conflict learning works in a fundamentally different way to previous learning rules by leveraging inhibition as an error signal to dynamically adjust plasticity. Though many existing techniques built upon Hebbian learning, such as those derived from STDP (spike timing-dependent plasticity, Song, Miller and Abbott, 2000) or BCM learning (Bienenstock, Cooper & Munro, 1982), have some method to explicitly control synaptic weakening (e.g., based on signal timing for STDP or comparisons to long term activation averages for BCM), inhibition only indirectly affects learning by lowering activation. Our successful application of the rule to learning models of orientation selectivity as well as border ownership serves as a prediction that modulatory connections in the brain require inhibition and competition to play a bigger role in the dynamics of neural plasticity and activation.

2. Modulatory connections

Modulatory connections are the primary motivation for the development of conflict learning. They are found extensively in feedback projections related to visual processing, for example from visual cortex to the thalamus (Cudeiro & Sillito, 2006; Jones et al., 2012, 2015), from higher visual areas to primary visual cortex (Callaway, 2004; Hupe, James, Girard, Lomber, Payne et al., 2001), as well as from posterior parietal cortex to V5/MT (Friston & Büchel, 2000). Top-down modulatory influences also play a role in phenomena such as attention (Baluch & Itti, 2011; Beuth & Hamker, 2015; Yantis, 2008), object segmentation (Roelfsema, Lamme, Spekreijse & Bosch, 2002), and object recognition (Bar et al., 2006). Attention is a modulatory effect and has the greatest impact on already active representations (Buschman & Kastner, 2015). Modulatory feedback, used in much the same way as in our border ownership experiment, has been used to construct a model of attention that replicates numerous observed attentional effects on both firing rates and receptive field structure (Miconi & VanRullen, 2016).

Modulatory connections can alter the existing activation of a neuron, but cannot cause activity in isolation; they must work in conjunction with driving inputs (Brosch & Neumann, 2014b). We can observe this distinction mathematically by first looking at the activation function for an artificial neuron, which is typically modeled by some function of its weighted inputs:

$$x_j = f\left(\sum_{i \in \text{input}} x_i w_{ij}\right) \quad (1)$$

where w_{ij} is the weight between neurons i and j and x_i is the activation of neuron i .

However, as modulatory connections are defined as those that do not directly drive the activation of a neuron, their effect must be distinguished from driving connections, which, in similar fashion to Brosch and Neumann (2014b), we formalize as:

$$x_j = f(D_j + g(D_j, M_j)) \quad (2)$$

where $D_j = \sum_{i \in \text{driving}} x_i w_{ij}$ and $M_j = \sum_{i \in \text{modulatory}} x_i w_{ij}$. g is a monotonically increasing function with respect to D_j and $D_j = 0$ implies that $g(D_j, M_j) = 0$. Typically, g is a simple product between D_j and M_j (e.g., Bayerl and Neumann, 2004; Brosch and Neumann, 2014a; Roelfsema et al., 2002), hypothesized to be implemented biologically by backpropagation-activated coupling (Larkum, 2013).

When feedforward inputs are taken to be driving and feedback to be modulatory, it can be said that feedback is gated by feedforward, an effect noted by Larkum (2013). Roelfsema et al. (2002) discuss the idea of gating in detail and use it to support a model of figure-ground segregation. This gating allows networks to integrate feedback without struggling to balance it against feedforward input or incurring spurious top-down-driven activation. The physiological mechanics of modulation have been best studied in relation to the thalamus, with a recent review by Varela (2014) showing that modulatory input is extensive and heterogeneous in regard to origin, neurotransmitter, and function. Brosch and Neumann (2014b) discuss the evidence for the potential physiological implementation of modulatory feedback while developing a network-level circuit model for feedforward and feedback interaction.

2.1. Hebbian learning and modulatory connections

Traditional Hebbian based learning rules adapt weights based on some function of the coincidental firing of pre and postsynaptic neurons:

$$\Delta w_{ij} = f(w_{ij}, x_i * g(x_j)) \quad (3)$$

Hebbian learning in its most basic formulation has no mechanism to bound weight growth, making it trivially unstable. For our purposes we use a formulation of Hebbian learning that includes a normalization component for stability, adapted from Stevens et al. (2013):

$$\Delta w_{ij} = \frac{w_{ij} + \eta x_i x_j}{\sum_k (w_{kj} + \eta x_k x_j)} - w_{ij} \quad (4)$$

where η is the learning rate. This weight update, and its normalization, are applied independently to driving and modulatory connections (i.e. all w_{ij} are the same connection type).

To better understand why such a Hebbian rule is not suitable for learning modulatory connections, let us look at the dynamics of a minimal network with two competitive neurons, illustrated in Fig. 1. In this context, competitive means that the neurons are connected such that more active neurons inhibit the activation of those less active through lateral connections. The desired state of this network is to have each competing neuron develop a strong connection to a unique source of modulatory input. It should be noted that this end state is considered desired due to its computational usefulness as a source of top-down information rather than a direct extrapolation from biology.

We can imagine this network as, for example, a simple attention network concerned with detecting apples or oranges in its input. The modulatory connections act as attentional biases towards either apples (M_1) or oranges (M_2). Though one fruit may be desired over the other (e.g., searching for a specific fruit; M_1 active versus M_2), the network has no control over what is present in its input. Features related more to apples (N_1) or to oranges (N_2) may be active regardless of the bias signal, even occurring simultaneously.

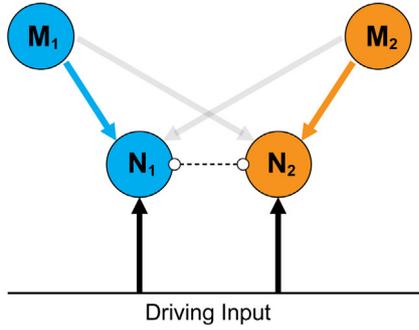


Fig. 1. A simple network with modulatory connections. Neurons N_1 and N_2 receive identical driving input and compete over input from two modulatory neurons, M_1 and M_2 . The colored connections show the desired state of the network, where each competing neuron has learned a unique source of modulatory input. The dashed connection represents lateral inhibition. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

This presents a problem to learning if a pure correlation based rule, like Hebbian learning, is to be used, as the top-down bias is equally correlated with each bottom-up driving input. Learning a unique source of modulatory input is desirable because it allows the attentional biases to affect only the features with which they are semantically associated. With this in mind, let us analyze how this Hebbian learning rule behaves in this network.

The activity of a neuron in the network can be expressed using (2) with a product for $g()$ along with adding divisive inhibition (Carandini & Heeger, 2012) for competition (following Brosch and Neumann, 2014b; Stevens et al., 2013) as well as a noise term:

$$x_j = \frac{D_j + D_j M_j + \epsilon}{1 + \text{Inhib}_j}. \quad (5)$$

We are interested in the dynamics of the network once it has reached the desired state. Let us assume that the neurons have each already learned associations to a unique modulatory input, such that $w_{M_1 N_1} = w_{M_2 N_2} = w_{\max}$ and $w_{M_2 N_1} = w_{M_1 N_2} = w_{\min}$. Because the weights are normalized (see (4)), this configuration implies that $w_{\min} + w_{\max} = 1$.

Without loss of generality, assume that M_1 is highly active while M_2 is inactive, resulting in M_1 sending strong feedback to N_1 . Because of that feedback, N_1 will become more active than N_2 regardless of whether or not it was more active prior. N_2 is then inhibited by N_1 , but because it receives the same driving input, it remains at a lower but non-zero activation. Formally:

$$x_{M_1} > 0 \quad \text{and} \quad x_{N_2} > 0$$

substituting this into (4) gives:

$$\begin{aligned} \Delta w_{M_1 N_2} &= \frac{w_{\min} + \eta x_{M_1} x_{N_2}}{w_{\min} + w_{\max} + \eta x_{M_1} x_{N_2} + 0} - w_{\min} \\ &= \frac{w_{\min} + \eta x_{M_1} x_{N_2}}{1 + \eta x_{M_1} x_{N_2}} - w_{\min} \end{aligned} \quad (6)$$

letting $\alpha = \eta x_{M_1} x_{N_2}$,

$$\begin{aligned} \Delta w_{M_1 N_2} &= \frac{w_{\min} + \alpha}{1 + \alpha} - w_{\min} \\ &= \frac{w_{\min} + \alpha - w_{\min}(1 + \alpha)}{1 + \alpha} \\ &= \frac{(1 - w_{\min})\alpha}{1 + \alpha} \end{aligned} \quad (7)$$

Because $\alpha > 0$ and $1 > w_{\min} \geq 0$, $\Delta w_{M_1 N_2} > 0$. Thus $w_{M_1 N_2}$ is increasing and the system is not in a steady-state. This implies that even if this Hebbian learning rule managed to reach the desired

state, it would not be in equilibrium and would be disrupted by any input.

Compared to this simple example, the modulatory inputs in more general networks will be populations of correlated neurons, and the competing neurons may not all receive identical driving input. Distinct populations are assumed to be weakly correlated with each other (otherwise they would be the same population). The core challenge of learning modulatory connections, however, can be captured by this example using two neurons driven by identical input competing over two independent modulatory inputs.

Implementations of Hebbian learning that restrict weight growth through means other than weight re-normalization, such as the Generalized Hebbian Algorithm (Sanger, 1989), which is closely related to Oja's rule (Oja, 1982), or the BCM rule, which uses an adaptive threshold based on expected average activation to adjust the sign of the weight update, can also be shown to be either unstable or not guaranteed to reach the desired state of this network. We will revisit and analyze these two variations of Hebbian learning in Appendix A after introducing conflict learning in the next section.

3. Introducing conflict learning

Conflict learning was developed to address the demonstrated instability of Hebbian learning rules in the context of modulatory connections, and can be intuitively described as a rule that assigns a unique population of correlated modulatory inputs to each neuron competing over those inputs. It is a general learning rule composed of three conceptually simple, physiologically plausible mechanisms: adjusting plasticity based on the activation of strongly learned connections, using inhibition as an error signal to explicitly unlearn connections, and exploiting several timescales. These concepts are formalized by the following equations:

1. Spreading: Neurons are restricted to increasing weight on only those connections that overlap with their existing preferred stimulus – thus causing a smooth spreading through feature space. This is accomplished using a coefficient applied to the weight update, equal to the maximum activation amongst a neuron's strongly learned connections:

$$\kappa_i = \max_{j | (w_{ij}(t) > \frac{1}{2} \max_j w_{ij}(t))} x_j \quad (8)$$

where strongly learned connections are those whose weight exceeds half the strength of the largest weight amongst that individual neuron's connections.

2. Unlearning: Conflict learning treats inhibition as an error signal indicating that the inhibited neuron has mistakenly strengthened any currently active connections. A neuron competing with its neighbors via inhibition exerts pressure on those neurons to unlearn the connections driving its activation, while receiving reciprocal pressure to unlearn the connections that drive its neighbors. The amount of inhibition a neuron receives is used to interpolate between a positive and negative associative weight update:

$$\delta_{ij} = (1 - \text{Inhib})\eta x_i x_j \kappa_i - \text{Inhib} * \beta \eta x_i x_j \quad (9)$$

where β (set to 1 in all experiments) can be used to control the rate of learning versus unlearning. The interpolation between learning and unlearning is irrespective of activation strength and depends only upon the amount of inhibition received.

3. Short and Long-Term (SLT): Connection weights are adjusted on a short-term and long-term timescale, striking a balance between initial exploratory learning and long-term exploitation of a learned pattern. The short-term weight w_{ij}

adjusts rapidly to the current stimulus, but decays towards and fluctuates around the more stable, slowly adapting long-term weight w_{ij}^{ltm} . The only visible weight for a neuron is its short-term weight; long-term weights are internal and only observed via their effect on short-term weights. The entire neuron weight update process has four steps:

- (a) Compute short-term weight updates δ_{ij} .
- (b) Move long-term weights towards short-term weights:

$$w_{ij}^{\text{ltm}}(t+1) = (1 - s_{\text{ltm}})(w_{ij}(t) + \delta_{ij}) + s_{\text{ltm}}w_{ij}^{\text{ltm}}(t). \quad (10)$$

- (c) Move short-term weights towards long-term weights:

$$w_{ij}(t+1) = (1 - s_{\text{stm}})(w_{ij}(t) + \delta_{ij}) + s_{\text{stm}}w_{ij}^{\text{ltm}}(t+1). \quad (11)$$

- (d) Normalize short- and long-term weights independently.

where s_{ltm} and s_{stm} are smoothing factors, described in [Appendix B.2.2](#).

Conflict learning uses the same neuron activation principles as GCAL ([Stevens et al., 2013](#)), described in [Appendix B.1](#). It should be noted that the above equations, although conceptually grounded, are not directly fit to experimental data. The intent of this formulation is to demonstrate that these concepts, when used together, provide a stable and plausible way to learn in networks with modulatory connections that could exist in some fashion in actual neurons. Although weight re-normalization (3d) is not strictly biologically plausible (see [Turrigiano and Nelson \(2004\)](#) for more viable alternatives), it ensures that weights are bounded in a computationally amenable fashion, and furthermore is used in the weight update equation for GCAL.

An in depth discussion of each component of conflict learning is provided after the experiments, in [Section 5](#), using the results to address the components' contributions towards learning and their biological plausibility.

3.1. Conflict learning and modulatory connections

We can now revisit the simple network of [Fig. 1](#) and see how conflict learning resolves the observed stability problems of the Hebbian learning rules. Recall the earlier argument ([Section 2.1](#)), which showed that the analyzed Hebbian learning rules are not stable in the desired state of the network. Specifically, we noted that $w_{M_1N_2}$ had a non-zero update. This is not the case when the conflict learning rule is used instead.

Assuming the same weight configuration as for the Hebbian rule, if M_1 and N_1 are active, $x_{N_1} > x_{N_2}$, and thus $\text{Inhib}_{N_1} = 0$ and $\text{Inhib}_{N_2} = 1$. Additionally, because N_1 has an active strongly learned connection, $\kappa_{N_1} = 1$ while N_2 has no strongly learned active connections, so $\kappa_{N_2} = 0$. For simplicity we use 1 and 0 for the values of κ and Inhib , though the sign of the update remains the same so long as $(1 - \text{Inhib}_{N_2})\kappa_{N_2} > \text{Inhib}_{N_2}\beta$ holds. Substituting all of this into the short-term weight update ([9](#)) gives:

$$\begin{aligned} \delta_{M_1N_2} &= (1 - (1))\eta x_{M_1} x_{N_2}(0) - (1) * \beta \eta x_{M_1} x_{N_2} \\ &= -\beta \eta x_{M_1} x_{N_2} < 0 \end{aligned} \quad (12)$$

Since $w_{M_1N_2}$ already has a value of w_{min} , the effective negative weight update applied will be 0, much like the effective positive weight update for $w_{M_1N_2}$ will be 0 because it is already at w_{max} . Although N_2 is still partially active, it is being inhibited by N_1 , so it performs explicit unlearning towards M_1 instead of positive learning like in the Hebbian case. This same procedure can be

applied to the other three feedback connections in this example, and in each case the weight update will be 0 or restricted to 0 by the weight value range. Since all of the connection weights maintain their values, the system is at equilibrium and can maintain this steady state. Knowing that conflict learning is stable in the desired state, we can consider its behavior in the other possible states of the network and how the system transitions from an initial unlearned state to the desired stable state.

The network has five functionally distinct states of interest as seen in [Fig. 2](#). (1) The initial state, where no connections have become strongly learned (OSL). (2) A strongly learned connection between one competitive neuron and one modulatory neuron (1SL). Two strongly learned connections, either (3) one competitive neuron with a strongly learned connection to both modulatory inputs (2SL-Split), (4) one modulatory neuron with strongly learned connections between both competitive neurons (2SL-Shared), or (5) unique strongly learned connections between modulatory and competitive neurons (2SL-Desired).

We performed 30 repeated simulations of this simple network to illustrate the trajectory taken by both the considered Hebbian learning rule and conflict learning through the state space (see [Appendix C.1](#) for experimental procedures). [Fig. 2](#) shows the outgoing transition probabilities as well as the percentage of time spent in each state for both learning rules. This demonstrates that the Hebbian learning rule, which cannot prevent both competitive neurons from performing learning, immediately transitions into the 2SL-Shared state before entering an oscillation between 2SL-Shared, 2SL-Desired, and 1SL. The Hebbian rule cannot enter the 2SL-Split state because this state requires one neuron to perform learning while the other does nothing. Conflict learning, as shown in ([12](#)), is capable of performing positive learning on a competitive neuron in isolation, due to its spreading and unlearning components. The spreading component is chiefly responsible for preventing the system from entering the 2SL-Split state. The unlearning and SLT components are similarly responsible for transitioning the network out of the 2SL-Shared state, were it ever to be in that state.

A case by case analysis of the transitions made or avoided by conflict learning can be found in [Appendix A.1](#). Using the nomenclature for the states introduced here, additional analysis for simulating two additional variations of Hebbian learning, the Generalized Hebbian and BCM learning rules, is provided in [Appendices A.2](#) and [A.3](#), respectively.

4. Network modeling results

In contrast to the simple network with two competitive neurons, we now focus on large scale (several thousand neurons) neural networks. We test conflict learning by learning a model of border ownership as well as a model of orientation selectivity. The border ownership network relies on modulatory feedback for proper operation, whereas the orientation selective network demonstrates that conflict learning is a general learning rule also applicable in contexts lacking modulatory connections.

Conflict learning is compared against an implementation of GCAL ([Bednar \(2012\)](#); [Stevens et al. \(2013\)](#); threshold adjustment is implemented differently, see [Appendix B.2.1](#) for full implementation details), a learning rule that uses purely Hebbian logic to adjust its weights, increasing them when pre and postsynaptic neurons are simultaneously active. Throughout the rest of this work, we will often refer to GCAL as the ‘‘Hebbian learning rule’’ to emphasize the associative nature of its weight update. GCAL is able to achieve biologically plausible results in applications such as learning V1-like orientation selective maps by way of adjusting neuron activation through contrast normalization and adaptive thresholds ([Stevens et al., 2013](#)). For all experiments, both rules use identical activation functions, activation thresholds, and connection patterns, only differing in how their weights are adjusted.

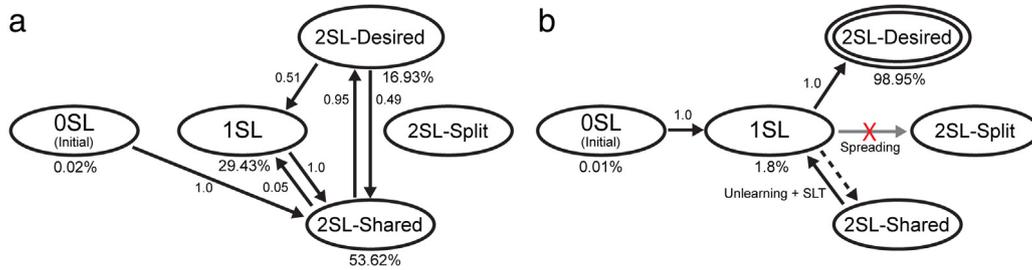


Fig. 2. State diagram for the simple network of Fig. 1. This diagram shows the progression of the network from an initial unlearned state (0SL) to the desired state of each competing neuron learning a unique modulatory input (2SL-Desired). Outgoing transition probabilities as well as the percentage of time spent in each state are shown for both (a) Hebbian learning and (b) conflict learning, based on simulation. Conflict learning enters and remains in the 2SL-Desired state, having no outgoing transitions from 2SL-Desired. By contrast, Hebbian learning oscillates between 2SL-Desired, 1SL, and 2SL-Shared. The components of conflict learning essential for specific transitions are labeled. The spreading component prevents the network from transitioning from the 1SL to the 2SL-Split state. Although the simple network under conflict learning cannot make the transition from 1SL to the 2SL-Shared state (dashed arrow), this transition is possible in general, and made unstable by the unlearning and SLT components.

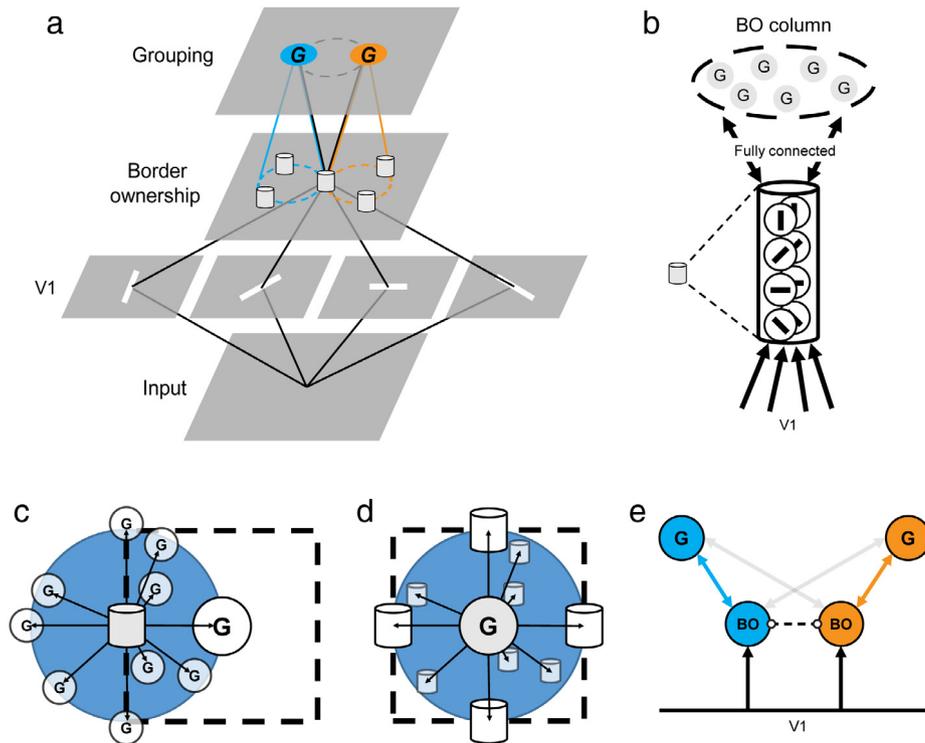


Fig. 3. Border ownership model architecture. (a) Diagram of full architecture. A V1-like layer consisting of Gabor filters processes the input at four orientations (0, 45, 90, and 135°). Each orientation neuron provides input to two border ownership cells, which are connected laterally to six others (for the three remaining orientations) at the same retinotopic location within a column in the Border Ownership layer. The grouping layer pools BO column activation, receiving input from all columns in a local receptive field. The grouping layer additionally sends feedback to those same cells. (b) Diagram of a single BO column. Column contains eight competing neurons, two for each orientation, and internally have lateral inhibitory connections between each neuron. They also receive feedback from a local receptive field in the grouping layer. (c and d) The effects of an example stimulus (dotted square, actual experiment uses solid input) on BO columns (cylinders) and grouping cells (circles labeled G). (c) Feedforward connections from the perspective of a BO column. The column sends feedforward input to all grouping cells in its receptive field, but only the grouping cell receiving input from multiple columns is highly active (indicated by increased size). (d) Feedback connections from the perspective of a grouping cell. Feedback is sent to all BO columns within its receptive field, but only those along the boundary of the object will be highly active. (e) Detailed relationship of competition between two BO neurons with the same orientation. Each BO neuron eventually learns to project to and receive feedback from a grouping cell on only one side of its orientation.

This section focuses on reporting the results of the experiments; full technical details on the experimental procedures is provided in [Appendix C](#). Intuition and further analysis of how each component of conflict learning gives rise to the results shown is provided after the results in [Section 5](#).

4.1. Border ownership

The primary benefit of conflict learning is its ability to learn in networks with modulatory feedback, a feature that allows it to be used to learn a model of border ownership. As border ownership (BO) is a less familiar and more complicated process than

orientation selectivity, it is worth briefly revisiting its putative architecture (illustrated in [Fig. 3](#), also see the experimental methods in [Appendix C.2](#)) to fully appreciate the results. The model we develop is a derivative of the feedback model of [Craft et al. \(2007\)](#), which as mentioned in the introduction, is one of multiple models capturing the observed behavior of actual border ownership neurons.

BO neurons are identified not just by an orientation, but also by a polarity, which indicates to which side of their orientation the figure (or background) lies ([Zhou et al., 2000](#)). The key challenge is to develop receptive fields such that each BO neuron responds to a single orientation with a single polarity, with full coverage over

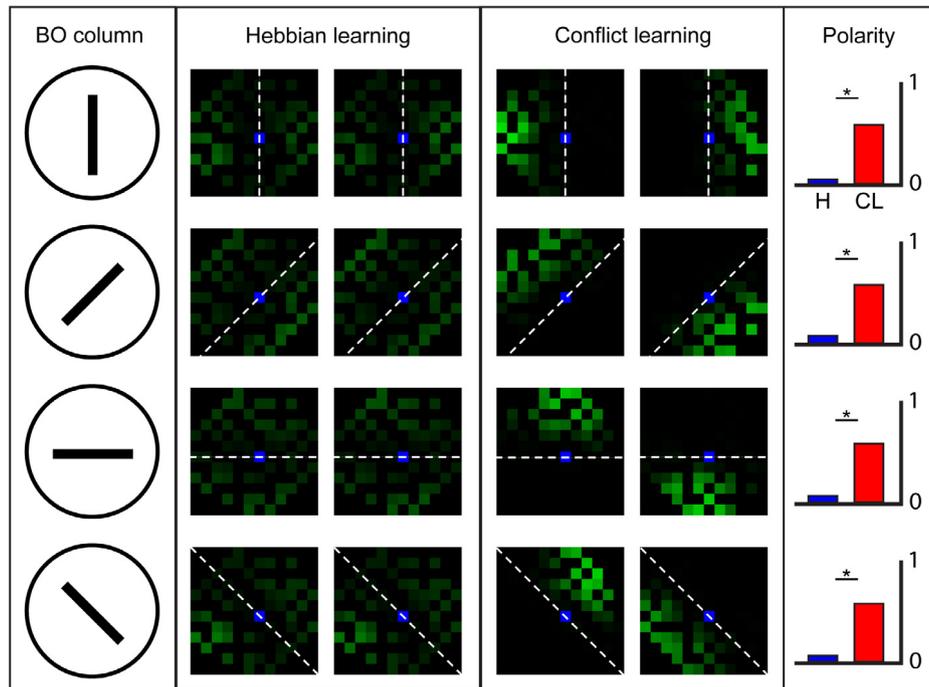


Fig. 4. Learned Feedback Receptive Fields for BO Neurons. Receptive fields are shown for all eight neurons of a single representative BO column for both Hebbian and conflict learning rules. Each row of the figure represents a different orientation. Each BO neuron is marked by a blue pixel, and green pixels show feedback connections from grouping neurons, with brightness corresponding to weight strength. Polarity represents the average degree to which every pair of neurons in a network learns feedback from grouping neurons on opposite sides, a necessary requirement for consistent border ownership assignment. The conflict learning network successfully learn pairs of competing polarity BO neurons without any a priori information regarding BO or grouping cells spatial position. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

all orientations and polarities. In our model, this relies on learning feedforward and modulatory feedback connections between columns of BO neurons and a layer of so-called grouping neurons which pool over multiple BO columns, integrating non-local information. Learning these connections is especially challenging because the multiple BO neurons that exist for each orientation, destined to develop a specific polarity, must learn consistent and opposite connection patterns. The network accomplishes this task purely through experience, with no a priori spatial information – not only are feedforward and feedback weights initially uniform, but BO neurons within a column must also learn to specialize their inhibitory lateral connections, a necessary requirement for competition. While many other models of border ownership require explicit features for junction (e.g., L, T) detection, our learned model requires only edge information.

Note that not all components of this model have been directly observed in the brain. Although BO neurons and their responses to various stimuli have been recorded (Zhou et al., 2000), grouping neurons have yet to be explicitly discovered (Craft et al., 2007). Grouping neurons can thus be seen as a computational generalization of a more complicated grouping process, for which there is mounting evidence (e.g., Martin and von der Heydt, 2015; Wagatsuma, von der Heydt and Niebur, 2016). This model is nonetheless a good approximation of the current understanding of border ownership circuits. Additionally, the structure of the border ownership network fits within a standard model of computation in visual cortex: it consists of competition followed by grouping, with increasing receptive field size. This is reminiscent of alternating simple and complex cells (Wiesel et al., 1963), which have formed the basis of many models of visual cortex (e.g., Fukushima, 1980; Serre, Wolf and Poggio, 2005). The connection from edge responsive neurons (input in the model) to border ownership neurons is a simplification (input in the model); we imagine a more realistic circuit would have edge or contour responsive neurons directly compete with each other over border ownership polarity.

4.1.1. Results

The learned feedback receptive fields for a representative BO column taken from fully trained networks are shown in Fig. 4, and the feedforward and lateral receptive fields are shown in Fig. 5 (the full details of training and other experimental procedures can be found in Appendix C.2). Under conflict learning, each neuron within a BO column learns to associate with grouping feedback occurring on only one side of its orientation, with all orientations and polarities represented. Additionally, the two BO neurons associated with each orientation learn to become competitive with each other and learn opposite sides of feedback. This occurs because the opposite sides of grouping feedback come from distinct populations of grouping neurons, and conflict learning, as was shown in Section 3.1, strives to associate one competing neuron to each population of modulatory input. The Hebbian learning based rule, however, is unable to develop this partitioning of modulatory feedback amongst competing neurons. The two BO neurons for each orientation learn the same receptive fields as each other, causing them to be unable to reliably associate with objects occurring on a particular side of their orientation. When a stimulus is presented to these neurons, the winner will be chosen randomly instead of being chosen based on any border ownership information.

Along with the sampled receptive fields, the average polarity score for BO neurons of each orientation is shown. This score represents the degree to which a competing pair of BO neurons learns feedback on opposite sides (see Appendix C.2). These averaged scores, computed from all pairs of BO neurons, demonstrate that the pictured examples are representative of the whole network.

Fig. 6 shows the results of running the trained conflict learning network on common stimuli from the border ownership literature. As the network was trained on single presentations of squares (see Appendix C.2), every shape presented here is one to which the network has never been exposed. The network in its current

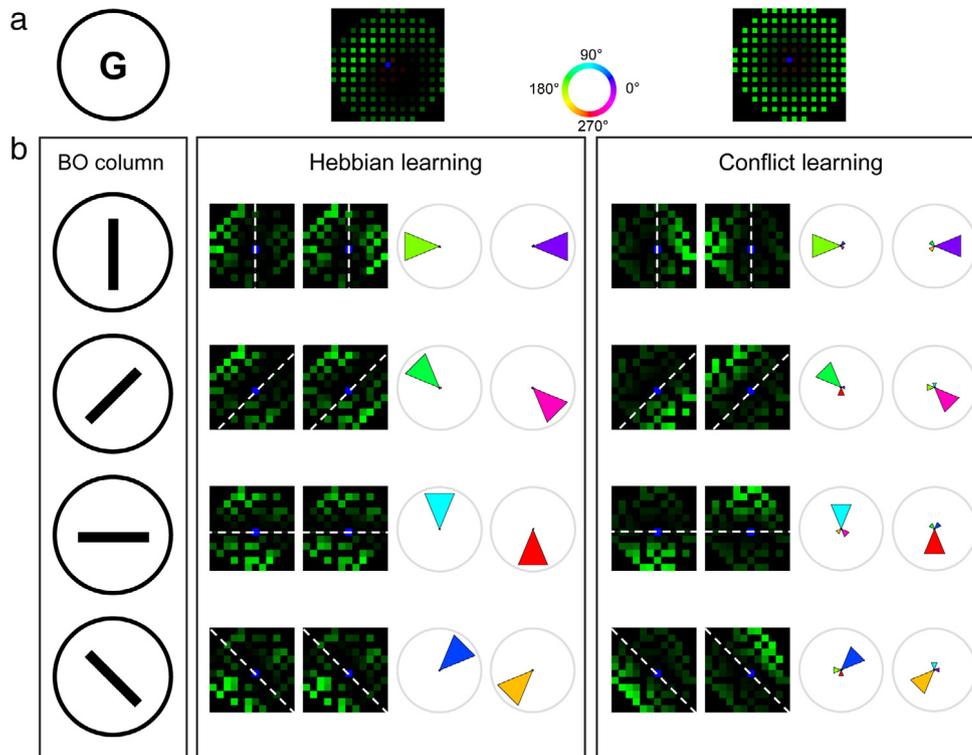


Fig. 5. Learned feedforward and lateral receptive fields for BO and grouping neurons. (a) Feedforward receptive fields for a grouping neuron, shown for both learning rules. Successful learning entails a ring-like pattern of strong connectivity. (b) As in Fig. 5, the results are organized by the orientation of the BO neuron. For each orientation, the learned outgoing feedforward projections are displayed first followed by a radial graph of the corresponding learned lateral inhibition strength for the same neurons. Lateral connections project to other neurons within the BO column, colored by the preferred polarity of the inhibited neuron. For example, a red polarity corresponds to inhibition towards a horizontal selective BO neuron with a preference for objects in the lower half of its receptive field. Under conflict learning, BO neurons learn to primarily inhibit the other neuron sharing their orientation, as well as applying a small amount of inhibition to immediately adjacent orientations with overlapping polarities. This pattern of inhibition not only ensures the creation of competing pairs of BO neurons, but also a winner-take-all like behavior amongst all orientations in a column. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

implementation has limited scale invariance, demonstrated by the weak response at the vertices of the triangle input (Fig. 6e). The responses to the tiled squares (Fig. 6a), the C pattern (Fig. 6d), and the rounded squares (Fig. 6c) are especially interesting because local information may favor a globally incorrect polarity assignment. The network, in all cases, is able to use feedback to correct ambiguous feedforward input in order to reach the correct assignment of border ownership. To our knowledge, this is the first time a border ownership network has been learned, enabled by the new conflict learning rule.

Finally, we investigate the contribution of each component of conflict learning as it applies to learning the modulatory feedback connections in the border ownership network. Fig. 7 shows receptive fields taken from a vertically oriented BO pair for all variations of rules tested. The receptive fields were chosen to be exemplars of common failures (if they existed) for the various configurations. Histograms of polarity scores over all vertical BO neurons show typical network-level results. In Fig. 7i we compare the median score across all configurations, showing that conflict learning receives benefit from the amalgamation of all of its components. The results demonstrate that there is a non-linear relationship between the introduction of a rule component and its effect on the polarity score. However, we can still extract some general conclusions with respect to the polarity score: while unlearning on its own is very influential (C), the unlearning and spreading components complement each other and together (G) account for most of the improvement over Hebbian (A). The SLT component, by slowly transitioning the network to reflect long-term statistics, appears to have the effect of eliminating outliers and reducing the variance of the distributions (e.g., histograms B versus F, C versus E,

and G versus H). Additional discussion on the contribution of each component follows in the discussion (Section 5).

4.2. Orientation selectivity

We next apply conflict learning to a problem that can be seen as a baseline for self-organizing networks of the brain – orientation selectivity. The network, seen in Fig. 8, consists of an input layer, a center-surround layer, and an output layer, like that used to demonstrate the properties of GCAL (Stevens et al., 2013). The connections between the input layer and the center-surround layer are fixed; all learning in this network takes place between the center-surround neurons and the output neurons. The network has no modulatory connections, such that the activation equation for neurons reduces to (1). The desired goal of learning in this network is to develop output neurons which are orientation selective over all possible input orientations. Detailed information on the network architecture, training, and experimental procedures are provided in Appendix C.3.

4.2.1. Results

The primary goal of this experiment is to demonstrate that the conflict learning rule, even when applied to networks lacking modulatory feedback and compared against a learning rule tailored for such an environment, produces similar biologically consistent output.

Fig. 9a shows the output neurons, for both learning rules, colored by orientation selectivity after training on oriented bar input. The learned maps show an arrangement that mimics physiological maps of orientation selectivity in mammalian cortex

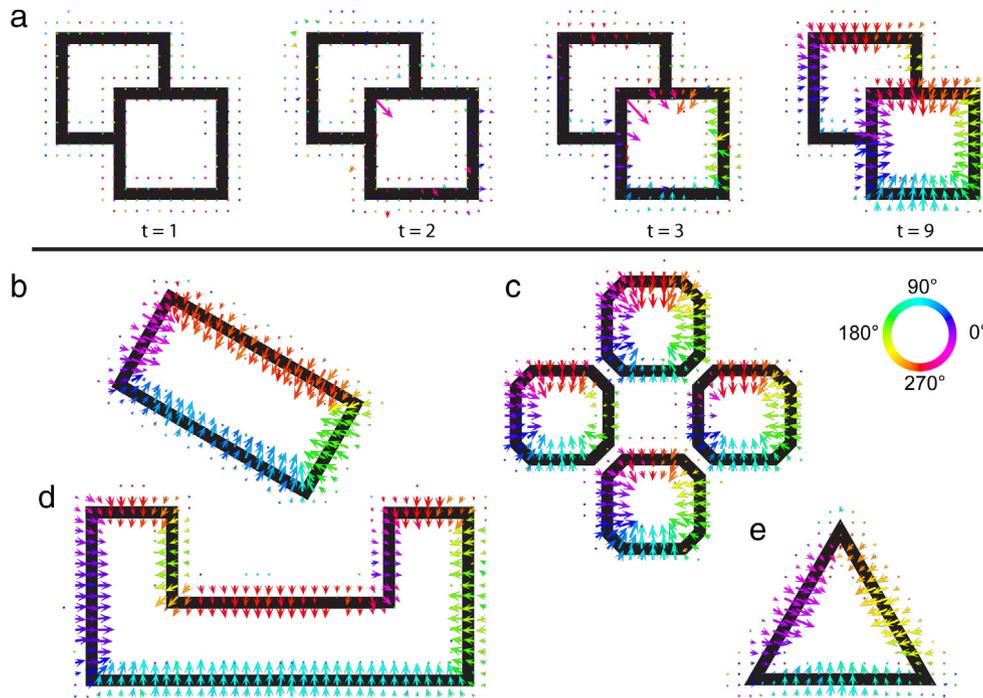


Fig. 6. Border Ownership Assignments by a Network Trained with Conflict Learning. Black lines represent the stimulus and colored arrows represent BO assignments at those locations. Each BO neuron is assigned a direction vector based on its learned polarity. Assignments are made by summing these direction vectors, weighted by activation. All results are taken from a fully learned network naive to these example inputs. The network has complete position and orientation invariance. (a) The progression of BO assignment over time. Feedback begins to arrive in iteration 3. (b–e) Settled (iteration 9) assignments for various stimuli. (c and d) These shapes have locally ambiguous border ownership assignments that are resolved through modulatory feedback from the grouping neurons. (e) The network is not fully scale invariant because the BO to grouping neuron connections exist only at a single radius, resulting in the corners being weakly activated. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

(e.g., pinwheels, which are singularities where orientation preference increases clockwise or counterclockwise; see Chapman, Stryker and Bonhoeffer (1996)). To quantify this subjective similarity, the pinwheel density metric of Stevens et al. (2013) is computed for the maps. A pinwheel density of π pinwheels per unit hypercolumn area (see Appendix C.3) has been found to be consistent across a number of mammalian species (e.g., tree shrew, galago, cat, and ferret, see: Kaschube, Schnabel, Löwel, Coppola, White et al., 2010; Keil et al., 2012), and may be a fundamental constant of map organization (Stevens et al., 2013). Both learning rules result in pinwheel densities within 3% of π .

In testing conflict learning, we also observed noteworthy behavior when varying amounts of noise were injected into the input of the system. Fig. 9 also shows the results of simulating the orientation selective network for both learning rules under varying amounts of Gaussian noise applied to the input neurons (by adjusting their activation noise term ϵ ; see Appendix C.3 for details). Fig. 9b shows an increased resistance to the effects of noise in the conflict learning results. Hebbian learning more quickly succumbs to a significant drop in the quality of learned receptive fields compared to conflict learning, which only begins to be affected by noise at very high standard deviations. The scoring metric for selectivity is based on how well a receptive field can be represented by any possible Gabor function for all neurons in the network (Olshausen & Field, 1997). Real neurons are subject to many more sources of noise and variability than is present in our modeling, and handling that noise is a fundamental requirement for the nervous system (Faisal, Selen & Wolpert, 2008). We discuss reasons why conflict learning is less affected by noise in the discussion section.

Using the same stability metric as Stevens et al. (2013), we compare how similar learned receptive fields are at any given time to the final state of the network (Fig. 9c). Conflict learning reaches a higher plateau of stability at earlier iterations compared to Hebbian learning. As stability may be important for the development of

downstream brain regions (Stevens et al., 2013), earlier stability could decrease the delay between a reliable orientation selective representation and further visual processing. Additional experiments comparing stability across a greater number of iterations did not show any appreciable difference in the time it took to reach stability or the final values. When looking at stability over increasing levels of noise, we again see a resistance to noise in conflict learning that only gives way at high standard deviations.

5. Discussion

Typically a learning rule is devised with a specific activation function in mind, so it may not seem surprising that the Hebbian learning rule we compare against was unable to learn a model of border ownership dependent on modulatory connections. However, the orientation selective network, in which there is no source of modulatory input, served as a comparison of the two learning rules in a setting where the activation function was as expected by the Hebbian rule, yet was still compatible with conflict learning, which was designed around the presence of modulatory input. In Section 2.1, Appendices A.2 and A.3, we demonstrated that unlike conflict learning, for even a minimal network with modulatory connections, neither the normalized Hebbian rule, the Generalized Hebbian Algorithm, nor BCM is capable of stably learning modulatory weights.

We suggest that this is because all of these variants of Hebbian learning are based on a core principle of associative learning, which alone seems incompatible with modulatory input. Our computational experiments suggest that a synapse does not have enough information as to how a weight should be adjusted using only incoming activation compared with the output activation of the cell. Even learning rules like BCM, which control plasticity via an adaptive threshold based on expected activation, do not solve the

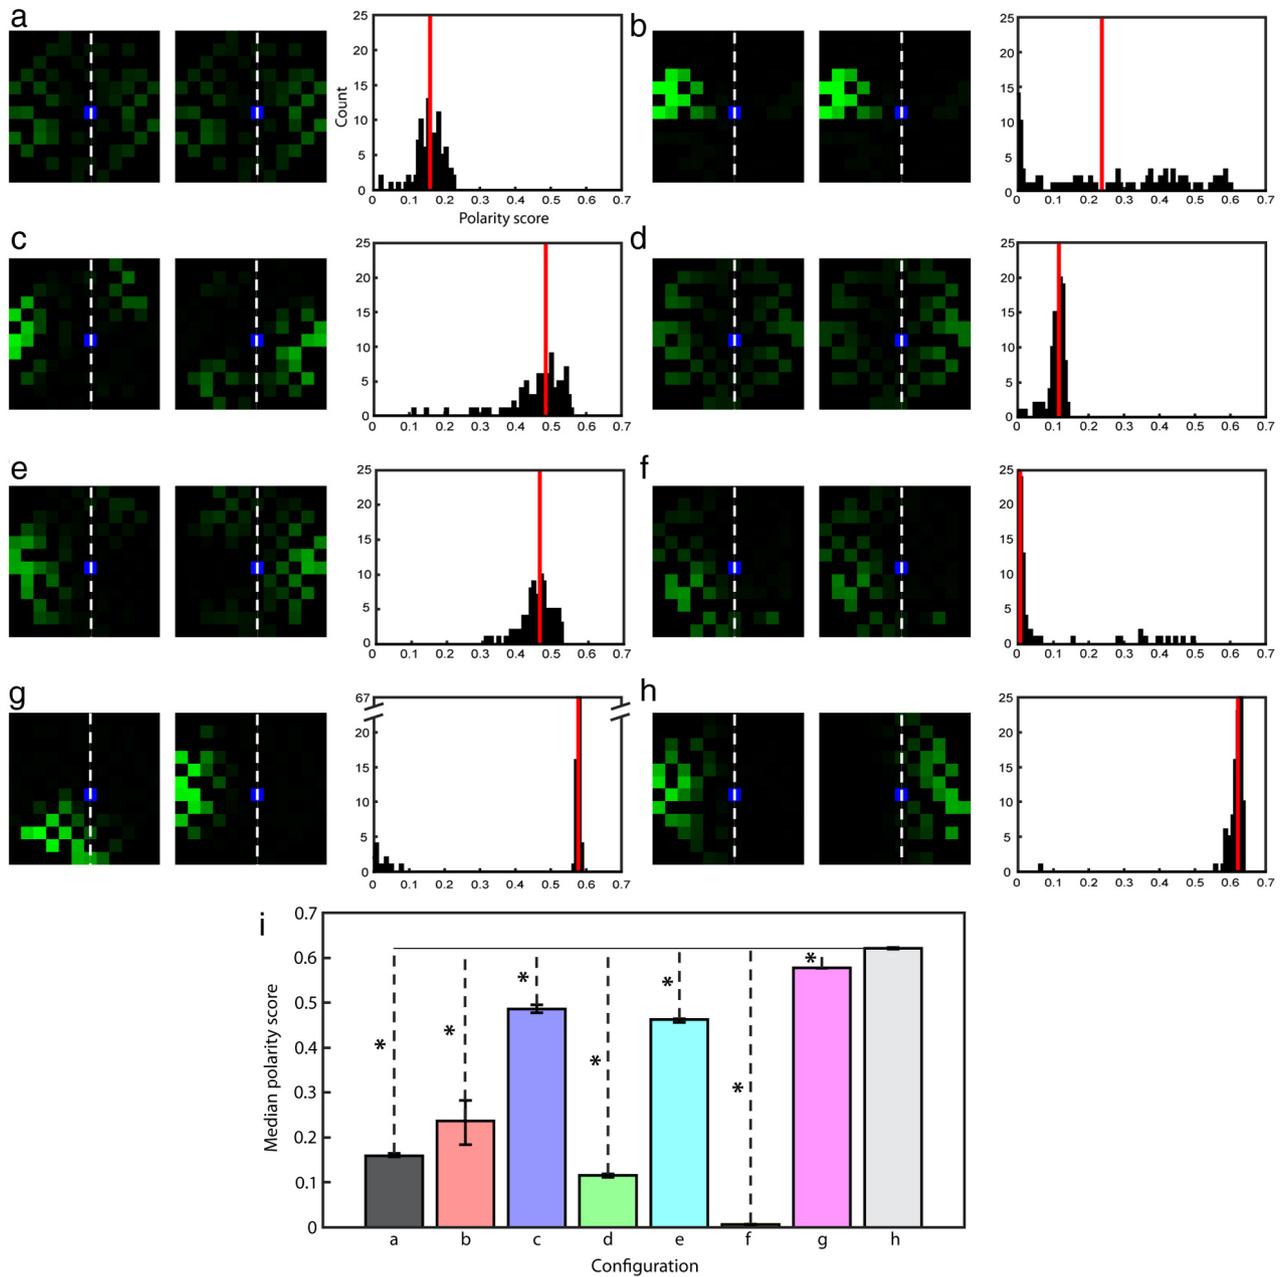


Fig. 7. Contribution of rule components. (a–h) Representative receptive fields from a vertical BO neuron pair taken from various configurations of conflict learning as well as Hebbian learning. Histograms depict the polarity scores of all vertical BO neurons for a given configuration, with the median denoted by a red line. Configurations are: (a) Hebbian, (b) spreading component only, (c) unlearning component only, (d) SLT component only, (e) conflict learning without spreading, (f) conflict learning without unlearning, (g) conflict learning without SLT, (h) full conflict learning. (i) Median scores for (a–h) with error bars indicating 95th percentile cutoffs. Conflict learning (h) is significantly higher with respect to all other configurations (a–g). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

problem, because they do not draw on any additional sources of information. We hypothesize that additional control signals are required to support modulatory connections, where the incoming activation may be coincident with the firing of the cell, but not relevant. Conflict learning uses two additional sources of information for these signals: the activation of strongly learned synapses within the cell, and inhibitory input driven by competing neurons. Strongly learned connections identify relevant firing, while inhibition partitions firing by indicating that a neuron is losing a local competition amongst connected neurons.

We demonstrated through computational models that using inhibition as a control signal results in a partitioning of correlated firing in modulatory input amongst competing neurons. Our results (e.g., Fig. 7) suggest that lowering activation through inhibition

is insufficient to prevent unwanted learning from taking place – inhibition must actively drive the partitioning of modulatory input through unlearning.

Additionally, we also demonstrated that restricting learning based on the activation of strongly learned connections results in a successful clustering of correlated firing amongst modulatory input to an individual neuron. This behavior is complementary to the partitioning performed by the inhibitory control signal, resulting in neurons which compete over correlated firing of incoming connections, regardless of whether they are sourced from driving or modulatory input.

These two components of conflict learning, together with short- and long-term learning, will be discussed and related to experiment in detail in the following section of the discussion.

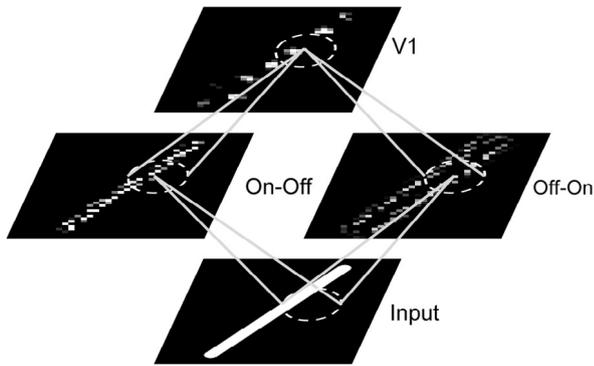


Fig. 8. V1-like feedforward network. Center-surround layers perform a difference-of-Gaussian like computation either preferring the center (On-Off) or the surround (Off-On). The orientation selective layer receives input from both On-Off and Off-On neurons and forms lateral connections within some radius. The connections between center-surround layers and the orientation selective layer are learned. Figure depicts actual model responses from a learned network.

Although it may be the case that a different learning rule could govern driving versus modulatory connections, we think there is some elegance in a single set of principles being compatible with both types of excitatory connections. Conflict learning does not directly address the plasticity of inhibitory connections, which likely do operate with a different set of mechanisms. In fact, conflict learning cannot be used for learning inhibitory connections because of its reliance on inhibition as a control signal (see Appendix B.2.2).

5.1. Analyzing the rule

The results demonstrate that for certain patterns of connections and firing, traditional Hebbian learning mechanisms are ill-suited for adapting synaptic weights. This was seen directly in learning a model of border ownership, where only conflict learning was able to properly learn the required modulatory feedback connections to perform the computation correctly. Additionally, conflict learning operates in a biologically consistent manner even in situations lacking these types of connections, with the pinwheel density of the learned orientation selective network matching biology as well as other learning rules. The orientation selective network experiments also show interesting properties with regard to increased stability and robustness to noise. All of these results are a product of the three complementary components that make up conflict learning, introduced in Section 3, which we now discuss in detail.

5.1.1. First component – spreading

The first component of conflict learning states that neurons cannot strengthen their connection weights unless an already strongly learned connection is currently active. In the border ownership experiments, the spreading component helps prevent neurons of a border ownership pair from associating with grouping neurons on both sides of their oriented edge. While populations of grouping neurons on both sides are individually co-active with a BO neuron, there is little to no correlation between the firing of the distant populations themselves. A Hebbian neuron cannot detect this distinction, whereas a conflict learning neuron can. This is illustrated most clearly in the learned receptive fields of the border ownership experiment, seen in Fig. 4, as well as by the simple network of Section 3.1.

Spreading is similar to the concept of associative LTP (long-term potentiation), where the strong firing of a learned synapse supports the strengthening of a weaker one (Linden & Connor,

1995; Shouval, Samuel & Wittenberg, 2010). There has been discussion on the spatial requirements (Engert & Bonhoeffer, 1997) as well as temporal constraints (Levy & Steward, 1983) of synapses involved in associative LTP, suggesting that it is both a spatially and a temporally local process. Since we do not model the physics of our synapses, we use only a temporal constraint. This means that once a neuron has begun to associate with certain connections, any further connections it strengthens must co-occur with the existing ones, which forces connection weights to smoothly spread outward through feature space from an initially learned pattern. In situations where initial conditions allow competing neurons to learn the same set of connections (analogous to the 2SL-Shared state described in Section 3.1), the spreading component, if used without the unlearning component, would make it impossible for the neurons to disentangle their learned features. In Fig. 7b, the two BO neurons are correctly learning on only one side of the boundary, but have no mechanism to prevent them from learning and spreading to the same features. This effect is exacerbated when combined solely with the long-term statistics used by the SLT component (Fig. 7f), which compounds the mistaken initialization over time.

Our method of labeling connections within a single neuron as strongly learned (8) is a simple abstraction intended to capture the behavior, but not the exact biological implementation, of such a mechanism. It has been demonstrated that the soma can back-propagate signals to its dendrites for the purpose of manipulating thresholds (e.g., Larkum, 2013) and that individual dendrites display a wide array of active properties such that one synapse can affect the behavior of many others (Major, Larkum & Schiller, 2013). Such mechanisms could also be responsible for the manipulation of a learning threshold affecting synaptic plasticity. Therefore the spreading component, in a real neuron, would likely be implemented through a variety of adaptable thresholds as opposed to the simple activation strength based product that we employ.

In the context of modulatory connections, the spreading component is essential to enabling a neuron to identify a population hidden within the many correlated activations of its inputs. In networks without modulatory feedback, the spreading component gives increased resistance to the effects of noise (Fig. 9) by lessening the impact of spurious activation as it is unlikely to consistently coincide with the strongly learned connections.

5.1.2. Second component – unlearning

In conflict learning, inhibition, in addition to reducing the activation of a neuron, causes the neuron to directly unlearn its active connections. This is in contrast to a typical Hebbian learning rule which still allows positive learning to take place, dependent on the activation. It is also distinct from examples of explicit synaptic weakening in BCM-like rules or STDP, which use activation or timing to control the unlearning. In conflict learning, a neuron can be strongly active but still undergo unlearning if its inhibitory input is high enough. In the border ownership experiments, inhibition primarily occurs between pairs of border ownership cells competing over feedback from grouping cells on either side of their local oriented boundary. Unlearning helps correct mistaken assignments within a BO pair, ultimately resulting in a near even split along the polarity boundary (Fig. 4). Mistaken activation close to the boundary will be frequently contested and thus unlearned by both cells in the pair.

There is significant evidence of complex interactions between inhibition and excitation in the brain. Wang and Maffei (2014) found that inhibition controlled the sign of excitatory plasticity in rat visual cortex, which is remarkably similar to our unlearning component, via crosstalk between inhibitory and excitatory signaling. Fino, Paille, Cui, Morera-Herreras, Deniau et al. (2010) found that the presence or lack of inhibition could reverse the classic

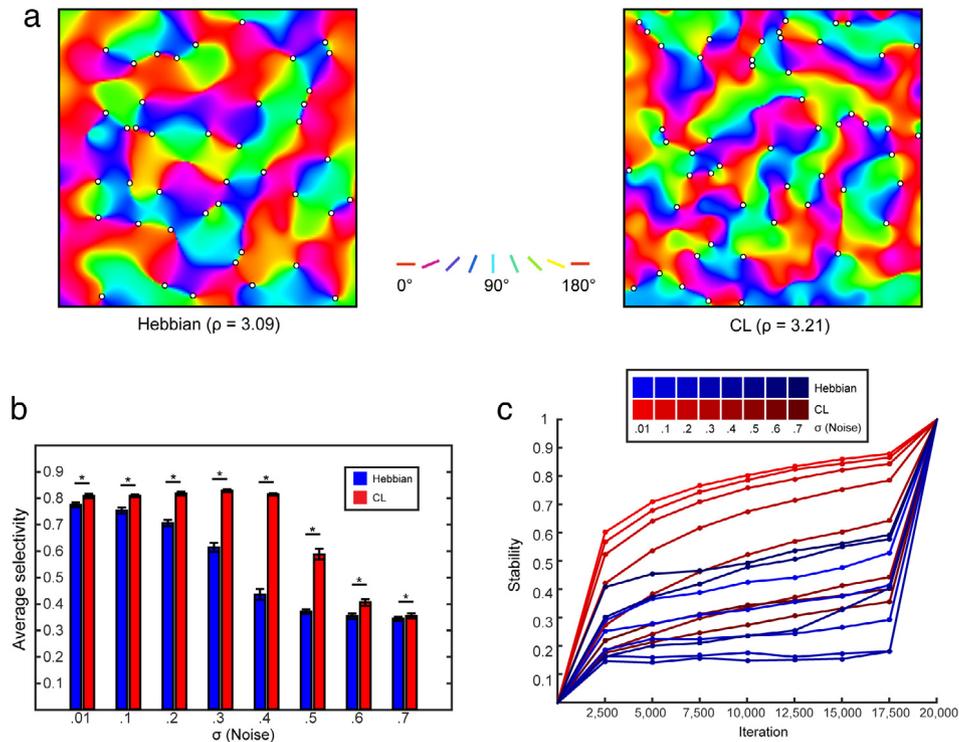


Fig. 9. Orientation selectivity results. (a) Orientation maps for both learning rules, colored according to the preferred orientation of each neuron. Pinwheel locations are determined algorithmically and denoted by white circles. Both learning rules result in a biologically plausible pinwheel density within 3% of π . (b) Average selectivity for both learning rules while training with input data corrupted with increasing amounts of Gaussian noise. Selectivity is based on how well a neuron's receptive field can be modeled by any Gabor function. (c) Stability for both learning rules as a function of learning iteration for a range of input noise values. Stability is based on the correlation between the current and final (iteration 20,000) maps. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

STDP window, causing either LTP or LTD (long-term depression) to occur. Additionally, in a recent review on inhibitory plasticity, Vogels et al. (2013) emphasize the increasing evidence that excitation and inhibition are deeply intertwined, with inhibition potentially providing a mechanism that allows selective learning to occur.

Unlearning through inhibition allows one neuron to force another to unlearn common connections between the two, causing the inhibited neuron to return to an initial unlearned state, at which point it is possible to learn a different population of input. This was first seen in analysis of the simple network (Appendix A.1), where the unlearning component is the primary mechanism by which a 2SL-Shared state is made unstable. A consequence of this component is that when a neuron competes for features it actively weakens competitors, leading to a greater separation in feature space (weight values) between the neurons (Fig. 7c). When combined with the spreading component in competitive groups of neurons (such as the mutually inhibitory groups of neurons in a column), neurons learn in a smooth yet competitive fashion (Figs. 7g and 7h). The neurons identify populations in the feature space and slowly expand their receptive fields until they have no more correlated connections to learn or they are faced with competition from another neuron. In the orientation selectivity experiment, unlearning enforces a greater difference in connection weight strength between the features learned by each neuron, meaning responses are more stable and higher levels of noise can be introduced without confusing the input pattern.

5.1.3. Third component – short and long term

In conflict learning, all neurons have an externally visible short-term weight as well as an internal long-term weight. The two weights constantly pull on each other until they settle to the same value, with the rates at which they move towards each other controlling how quickly a neuron adapts its weights and how steadfast

it becomes in its decisions. This short- and long-term learning, or SLT, allows neurons to quickly associate with populations in their input while remaining sensitive to long-term trends. In the border ownership network, this ability to be initially flexible but stable in the long run leads to more neurons learning significantly better separation along BO neuron boundaries (Fig. 7i). We found SLT to be especially beneficial for feedforward connections, where capturing long-term statistics is useful (e.g., BO feedforward connections). SLT, used alone, works essentially like Hebbian learning (Fig. 7d), but when combined with the other two portions of the rule, leads to a significant improvement and consistency of learned receptive fields (Fig. 7h). This increased consistency can also be seen in Fig. 7e compared to Fig. 7c, which differs only by the inclusion of SLT.

The physiological underpinnings of multi-timescale learning are notably discussed by Zucker and Regehr (2002), who review the dynamics of short-term learning, Abbott and Nelson (2000), who review synaptic redistribution and the interplay between short- and long-term potentiation, and Grossberg (2013), throughout his extensive development of adaptive resonance theory.

5.2. Implications for plasticity

Our results, accompanied by physiological evidence for the mechanisms we have described, suggest that similar mechanisms are likely used in the brain for the learning of modulatory connections. By acting as an error signal to instigate unlearning, inhibition can dynamically alter plasticity and encourage diversification amongst competing neurons, and by requiring strongly learned connections to be active for learning, spreading allows for the detection of correlated clusters of activation within non-driving inputs. Our model of primary visual cortex shows that such mechanisms do not interfere with learning in more traditional contexts

lacking modulation. We therefore predict that neurons likely have the key mechanisms of conflict learning: the ability to adjust their plasticity based on a concept of synaptic strength, and the usage of inhibition as a control signal for unlearning. These concepts could be tested in actual neurons with a series of simple experiments on single neurons. For all of these proposed experiments, we assume that a single neuron has learned a preferred stimulus such that it has increased synaptic strength towards the inputs associated with that stimulus. Inhibition is assumed to originate through interaction with other neurons (e.g., inhibitory interneurons [Markram, Toledo-Rodriguez, Wang, Gupta, Silberberg et al., 2004](#)). Conflict learning predicts that inhibition has additional effects on plasticity if its presence lowers the activation without completely suppressing its firing.

If inhibition serves as a signal that unlearning should occur, the strength of the synapses associated with the learned input should decrease when inhibition is applied simultaneously to the driving input. As noted in Section 5.1.2, there is existing evidence that this is indeed a potential role for inhibition. A classical Hebbian theory, such as any of the rules discussed in Section 2.1, or STDP, would predict no decrease in synaptic strength in such a situation.

To establish the existence of a behavior similar to the spreading component, a new, independent source of input could be applied while artificially activating the neuron. Conflict learning predicts that the lack of activation of the already learned input will prevent or significantly impair the learning of the novel input. Existing Hebbian rules predict that synaptic strength towards the new input should increase unimpeded.

Finally, the interaction of these two components could be tested in a combination of the two experiments. While driving the neuron via its preferred stimulus and supplying a sufficient source of inhibition, additionally apply a new, independent source of input. In this situation, conflict learning predicts that the neuron will not increase synaptic strength towards the novel input, even though it is presented simultaneously to the learned input. This prediction arises from the proposed role of inhibition, which in this situation would cause all active inputs to the neuron to have their synaptic strength decreased. A classical Hebbian rule here would predict that the synaptic strength to the novel input would increase.

Furthermore, if inhibition is indeed a necessary component for learning modulatory connections, it follows that modulatory connections (and thus a majority of feedback) must develop to maturity alongside inhibition. The balance between excitation and inhibition is a drawn-out process controlled by experience ([Froemke, 2015](#)), and a potential additional reason for this delayed maturation could be explained by a dependence between inhibition and feedback.

5.3. Learning border ownership

As mentioned in 4.1, the border ownership network architecture we present is not fully drawn from physiological observations. Our results do not definitively rule out that a Hebbian based rule, under some alternative network configuration, could reproduce the behavior of border ownership. However, given the prevailing theory that the computation of border ownership is dependent upon feedback ([Kogo & van Ee, 2014](#)), along with the argument presented in Section 2 demonstrating Hebbian learning's incompatibility with modulatory connections, it seems unlikely that a Hebbian learning rule could learn a feedback-based border ownership network. Additionally, through our experiences developing conflict learning, we believe that any network configuration compatible with purely Hebbian learning would be overwhelmingly complex and likely not support stimulus driven learning.

As briefly discussed in 4.1, the network architecture used here, although only applied to border ownership, is not specifically tied to computing this one feature. The network has no a priori information about borders or specific relationships between neurons. Border ownership is instead an emergent property of the network given competition over orientation responses coupled with higher level grouping. A deepened hierarchy composed of the same type of competition and grouping may potentially lead to the computation of higher level features more akin to proto-objects (for a discussion of proto-objects, see [von der Heydt, 2015](#)), and is a target for future work.

6. Conclusion

In developing conflict learning, we have shown how existing mechanisms already found in the brain can interact together to provide substantial benefits in learning and allow the learning of modulatory connections. We have demonstrated the effectiveness of conflict learning by showing, for the first time, how a model of border ownership might be learned through experience. This new rule could additionally be beneficial for modeling many brain functions, including figure-ground segmentation, top-down attention, and object recognition, which may all benefit from top-down modulation. As we uncover more details of the mechanisms governing neural plasticity, models capable of incorporating this new information, such as conflict learning, become increasingly necessary.

Acknowledgments

This work was supported by the National Science Foundation (Grant Numbers CCF-1317433 and CNS-1545089), the Army Research Office (W911NF-12-1-0433), and the Office of Naval Research (N00014-13-1-0563). The authors affirm that the views expressed herein are solely their own, and do not represent the views of the United States government or any agency thereof.

Appendix A. Stability of modulatory connections

This section details the possible transitions that occur in the network of [Fig. 1](#) using various learning rules. The nomenclature used for the various states of the network is the same as introduced in Section 3.1.

A.1. Conflict learning transitions

Traditionally this type of stability analysis is performed by analyzing the properties of the Jacobian. The discontinuous nature of the spreading component (8) of conflict learning, which is caused by the categorization of neurons as strongly learned or not, precludes writing a single equation for the individual components of the Jacobian. Given the analyzed network, this would mean creating a distinct Jacobian for each state and categorization of neurons, which would only serve to complicate the presented analysis. We instead continue in the same fashion as in Sections 2.1 and 3.1.

Transitions out of OSL

When the network is in its initial unlearned state OSL, there is no association between modulatory input and competitive neurons, so regardless of which neuron wins or which modulatory input is active, the update occurs in the same fashion. Without loss of generality, let N_1 and M_1 be the active neurons. The updates from M_1 are then:

$$\delta_{M_1 N_1} = \eta x_{M_1} x_{N_1} > 0 \quad (\text{A.1})$$

$$\delta_{M_1 N_2} = -\beta \eta x_{M_1} x_{N_2} < 0 \quad (\text{A.2})$$

which transitions N_1 into being strongly learned towards M_1 . Connection weights from M_2 are unchanged because M_2 is inactive.

Transitions out of 1SL.

Once in a state where one of the competing neurons has a strongly learned connection, there are four possible scenarios of activation. We will again assume, without loss of generality, that N_1 has strongly learned connections from M_1 , and that N_2 has no strongly learned connections.

- N_1 and M_1 active: $\delta_{M_1N_1}$ is the only positive update, so the weight changes proceed as they did under the same conditions in the initial state, keeping the network in the 1SL state.
- N_1 and M_2 active: Because N_1 is strongly learned towards M_1 , κ_{N_1} will be 0 as M_1 is inactive. Thus none of N_1 's weights can change and the network remains in the same state, the spreading component preventing the network from transitioning into the 2SL-Split state. N_2 receives inhibition from N_1 , causing it to unlearn towards the active modulatory neuron M_2 , which results in no effective change as $w_{M_2N_2}$ is already w_{\min} .
- N_2 and M_1 active: In this simple example, the existing feedback from the strongly learned connection between N_1 and M_1 overrides the driving input to N_2 , so N_1 becomes active and N_2 inactive, which we have already seen results in no change to the state.
- N_2 and M_2 active: N_2 has no strongly learned connections, thus $\kappa_{N_2} = 1$ and N_2 can learn towards the active modulatory input M_2 . N_2 becomes strongly learned towards M_2 and the network enters the 2SL-Desired state.

In more complex networks, it is possible to transition from the 1SL state to the 2SL-Shared state. In these networks, in place of a single neuron, modulatory input comes from correlated populations of neurons. Depending on the activation of the population, a particular competitive neuron may only be able to learn a subset of connections to a population while one of its competitors learns a different subset. Alternatively, there may be overlap between populations of modulatory inputs, meaning that some of the neurons that are learned belong to both populations, resulting in a sharing of strongly learned connections.

Transitions out of 2SL-shared.

When a network is in this state, more than one competitive neuron has a strongly learned connection to the same modulatory population. In this case, the unlearning component in conjunction with the SLT component work to make this an unstable state and move the network back to 1SL: the less active competitive neuron will actively unlearn its connection to the active population, while the more active one strengthens its connection. Over time this will result in one of the neurons losing its strongly learned status to that population, allowing it to return to an initial unlearned state. The SLT component allows initial changes to happen quickly and creates momentum via long-term statistics once one neuron begins to consistently win versus the other.

Consider the behavior of the simple network of Fig. 1 if placed into the 2SL-Shared state: because N_1 and N_2 both have strongly learned connections to M_1 , the following applies identically to either N_1 or N_2 :

- If M_1 becomes active, either N_1 or N_2 will be more active, depending on noise. The winner will update its weights further towards M_1 while the loser will unlearn its weights towards M_1 . If N_1 were the winner, this differential in weight value will cause N_1 to win versus N_2 in future cases of M_1

being active, maintaining these weight updates until the system returns to the 1SL state:

$$\begin{aligned} \delta_{M_1N_1} &= (1 - (0))\eta x_{M_1} x_{N_1} (1) - (0) * \beta \eta x_{M_1} x_{N_1} \\ &= \eta x_{M_1} x_{N_1} > 0 \end{aligned} \quad (\text{A.3})$$

$$\begin{aligned} \delta_{M_1N_2} &= (1 - (1))\eta x_{M_1} x_{N_2} (1) - (1) * \beta \eta x_{M_1} x_{N_2} \\ &= -\beta \eta x_{M_1} x_{N_2} < 0. \end{aligned} \quad (\text{A.4})$$

- If M_2 becomes active, neither N_1 nor N_2 will perform positive learning because they are strongly learned towards M_1 and $\kappa_{N_1} = \kappa_{N_2} = 0$.

A.2. Generalized Hebbian algorithm

The Generalized Hebbian Algorithm (GHA) can be shown to be unstable for the network of Fig. 1 using the same procedure as was used for the normalization based Hebbian learning rule (4). GHA adjusts weights as follows:

$$\Delta w_{ij} = \eta \left(x_j x_i - x_j \sum_{k=1}^j w_{ik} x_k \right). \quad (\text{A.5})$$

Here we assume most of the same network assumptions as 2.1. This means the network is already in the desired state and $w_{M_1N_1} = w_{M_2N_2} = w_{\max}$ and $w_{M_2N_1} = w_{M_1N_2} = w_{\min}$. However, we assume M_2 and N_2 will be highly active instead of M_1 and N_1 . Now, replacing (4) with (A.5) yields:

$$\Delta w_{M_2N_1} = \eta (x_{N_1} x_{M_2} - x_{N_1} (w_{\min} x_{N_1})). \quad (\text{A.6})$$

As we are only interested in the sign of $\Delta w_{M_2N_1}$, and because $\eta > 0$ and $x_{N_1} > 0$, we have:

$$\begin{aligned} \text{sgn}(\Delta w_{M_2N_1}) &= \text{sgn}(\eta (x_{N_1} x_{M_2} - x_{N_1} (w_{\min} x_{N_1}))) \\ &= \text{sgn}(x_{M_2} - w_{\min} x_{N_1}). \end{aligned} \quad (\text{A.7})$$

Because M_2 is highly active while N_1 is being inhibited, $x_{M_2} > x_{N_1}$. Considering this along with the fact that $1 > w_{\min} \geq 0$, it must be true that $\Delta w_{M_2N_1} > 0$, indicating that the system is not in a steady-state.

Results for simulating this learning rule for the network of Fig. 1 can be seen in Fig. A.10a. The simulation confirms that 2SL-Desired is not a stable state for this learning rule and shows that the network enters an oscillation between multiple states.

A.3. BCM

We can also demonstrate that the BCM rule, another variant of Hebbian learning, is not guaranteed to converge to the desired state of the network of Fig. 1. BCM uses a Hebbian update modulated by a dynamic threshold to control explicit synaptic weakening:

$$\Delta w_{ij} = \eta x_i x_j (x_j - \theta_j) \quad (\text{A.8})$$

where θ_j is the expected value (long-term average) of x_j^2 .

The value of θ_j directly controls whether this rule is stable in the desired state of the simple network. Let us assume that the network is in the desired state 2SL-Desired, and M_1 is the current active modulatory neuron, implying $x_{N_1} = x_{\text{active}}$, $x_{N_2} = x_{\text{inhibited}}$, and $x_{\text{active}} > x_{\text{inhibited}}$. The sign of each weight update is then dependent solely on the $(x_j - \theta_j)$ term of (A.8). For the system to remain in the stable state, $\Delta w_{M_1N_1} \geq 0$ and $\Delta w_{M_1N_2} \leq 0$ must hold, as these updates maintain the same assignment of strongly learned connections. The system must simultaneously satisfy the case when M_2 is the active modulatory neuron, which sets up a similar set of requirements: $\Delta w_{M_2N_1} \leq 0$ and $\Delta w_{M_2N_2} \geq 0$.

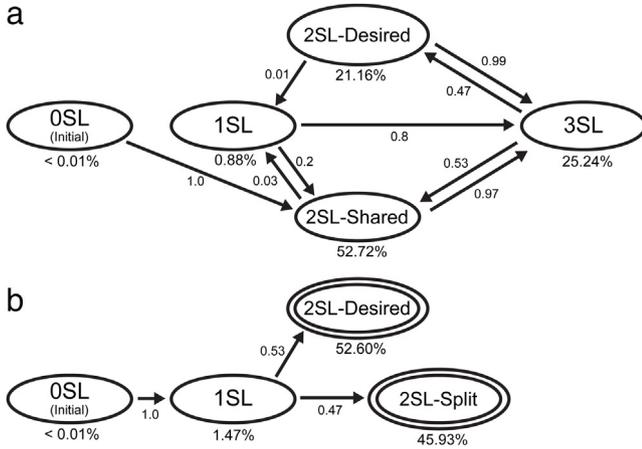


Fig. A.10. State diagram for the simple network of Fig. 1 for the Generalized Hebbian Algorithm (Sanger’s Rule) and BCM. This diagram shows the progression of the network from an initial unlearned state (0SL) to the desired state of each competing neuron learning a unique modulatory input (2SL-Desired), much like Fig. 2 did for a normalized Hebbian learning rule and conflict learning. Outgoing transition probabilities as well as the percentage of time spent in each state are shown for both (a) the Generalized Hebbian Algorithm (GHA) and (b) BCM, based on simulation. States which were not reached by a learning rule have been omitted for clarity. The 3SL state corresponds to exactly three strongly learned connections between modulatory (M_1 and M_2 in Fig. 1) and competitive neurons (N_1 and N_2), regardless of which set of three connections is strongly learned. The 3SL state was not reachable by either rule shown in Fig. 2, either due to weight normalization or various components of conflict learning. Although GHA can spend time in the 2SL-Desired state, it is not stable in that configuration and oscillates between four different states. BCM is stable in the 2SL-Desired state, but is also stable in the 2SL-Split state. Using BCM, the transition out of 1SL is essentially random, and the system cannot reliably end up in the desired state.

Arranging all requirements and substituting x_{active} and $x_{\text{inhibited}}$ where appropriate, we get:

$$\begin{aligned} x_{N_1} = x_{\text{active}} &\geq \theta_{N_1} \\ x_{N_2} = x_{\text{inhibited}} &\leq \theta_{N_2} \\ x_{N_1} = x_{\text{inhibited}} &\leq \theta_{N_1} \\ x_{N_2} = x_{\text{active}} &\geq \theta_{N_2} \end{aligned} \quad (\text{A.9})$$

which is satisfied if and only if $x_{\text{active}} > \theta > x_{\text{inhibited}}$.

However, the BCM rule has another stable state which it can reach, 2SL-Split, which is the state where both modulatory neurons are associated with a single competitive neuron. Once in the 2SL-Split state, the competitive neuron with two strongly learned connections will always activate more strongly than and inhibit the other because it is receiving additional feedback input regardless of which modulatory neuron is active.

Let us investigate the dynamics of the network in the 1SL state, before it reaches either 2SL-Split or 2SL-Desired. Without loss of generality, assume N_1 has strongly learned connections from M_1 , and that N_2 has no strongly learned connections. Consider what happens when the threshold, θ , falls within the required bounds for stability in 2SL-Desired, such that the winning neuron with activation x_{active} will do positive learning, and the inhibited neuron with activation $x_{\text{inhibited}}$ will do negative learning. The interesting case is what happens when M_2 is the active modulatory neuron, which has no existing strongly learned connections (i.e. $w_{M_2N_1} = w_{M_2N_2} = w_{\text{min}}$). N_1 and N_2 thus receive identical input, so the winner is decided by noise. Due to the value of the threshold, the winner will increase its weight towards M_2 , and the loser will decrease its weight towards M_2 . If the winner happens to be N_1 , the system will transition into 2SL-Split. If N_2 wins, the system will transition to 2SL-Desired.

This result can be seen in the simulation results presented in Fig. A.10b, where the network under the BCM rule has two terminal

states: 2SL-Desired and 2SL-Split. To achieve this, we specifically initialized the adaptive threshold to a value between the activation values of activation and inhibition for the network.

Appendix B. Learning rule details

B.1. Activation

For all experiments, all model neurons use the same activation function regardless of learning rule. A neuron j has a continuous firing rate x based on integrating weighted inputs:

$$x_j = f\left(\frac{\text{FF} + \text{Lat} + (\text{FB} * \text{FF}^2) + \epsilon}{1 + \text{Inhib}}, \theta_j\right) \quad (\text{B.1})$$

where FF, Lat, and FB represent the sum of weighted inputs of all excitatory (weight $w \geq 0$) feedforward, lateral, and feedback inputs, respectively. Each sum is calculated as: $\sum_{i \in \text{type}} w_{ij} x_i$, where w_{ij} is the weight between neurons i and j . Note that feedback is gated by feedforward input; it cannot activate a neuron in the absence of feedforward driving input.

Inhib is calculated by taking the weighted sum of all inhibitory inputs from more strongly active neurons.

ϵ is a noise term sampled from a normal distribution: $\mathcal{N}(0, \sigma_{\text{noise}}^2)$.

$f(x, \theta)$ sets the output to zero if it is less than a threshold value. Thresholds are updated whenever a neuron is active and not inhibited:

$$\theta = \begin{cases} s * \text{FF} + ((1 - s)\theta) & \text{FF} \geq \theta_{\text{ff}} \text{ and } \text{Inhib} < \theta_{\text{inhib}} \\ 0 & \text{else} \end{cases} \quad (\text{B.2})$$

where s is a smoothing parameter, θ_{ff} a threshold for considering a neuron active, and θ_{inhib} a threshold for considering a neuron inhibited.

Thresholds are further bound between a minimum (θ_{min}) and maximum (θ_{max}) value. The minimum is set such that the noise term ϵ is unlikely to spuriously activate the neuron.

B.2. Learning

In the experiments, each model neuron, under either learning rule, learns each type of connection (i.e. feedback, feedforward, and lateral) independently.

B.2.1. Hebbian learning

Our experiments use a slightly modified version of GCAL (Bednar, 2012) where the threshold works as described in Appendix B.1, instead of a global target activation based threshold such as that described in Stevens et al. (2013). This change to the threshold resulted in better performance and easier system tuning for both of our experiments. The rule is otherwise the same, using purely Hebbian logic, i.e. (4), to determine weight updates, and the activation function described above (Appendix B.1).

B.2.2. Conflict learning

Conflict learning neurons adapt their weights as described in Section 3. Neurons additionally have an accumulator of lifetime short-term weight updates which is used for computing the smoothing factor s_{ltm} for the long-term weight update:

$$\text{acc}_{ij}(t + 1) = \text{acc}_{ij}(t) + \delta_{ij}. \quad (\text{B.3})$$

The smoothing factor for the long-term update, s_{ltm} , is computed by comparing this neuron’s proportion of long-term weight against its proportion of lifetime accumulator value (normalized $w_{ij}^{\text{ltm}}(t)$ versus $\text{acc}_{ij}(t + 1)$). When the $w_{ij}^{\text{ltm}}(t)$ update would move the long-term weight proportion towards that of the accumulator, s_{ltm} is decreased, proportional to the remaining distance between

them. In cases where the w_{ij}^{ltm} update would move the proportion away from the accumulator, s_{ltm} is increased.

The smoothing factor for the short-term update, s_{stm} , is constant, with smaller values preferring the short-term weight.

Short-term and long-term weights are divisively normalized independently, as in the Hebbian update. Weights initially start lower than their allowed totals and are not normalized until they have grown to exceed it.

The full conflict learning rule is not used for learning inhibitory connections as they serve as control signals for the rule itself. Instead, these connections have a single weight based upon a normalized lifetime accumulation of weight updates:

$$\text{acc}_{ij}^{\text{lat}^-}(t+1) = \text{acc}_{ij}^{\text{lat}^-}(t) + x_i x_j w_{ij}^{\text{lat}^-} (1 - \text{Inhib}) \quad (\text{B.4})$$

$$w_{ij}^{\text{lat}^-}(t+1) = \frac{\text{acc}_{ij}^{\text{lat}^-}(t+1)}{\sum_k \text{acc}_{kj}^{\text{lat}^-}(t+1)}. \quad (\text{B.5})$$

Appendix C. Experimental methods

C.1. Simple network

The simple network of Fig. 1 is used to demonstrate the instability of variants of Hebbian learning when modulatory connections are present. The results for simulating the network for both weight re-normalization Hebbian learning and conflict learning are presented in Fig. 2. Additional simulations using the Generalized Hebbian Algorithm and the BCM rule are presented in Fig. A.10.

All connection weights are fixed to 1 with the exception of incoming modulatory input to the competitive neurons, which adjust their weights using the learning rule being tested. Competition is implemented through lateral inhibition between neurons N_1 and N_2 . All tested learning rules use the same network parameters.

The results are averaged over 30 simulations. Each simulation contains 100 presentations of input, which each consists of a uniformly random modulatory input being active (activation set to 1) while the driving input to both competitive neurons is simultaneously active (set to 1). The non-active modulatory input is set to 0. The network is exposed to each presentation for 100 iterations, followed by 10 iterations of all zero-valued input before the next presentation. State transitions are computed based on the state the network is in before and after the presentation of an input.

A connection is considered strongly learned if it meets the conditions of the spreading component of conflict learning (8).

C.2. Border ownership network architecture

To analyze our learning rule in feedback contexts, we focus on a model of border ownership similar to that developed by Craft et al. (2007). The network is organized into four layers of cells arranged retinotopically: the input, orientation selective, BO, and grouping layer (Fig. 3). For the main BO experiments, we used 40×50 grids of cells. As training time scales with network size, the smallest network that would still allow interesting stimuli to be presented was used.

The network is given grayscale input. The orientation layer receives input from the input layer, and uses fixed log Gabor filters (Field, 1987), parameterized by $\theta_{\text{gabor}} = \pi$ and $f = \sqrt{\pi}$, to compute four orientation maps (0, 45, 90, and 135°) representing a simplified V1-like layer. There are four orientation selective neurons per grid space, giving a total of 8,000 neurons.

The output of each of the four angles in the orientation layer below provides input to two BO neurons at the same retinotopic

location. These BO neurons are grouped into a column at each location with eight neurons, for a total of 16,000 BO neurons. The neurons within a column have inhibitory lateral connections, initially with equally distributed weight. BO neurons in a column have no concept of their physical position relative to any other neuron, nor their border ownership polarity (i.e. left/right, up/down, etc.). Initially, without any learning for the lateral connections, neurons within a column are unaware of the neuron with which they will most directly compete to form a BO pairing.

Each BO neuron provides feedforward input to all grouping cells within a radius r of its retinotopic position, and receives feedback from the same set of grouping cells. This radius r determines the scale of objects that can be handled by the network. Both the feedforward and feedback connections between these two layers are learned. The grouping layer is much more sparsely populated than both the input and border ownership layers, with roughly 1000 neurons placed randomly using a Poisson-disc algorithm (Bridson, 2007). Finally, there are lateral connections between grouping neurons in a center-surround fashion, extending to $0.6r$ for excitation, $3r$ for inhibition.

Training involves the repeated presentation of a moving square with length 10, chosen to be slightly smaller than the grouping neuron receptive field diameter $2r$ (see Figs. 3c and d). Squares are given a random initial position, orientation, and scaled up or down by up to 10% in size. Once placed, squares move in a random linear path across the FOV until no longer visible. Each positioning of a square is presented for 10 time steps to allow the network to settle. The network is given a blank input for 10 time steps after the square is no longer visible. Training is terminated after 40,000 squares are presented, a sufficient amount to show a plateau in the polarity scoring metric, described next.

For evaluation, we compute a polarity vector for each BO neuron, which represents the strength and preferred polarity direction of a neuron. The polarity vector is calculated as the sum of retinotopic vectors, each from the BO neuron to one of the grouping neurons providing it feedback (scaled by weight strength), multiplied by 1 or -1 depending on which side of the BO neuron's orientation they fall. The median absolute difference between the magnitude of polarity vectors for BO neurons of opposite polarity is then aggregated across all neurons of each orientation preference, to give the overall polarity score shown in Fig. 4. Significance is established with a Wilcoxon signed rank test. The polarity vectors are also used in Fig. 6, where the polarity vectors from neurons in the same column are weighted by activation and summed together to provide the resulting response vectors.

Finally, to compare the Hebbian learning rule versus all possible variants of conflict learning, a smaller 30×30 network is trained for each configuration (Fig. 7), using the same methodology as the larger network. For each configuration, we use the vertical orientation as an exemplar, computing a histogram of polarity scores across all vertical BO neurons. The medians for each score are then compared and tested for significance with a Wilcoxon signed rank test.

C.3. Orientation selective network architecture

The orientation selective network (Fig. 8) has three layers: an input layer, a center-surround layer, and an orientation selective layer, like that used to demonstrate the properties of GCAL (Stevens et al., 2013). The center-surround layer consists of both on-off and off-on preferential cells. In order to avoid anti-aliasing issues and a bias towards diagonals caused by square pixels, the resolution of the input is scaled by some amount, s , for the center-surround convolution. On-off cells have a difference-of-Gaussians receptive field, with a sigma of 0.33s for the larger Gaussian and 0.4s for the smaller one. The receptive field for an off-on cell is the negation of an on-off cell.

Table C.1
Parameter listing.

Parameter	Value	Description
σ_{noise}	0.01	Standard deviation of noise distribution.
θ_{ff}	$4\sigma_{\text{noise}}$	Threshold of driving input for neuron to be considered active.
θ_{inhib}	0.2	Threshold of inhibition for neuron to be considered inhibited.
s	0.1	Smoothing factor for threshold update.
θ_{min}	$4\sigma_{\text{noise}}$	Minimum threshold value.
θ_{max}	0.5 or 0.7	Maximum threshold value. Larger value used for orientation selective network.
η	0.01 or 0.001	Learning rate. Lower value used by Hebbian learning.
β	1.0	Balances positive versus negative learning for conflict learning.

Orientation selective neurons receive feedforward input from a disc of center-surround neurons (on-off and off-on) within some radius r , initially with equally distributed weight. Learning these connections creates the orientation selective behavior of the neurons. Orientation neurons further have lateral connections in a center-surround fashion to promote grouping and competition. Excitation extends to $0.27r$, inhibition to $0.73r$. Center-surround and orientation selective neurons are placed randomly using the same Poisson-disc algorithm as used for grouping neurons in the border ownership experiment. The center-surround and orientation selective layers have approximately 1,600 ($800 + 800$) and 3200 neurons, respectively, depending on the randomness of the Poisson-disc algorithm.

Training involves the repeated presentation of an oriented line segment spanning the width of the input layer. Lines are given a random initial position, orientation, and are translated across the field of view (FOV) in a random direction until no pixel of the line can be seen. Each position is held for 10 time steps, which is sufficient for the network to settle. The network is given a blank input for 10 time steps after the line is no longer visible in the FOV. Training is terminated after 20,000 lines are presented, a sufficient amount of time to maximize the selectivity score for a non-noisy network.

Orientations are assigned to neurons by finding the best fitting Gabor function and taking its orientation and coefficient of determination (r^2) values, using the MATLAB library `knkutils` by Kendrick Kay. The orientation is used for the hue in generating the color maps, whereas the coefficient of determination is used for determining selectivity. Pinwheel density is computed on orientation maps using code adapted from Topographica (Bednar, 2015) using the methods described in Stevens et al. (2013).

For the noise and stability measurements, noise is introduced by adjusting the standard deviation of the neuron activation noise term, ϵ , to σ_{noise} in the input layer (see results Fig. 9 for noise values used). The noise score is the average of the r^2 coefficients across all orientation selective neurons. For stability, the scoring metric is identical to the metric used by Stevens et al. (2013). We perform a paired-sample t-test to test for significance.

C.4. Parameter listing and source code

Key parameters for the learning and activation functions for both Hebbian learning and conflict learning are displayed in Table C.1. All parameters were tuned for each experiment to maximize performance with both rules in mind.

The experiments were performed using a custom framework written in C++ explicitly for conflict learning, with some analysis of results performed using MATLAB or Python scripts. All learning rules tested were implemented in this same framework. Source code is available on the website for conflict learning (Grant, Tanner & Itti, 2016).

References

- Abbott, L. F., & Nelson, S. B. (2000). Synaptic plasticity: taming the beast. *Nature Neuroscience*, 3, 1178–1183.
- Baluch, F., & Itti, L. (2011). Mechanisms of top-down attention. *Trends in Neurosciences*, 34(4), 210–224.
- Bar, M., Kassam, K. S., Ghuman, A. S., Boshyan, J., Schmid, A. M., Dale, A. M., et al. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences of the United States of America*, 103(2), 449–454.
- Bayerl, P., & Neumann, H. (2004). Disambiguating visual motion through contextual feedback modulation. *Neural Computation*, 16(10), 2041–2066.
- Bednar, J. A. (2012). Building a mechanistic model of the development and function of the primary visual cortex. *Journal of Physiology-Paris*, 106(5), 194–211.
- Bednar, J. A. (2015). Topographica: building and analyzing map-level simulations from Python, C/C++, MATLAB, NEST, or NEURON components. *Python in Neuroscience*, 104.
- Beuth, F., & Hamker, F. H. (2015). A mechanistic cortical microcircuit of attention for amplification, normalization and suppression. *Vision Research*, 116, 241–257.
- Bienenstock, E. L., Cooper, L. N., & Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *The Journal of Neuroscience*, 2(1), 32–48.
- Bridson, R. (2007). Fast Poisson disk sampling in arbitrary dimensions. In *SIGGRAPH sketches* (p. 22).
- Brosch, T., & Neumann, H. (2014a). Computing with a canonical neural circuits model with pool normalization and modulating feedback. *Neural Computation*.
- Brosch, T., & Neumann, H. (2014b). Interaction of feedforward and feedback streams in visual cortex in a firing-rate model of columnar computations. *Neural Networks*, 54, 11–16.
- Buschman, T. J., & Kastner, S. (2015). From behavior to neural dynamics: An integrated theory of attention. *Neuron*, 88(1), 127–144.
- Callaway, E. M. (2004). Feedforward, feedback and inhibitory connections in primate visual cortex. *Neural Networks*, 17(5), 625–632.
- Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1), 51–62.
- Chapman, B., Stryker, M. P., & Bonhoeffer, T. (1996). Development of orientation preference maps in ferret primary visual cortex. *The Journal of Neuroscience*, 16(20), 6443–6453.
- Clopath, C., Büsing, L., Vasilaki, E., & Gerstner, W. (2010). Connectivity reflects coding: a model of voltage-based STDP with homeostasis. *Nature Neuroscience*, 13(3), 344–352.
- Craft, E., Schütze, H., Niebur, E., & Von Der Heydt, R. (2007). A neural model of figure-ground organization. *Journal of Neurophysiology*, 97(6), 4310–4326.
- Cudeiro, J., & Sillito, A. M. (2006). Looking back: corticothalamic feedback and early visual processing. *Trends in Neurosciences*, 29(6), 298–306.
- Engert, F., & Bonhoeffer, T. (1997). Synapse specificity of long-term potentiation breaks down at short distances. *Nature*, 388(6639), 279–284.
- Faisal, A. A., Selen, L. P., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, 9(4), 292–303.
- Field, D. J. (1987). Relations between the statistics of natural images and the response properties of cortical cells. *JOSA A*, 4(12), 2379–2394.
- Fino, E., Paille, V., Cui, Y., Morera-Herreras, T., Deniau, J.-M., & Venance, L. (2010). Distinct coincidence detectors govern the corticoatrial spike timing-dependent plasticity. *The Journal of Physiology*, 588(16), 3045–3062.
- Friston, K., & Büchel, C. (2000). Attentional modulation of effective connectivity from V2 to V5/MT in humans. *Proceedings of the National Academy of Sciences*, 97(13), 7591–7596.
- Froemke, R. C. (2015). Plasticity of cortical excitatory-inhibitory balance. *Annual Review of Neuroscience*, 38, 195.
- Fukushima, K. (1980). Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics*, 36(4), 193–202.
- Grant, W.S., Tanner, J. and Itti, L. (2016). Con ict learning source code. ilab.usc.edu/conflictlearning/ Accessed 03.07.16.
- Grossberg, S. (2013). Adaptive resonance theory: How a brain learns to consciously attend, learn, and recognize a changing world. *Neural Networks*, 37, 1–47.

- Hebb, D. O. (1949). *The organization of behavior: A neuropsychological approach*. John Wiley & Sons.
- Hupe, J.-M., James, A. C., Girard, P., Lomber, S. G., Payne, B. R., & Bullier, J. (2001). Feedback connections act on the early part of the responses in monkey visual cortex. *Journal of Neurophysiology*, *85*(1), 134–145.
- Jones, H. E., Andolina, I. M., Ahmed, B., Shipp, S. D., Clements, J. T., Grieve, K. L., et al. (2012). Differential feedback modulation of center and surround mechanisms in parvocellular cells in the visual thalamus. *The Journal of Neuroscience*, *32*(45), 15946–15951.
- Jones, H. E., Andolina, I. M., Shipp, S. D., Adams, D. L., Cudeiro, J., Salt, T. E., et al. (2015). Figure-ground modulation in awake primate thalamus. *Proceedings of the National Academy of Sciences*, *112*(22), 7085–7090.
- 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*(6007), 1113–1116.
- Keil, W., Kaschube, M., Schnabel, M., Kisvarday, Z. F., Löwel, S., Coppola, D. M., et al. (2012). Response to comment on universality in the evolution of orientation columns in the visual cortex. *Science*, *336*(6080), 413.
- Kogo, N., & van Ee, R. (2014). Neural mechanisms of figure-ground organization: Border-ownership, competition and perceptual switching. In *Handbook of perceptual organization*.
- Larkum, M. (2013). A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends in Neurosciences*, *36*(3), 141–151.
- Levy, W., & Steward, O. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience*, *8*(4), 791–797.
- Lim, S., McKee, J. L., Woloszyn, L., Amit, Y., Freedman, D. J., Sheinberg, D. L., et al. (2015). Inferring learning rules from distributions of firing rates in cortical neurons. *Nature Neuroscience*.
- Linden, D. J., & Connor, J. A. (1995). Long-term synaptic depression. *Annual Review of Neuroscience*, *18*(1), 319–357.
- Major, G., Larkum, M. E., & Schiller, J. (2013). Active properties of neocortical pyramidal neuron dendrites. *Annual Review of Neuroscience*, *36*, 1–24.
- Markov, N. T., Vezoli, J., Chameau, P., Falchier, A., Quilodran, R., Huissoud, C., et al. (2014). Anatomy of hierarchy: feedforward and feedback pathways in macaque visual cortex. *Journal of Comparative Neurology*, *522*(1), 225–259.
- Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., & Wu, C. (2004). Interneurons of the neocortical inhibitory system. *Nature Reviews Neuroscience*, *5*(10), 793–807.
- Martin, A. B., & von der Heydt, R. (2015). Spike synchrony reveals emergence of proto-objects in visual cortex. *The Journal of Neuroscience*, *35*(17), 6860–6870.
- Miconi, T., & VanRullen, R. (2016). A feedback model of attention explains the diverse effects of attention on neural firing rates and receptive field structure. *PLoS Computational Biology*, *12*(2), e1004770.
- Mihalas, S., Dong, Y., von der Heydt, R., & Niebur, E. (2011). Mechanisms of perceptual organization provide auto-zoom and auto-localization for attention to objects. *Proceedings of the National Academy of Sciences*, *108*(18), 7583–7588.
- Oja, E. (1982). Simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, *15*(3), 267–273.
- Olshausen, B. A., & Field, D. J. (1997). Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, *37*(23), 3311–3325.
- Qiu, F. T., Sugihara, T., & von der Heydt, R. (2007). Figure-ground mechanisms provide structure for selective attention. *Nature Neuroscience*, *10*(11), 1492–1499.
- Roelfsema, P. R., Lamme, V. A., Spekreijse, H., & Bosch, H. (2002). Figureground segregation in a recurrent network architecture. *Journal of Cognitive Neuroscience*, *14*(4), 525–537.
- Sakai, K., & Nishimura, H. (2006). Surrounding suppression and facilitation in the determination of border ownership. *Journal of Cognitive Neuroscience*, *18*(4), 562–579.
- Sanger, T. D. (1989). Optimal unsupervised learning in a single-layer linear feedforward neural network. *Neural Networks*, *2*(6), 459–473.
- Serre, T., Wolf, L., & Poggio, T. (2005). Object recognition with features inspired by visual cortex. In *Computer vision and pattern recognition, 2005. CVPR 2005. IEEE computer society conference on. Vol. 2* (pp. 994–1000). IEEE.
- Shouval, H. Z., Samuel, S., & Wittenberg, G. M. (2010). Spike timing dependent plasticity: a consequence of more fundamental learning rules. *Spike-Timing Dependent Plasticity*, 60.
- Song, S., Miller, K. D., & Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, *3*(9), 919–926.
- Stevens, J.-L. R., Law, J. S., Antolík, J., & Bednar, J. A. (2013). Mechanisms for stable, robust, and adaptive development of orientation maps in the primary visual cortex. *The Journal of Neuroscience*, *33*(40), 15747–15766.
- Supèr, H., Romeo, A., & Keil, M. (2010). Feed-forward segmentation of figure-ground and assignment of border-ownership. *PLoS One*, *5*(5), e10705.
- Turrigiano, G. G., & Nelson, S. B. (2004). Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, *5*(2), 97–107.
- Varela, C. (2014). Thalamic neuromodulation and its implications for executive networks. *Frontiers in Neural Circuits*, *8*, 69.
- Vogels, T. P., Froemke, R. C., Doyon, N., Gilson, M., Haas, J. S., Liu, R., et al. (2013). Inhibitory synaptic plasticity: spike timing-dependence and putative network function. *Frontiers in Neural Circuits*, *7*(EPFL-REVIEW-189448).
- von der Heydt, R. (2015). Figure-ground organization and the emergence of proto-objects in the visual cortex. *Frontiers in Psychology*, 6.
- Wagatsuma, N., von der Heydt, R., & Niebur, E. (2016). Spike Synchrony Generated by Modulatory Common Input through NMDA-type Synapses. *Journal of Neurophysiology*, jn-01142.
- Wang, L., & Maffei, A. (2014). Inhibitory plasticity dictates the sign of plasticity at excitatory synapses. *The Journal of Neuroscience*, *34*(4), 1083–1093.
- Widloski, J., & Fiete, I. R. (2014). A model of grid cell development through spatial exploration and spike time-dependent plasticity. *Neuron*, *83*(2), 481–495.
- Wiesel, T. N., Hubel, D. H., et al. (1963). Single-cell responses in striate cortex of kittens deprived of vision in one eye. *Journal of Neurophysiology*, *26*(6), 1003–1017.
- Yantis, S. (2008). The neural basis of selective attention cortical sources and targets of attentional modulation. *Current Directions in Psychological Science*, *17*(2), 86–90.
- Zenke, F., Agnes, E. J., & Gerstner, W. (2015). Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nature Communications*, 6.
- Zhaoping, L. (2005). Border ownership from intracortical interactions in visual area V2. *Neuron*, *47*(1), 143–153.
- Zhou, H., Friedman, H. S., & Von Der Heydt, R. (2000). Coding of border ownership in monkey visual cortex. *The Journal of Neuroscience*, *20*(17), 6594–6611.
- Zucker, R. S., & Regehr, W. G. (2002). Short-term synaptic plasticity. *Annual Review of Physiology*, *64*(1), 355–405.