Neuron, Vol. 17, 371–374, September, 1996, Copyright ©1996 by Cell Press

Synaptic Economics: Competition and Cooperation in Synaptic Plasticity

Minireview

Kenneth D. Miller
Departments of Physiology and Otolaryngology
Neuroscience Graduate Program
W. M. Keck Center for Integrative Neuroscience
Sloan Center for Theoretical Neurobiology at UCSF
University of California
San Francisco, California 94143-0444

There is abundant evidence that correlation-based rules of synaptic modification, which lead to the outcome "neurons that fire together, wire together," are followed in multiple biological systems (Shatz, 1990). Under these rules, synapses that work together are rewarded by being strengthened. Mechanisms of such cooperation are reasonably well understood in the case of NMDA-dependent long-term potentiation (LTP) in hippocampus (Bear and Malenka, 1994). However, synapses also have a competitive side: when some synapses grow stronger and prosper, others, which left to themselves would also have strengthened, instead weaken. As an example, consider the process of ocular dominance segregation (Miller and Stryker, 1990). Afferents representing the two eyes are initially intermixed in primary visual cortex, but then segregate into eye-specific patches through a correlation-based process. Assuming left eye inputs fire together and right eye inputs fire together, something must prevent both from wiring together on the same cells.

It is often imagined that LTD can account for such competition, since it provides a mechanism for synaptic weakening; but LTD as characterized thus far is not up to the task. Imagine repeatedly but independently tetanizing two pathways onto the same postsynaptic targets in hippocampus, a hippocampal model of afferent segregation. At least in the short term, the times over which LTP and LTD are studied, both pathways will end up strongly potentiated. Something more complicated is afoot here.

The fact of competition brings up the issue of scale: how correlated must neuronal firing be in order for the neurons to wire together? More generally, which patterns of pre- and postsynaptic activities lead to strengthening or weakening of synapses? The answer could be an absolute one, determined, for example, by fixed biophysical requirements for the opening of NMDA receptors. But the fact of synaptic competition suggests that the scale is determined relative to other variables, such as the activities of other competing neurons, the recent history of pre- and postsynaptic activities (termed "metaplasticity" by Abraham and Bear, 1996), or the current strengths or recent history of the synaptic weights. Indeed, it seems logical that evolution would endow a cell with a means of controlling the strength of the input that it receives, to ensure that it remains within a meaningful operating range. Such homeostatic control is not a part of our current understanding of LTP and LTD.

A recent paper in *Neuron* (Wong and Oakley, 1996) beautifully adds to the evidence that simple intuitions about correlation-based plasticity may lead us astray. In the LGN of the ferret, ON-center and OFF-center inputs segregate into separate layers during the third postnatal week. Blockade of NMDA receptors prevents this segregation (Hahm et al., 1991), suggesting an activity-dependent and correlation-based mechanism. Wong and Oakley (1996) studied spontaneous bursting activity in the developing retina. Using calcium imaging, they found that anatomically identified ON- and OFF-center cells burst together and indistinguishably before postnatal day (P)14. Beginning around P14, however, OFF-center cells come to burst three to four times more frequently than ON-center cells. Thus, ON- and OFF-center activity patterns become distinguishable just at the time at which their axons begin to segregate in the LGN.

However, during this period of segregation, the less frequent ON-center bursts tend to be coincident with OFF-center bursts. This raises an obvious question for a correlation-based rule: if coincident bursting leads to LTP of all the bursting inputs, why don't ON- and OFF-center neurons wire together? What instead leads to a competitive outcome, in which the two input types segregate onto different postsynaptic cells?

Here, I will review theoretical and experimental work that allow these results to be addressed without abandoning an explanation in terms of correlation-based plasticity. More generally, this framework may suggest directions for experimental work on the relative roles of cooperation and competition in correlation-based development.

Theoretical Approaches to Achieving Competition

Theoretical work on neural development can be briefly and heuristically summarized as follows (Miller, 1996). Cooperation and competition among the inputs to a single cell select the set of inputs received by the cell, and thus determine receptive field structure. Cooperation and competition between postsynaptic cells during correlation-based development determine the arrangement or maps of receptive field properties across the postsynaptic structure. Most theoretical work focuses on how these processes can account for the specifics of observed receptive fields and maps. Here, we focus instead on insights from theory into possible mechanisms of synaptic cooperation and competition.

We begin by considering development of the inputs to a single cell. Theory currently does not distinguish strongly between alternative mechanisms of cooperation, e.g., correlation-based strengthening and weakening of fixed synapses versus sprouting and retraction of synapses with correlation-based stabilization or destabilization. If each mechanism can explore the same set of possible connectivity patterns, each will converge to a similar pattern, one which roughly optimizes correlations under rules that define cooperation and competition.

However, theorists have long found it necessary to

add mechanisms of competition to correlation-based rules to achieve the outcomes observed biologically (reviewed by Miller and MacKay, 1994). Simple correlation-based rules lead to instability, owing to positive feedback (input activity drives the postsynaptic cell, causing potentiation, yielding stronger driving of the postsynaptic cell, causing more potentiation...). Furthermore, if afferents sort out onto different postsynaptic cells, and if the activity of each afferent is sufficient to achieve potentiation, how is it that some afferents (left eye) win and others (right eye) lose on any given cell? To solve these problems, correlation-based development typically must be subject to some constraint, for example, conservation of the total synaptic weight onto, or average activity level of, the postsynaptic cell.

Such constraints have one obvious problem: the total number of synapses and the overall activity levels in a tissue typically change greatly during development. However, overall changes in synapse number are activity independent in several systems (Bourgeois et al., 1989; Hayes and Meyer, 1989). This suggests that activity-dependent processes do not expand or contract the pool of synaptic resources, but simply allocate it. Hence, for now, we simplify by considering competition for some fixed level of resources.

To better understand theoretical approaches, it is helpful to write simple equations for a plasticity rule for a single postsynaptic cell. Let y be some measure of the activity of the postsynaptic cell, x_i a measure of the activity of the i^{th} input, and w_i the synaptic weight of the i^{th} input. The simplest correlation-based equation for the change in a synaptic weight, Δw_i, per unit time states that the weight should grow in proportion to the product of the postsynaptic activity level, y, and the presynaptic activity level, x_i. If we say that y must be above some threshold level θ_y to achieve LTP, and otherwise yields LTD; and allow a similar possibility for x_i, then we arrive at the equation:

$$\Delta w_i \propto (y - \theta_y)(x_i - \theta_x^i) \quad (1)$$

If both y and x_i are above their thresholds, LTP occurs; if one is below its threshold and the other above, LTD occurs. If both are below their thresholds, this rule would imply that LTP occurs, but this can be ignored; Δw_i can generally be set to 0 in this case without significantly altering the developmental outcome.

We now introduce a simple linear rule for postsynaptic activity, given by summing the input activities multiplied by their synaptic weights: $y = \Sigma_j w_j x_j$. Combining these equations, we obtain: $\Delta w_i \propto \Sigma_j [(x_i - \theta_x^i)x_j]w_j +$ (Terms that don't depend on w). The main point to notice about this equation is the quantity $(x_i - \theta_x^i)x_j$ that multiplies the weight w_j. We call the average value over time of this quantity the correlation C_{ij} between the activities of inputs i and j:

$$C_{ij} = \langle (x_i - \theta_x^i)x_j \rangle \quad (2)$$

Here, angle brackets indicate time averaging. C_{ij} tells quantitatively how the presence of one synaptic weight, w_j, affects the development of another, w_i, under the correlation-based rule.

Having defined the basic terms of discussion, we can now examine the methods theorists have developed for achieving competitive outcomes (Miller and MacKay, 1994).

Heterosynaptic or Associative LTD

Heterosynaptic LTD refers to depression of an inactive synapse owing to the activity of other inputs ($y > \theta_y$, $x_i < \theta_x^i$ in Equation 1), and associative LTD to the depression of an active synapse owing to the absence of postsynaptic activity (i.e., $y < \theta_y$, $x_i > \theta_x^i$ in Equation 1; Linden, 1994; Huerta and Lisman, 1995; Scanziani et al., 1996; Cummings et al., 1996). For such LTD to achieve competition, the synaptic strength lost through LTD must roughly equal the strength gained through LTP. In general, this is a delicate balance, easily upset by changes in input activity patterns or by ongoing activity-dependent changes in synaptic strength. Because of this fragility, it is unlikely that simple correlation-based LTD can adequately explain competition.

In terms of our simple model, the requirement for achieving competition through LTD is roughly that input activities have at least as much anticorrelation as correlation: $\Sigma_i C_{ij} \leq 0$ for all j. Two simple examples of such activity patterns are as follows: first, if correlations between inputs separated by a given distance are positive at short distances but negative at longer distances; and second, if there are two types of inputs (e.g., ON and OFF), such that correlations between two inputs of opposite type at a given separation are roughly equal and opposite to correlations between two same type inputs at that separation. Both of these are postulated to occur among ON- and OFF-center inputs to visual cortex at the time of simple cell development (Miller, 1994).

Rough Conservation of Total Synaptic Strength on the Postsynaptic Cell

Conservation of synaptic weight ensures that if some weights grow, others must correspondingly shrink. Heterosynaptic LTD can achieve this, but only with appropriate input correlations, as just discussed. Other mechanisms include: 1) If the presynaptic threshold θ_x^i increases sufficiently as the postsynaptic activity, y, or weight w_i (or both) increase, then weight conservation can be achieved. This is a "sliding" presynaptic plasticity threshold: the threshold presynaptic activity level θ_x^i changes value in response to changes in postsynaptic activity or weight. 2) Other homeostatic mechanisms might exist that maintain total postsynaptic strength near some desired level. Examples include competition among synapses for a finite resource, such as receptors or a trophic factor, cellular control of the number of synapses on a cell (e.g., Hayes and Meyer, 1988; Xiong et al., 1994) along with a limited range of strength for each synapse, or cellular regulation of the overall rates of synaptic growth or degradation.

Rough Conservation of Mean Postsynaptic Activity

If mean postsynaptic activity is maintained near some set point, then again an increase in some weights requires corresponding shrinkage in other weights. Activity conservation can be achieved in several ways: 1) If the postsynaptic threshold θ_y increases faster than linearly with the average postsynaptic activity $\langle y \rangle$ (i.e., $\theta_y \propto \langle y \rangle^p$ for $p > 1$), then the synaptic weights of the system will adjust to keep the average postsynaptic activity near a set point value (Bienenstock et al., 1982).

This is a "sliding" postsynaptic plasticity threshold. 2) If all excitatory synapses onto a cell are strengthened when the overall activity level of the cell decreases, and vice versa (Turrigiano et al., 1996, Soc. Neurosci., abstract), then activity conservation can be achieved. 3) If the excitability of the cell is regulated by activity, for example, through activity-dependent regulation of intrinsic conductances (LeMasson et al., 1993; Turrigiano et al., 1994), then postsynaptic activity can remain near some set point. Since the effect of a synaptic conductance is weighted by the excitability of the cell, such regulation of excitability is functionally similar to activity-dependent weight renormalization. 4) Other homeostatic mechanisms might exist that maintain average postsynaptic activity near some set point, for example, activity-dependent regulation of the amount of trophic factor available to synapses.

Interplay of Inhibitory Circuitry with These Factors

Inhibitory circuitry may play an important role in associative LTD by ensuring that a postsynaptic cell is inhibited when certain inputs to that cell are active. Inhibitory feedback can also provide an alternative means for activity-dependent regulation of the effective excitability of a cell. Finally, inhibitory circuitry, and plasticity thereof, could compensate for changes in average input activities, for example, by leaving output activities roughly invariant. This could itself provide a mechanism of competition (by suppressing, rather than eliminating, the "losing" inputs). This could also strongly interact with mechanisms that tend to control total postsynaptic strength or activity.

Cooperation and Competition between Different Postsynaptic Cells

Just as the mechanisms of competition among inputs to a single cell remain mysterious, so too do the mechanisms of both cooperation and competition between postsynaptic cells. Inter-cell cooperation refers to influences that lead two cells to develop similar sets of inputs; inter-cell competition refers to influences that lead two cells to develop differing patterns of inputs. A glance at columnar structure in cerebral cortex suggests the existence of cooperation among nearby cells (e.g., within 100 μm), and competition among more distant cells (e.g., at 200–400 μm).

Presumably, cooperation occurs through some type of positive influences between cells, whether via excitatory synapses, gap junctions (Peinado et al., 1993), or activity-dependent release and uptake of diffusible factors (Schuman and Madison, 1994; Bonhoeffer, 1996; Scanziani et al., 1996). Knowledge of competitive mechanisms is even more vague. Theorists typically assume a net inhibitory synaptic interaction between tissues separated by some distance (lateral inhibition). Competition can also be achieved by assuming that afferents with similar mean activities will each project a roughly equal amount of synaptic strength. For example, there might be competition within an afferent for a limited supply of vesicles or transmitter, or total projected synaptic strength might be limited by a trophic factor that afferents receive in proportion to their mean activity. Then an afferent that projects to some cells must withdraw from other cells, while every afferent must project somewhere. This has an effect very similar to lateral inhibition and achieves competition between postsynaptic cells (Miller and MacKay, 1994).

Lee and Wong (1996, Soc. Neurosci., abstract) have demonstrated that several of the simple competitive models described here can achieve ON/OFF segregation, using as inputs the spontaneous activity patterns observed by Wong and Oakley (1996) during P14–P21.

Possible Substrates of Competition

A number of recent results provide intriguing evidence as to possible mechanisms of competition among inputs to a single cell. Scanziani et al. (1996) found that LTP of some inputs in hippocampus leads to LTD of inputs in a neighboring volume, independently of the voltage or Ca^{2+} level of the postsynaptic cell. Such widespread LTD could, for example, ensure rough conservation of synaptic weight over a postsynaptic cell, if the net increase due to LTP were balanced by the net decrease due to LTD. The lack of dependence on postsynaptic activity might remove some of the fragility associated with correlation-based LTD.

Several homeostatic properties that may be relevant to synaptic competition have recently been demonstrated. Theoretical work (LeMasson et al., 1993) showed that activity-dependent modification of intrinsic membrane properties could maintain a neuron in a stable activity pattern, despite channel turnover, cell growth, and other perturbations. Turrigiano et al. (1994) demonstrated that behavior expected from this model is realized by lobster stomatogastric ganglion neurons in culture. Recently, in studies of cultured neurons from visual cortex, Turrigiano et al. (1996, Soc. Neurosci., abstract) found that prolonged blockade of activity by tetrodotoxin leads to an increase in the amplitude of CNQX-sensitive miniature synaptic currents (minis), while enhancement of activity by blockade of GABA-A inhibition leads to a decrease in mini amplitudes. The amplitude histograms are shifted multiplicatively, suggesting that synaptic strengths may be increased or decreased in proportion to their prior values. This would maintain the selectivity of a neuron for different input patterns, while simply raising or lowering the gain of the neuron's response. These results may be the first glimpses into a world of homeostatic regulation of electrical properties of neurons and synapses.

Several authors have recently reported circumstances in which a synaptic plasticity threshold appears to change as a function of the recent history of activity (Bear, 1995; Abraham and Bear, 1996; Kirkwood et al., 1996). The presence of a sliding presynaptic or postsynaptic plasticity threshold does not imply a competitive outcome. For example, an oft-studied plasticity rule is the covariance rule:

$$\Delta w_i \propto (y - \langle y \rangle)(x_i - \langle x_i \rangle).$$

Here, $\langle y \rangle$ is the average output cell activity, while $\langle x_i \rangle$ is the average activity of the i^{th} input. In this equation, the mean postsynaptic activity $\langle y \rangle$ constitutes a sliding postsynaptic plasticity threshold: it divides LTP from LTD for a fixed presynaptic activity level, and it slides

depending on the history of activity, so that if the postsynaptic cell is activated more often, its plasticity threshold rises accordingly. Similarly, $\langle x_i \rangle$ is a sliding presynaptic plasticity threshold. Yet the covariance rule does not achieve competition (Miller and MacKay, 1994). For a sliding plasticity threshold to achieve competition, at least within the framework of simple plasticity rules, one of the conditions described above must hold: either the presynaptic threshold must grow with increases in postsynaptic activity or in synaptic weight, or a postsynaptic threshold must increase faster than linearly with the average postsynaptic activity. The sliding thresholds of the covariance rule meet neither requirement.

The experiments that most directly suggest a sliding threshold (Bear, 1995; Kirkwood et al., 1996) assess the frequencies of presynaptic stimulation that yield LTP or LTD, and refer to the dividing line between LTP and LTD as the threshold frequency. Because pre- and postsynaptic activities covary in this protocol, it is unclear whether the threshold is one of presynaptic activity, of postsynaptic activity, or of some combination. Because these experiments involve global changes in biochemistry or neuronal activities, it is also unclear whether the threshold change is induced by presynaptic activity, postsynaptic activity, or something else. In several other experiments showing "metaplasticity" (activity-dependent changes in the conditions that yield LTP or LTD of a given strength), it has been established that the effect is input specific (Abraham and Bear, 1996), dependent on the pattern of presynaptic rather than postsynaptic activity. Such input specificity would largely rule out the possibility of achieving competition, since the activity of one input would have no effect on the plasticity of other inputs.

Mechanisms for cooperation and competition between postsynaptic cells are also an intriguing subject for experimental study. A number of mechanisms for nonsynaptic cross-talk between postsynaptic neurons have been uncovered in recent years (Peinado et al., 1993; Schuman and Madison, 1994; Bonhoeffer, 1996; Scanziani et al., 1996). All of these mechanisms seem to subserve inter-cell cooperation, rather than competition. That is, the net effect is to strengthen a synchronously firing set of inputs or to weaken less active inputs (or both), on all of the postsynaptic cells involved. To achieve inter-cell competition, a synchronous set of inputs must instead be strengthened on one group of cells while being weakened on another group of cells. The question of how competition is achieved, whether through inhibitory synaptic circuitry during development, through a rough set point for the total synaptic strength supported by an afferent, or through other mechanisms yet to be discovered, is a critical one.

What Is To Be Done?

Current knowledge of LTP and LTD are not sufficient for an understanding of correlation-based development and plasticity. Experiments, and further theoretical work, aimed at understanding the mechanisms of synaptic competition are needed. Experimental work on this topic is relatively new, but the current stirrings may portend a rapidly growing effort. As possible mechanisms are uncovered, it is important to test whether these mechanisms could achieve competition. Here, theory can play an important role, as illustrated above in the assessment of the role of sliding plasticity thresholds. One simple thought experiment, and perhaps a useful actual experiment, is that proposed in the introduction: repeatedly but independently giving LTP-inducing stimuli to two pathways in a hippocampal slice. In the short term, both pathways will be potentiated; perhaps on a longer time scale, mechanisms would be revealed that would reduce the synaptic strength of one pathway in favor of the other, as in ocular dominance segregation. Correlation-based learning is only one-half of the synaptic economy; the relentless drive of competition must also be understood.

Selected Reading

Abraham, W.C., and Bear, M.F. (1996). Trends Neurosci. *19*, 126–130.

Bear, M.F. (1995). Neuron *15*, 1–4.

Bear, M.F., and Malenka, R.C. (1994). Curr. Opin. Neurobiol. *4*, 389–399.

Bienenstock, E.L., Cooper, L.N., and Munro, P.W. (1982). J. Neurosci. *2*, 32–48.

Bonhoeffer, T. (1996). Curr. Opin. Neurobiol. *6*, 119–126.

Bourgeois, J.-P., Jastreboff, P.J., and Rakic P. (1989). Proc. Natl. Acad. Sci. USA *86*, 4297–4301.

Cummings, J.A., Mulkey, R.M., Nicoll, R.A., and Malenka, R.C. (1996). Neuron *16*, 825–833.

Hahm, J.O., Langdon, R.B., and Sur, M. (1991). Nature *351*, 568–570.

Hayes, W.P., and Meyer, R.L. (1988). J. Comp. Neurol. *274*, 539–559.

Hayes, W.P., and Meyer, R.L. (1989). J. Neurosci. *9*, 1414–1423.

Huerta, P.T., and Lisman, J.E. (1995). Neuron *15*, 1053–1063.

Kirkwood, A., Rioult, M.G., and Bear, M.F. (1996). Nature *381*, 526–528.

LeMasson, G., Marder, E., and Abbott, L.F. (1993). Science *259*, 1915–1917.

Linden, D.J. (1994). Neuron *12*, 457–472.

Miller, K.D. (1994). J. Neurosci. *14*, 409–411.

Miller, K.D. (1996). In Models of Neural Networks III, E. Domany, J.L. van Hemmen, and K. Schulten, eds. (New York: Springer-Verlag), pp. 55–78. (Available as ftp://ftp.keck.ucsf.edu/pub/ken/miller95.ps.Z.)

Miller, K.D., and MacKay, D.J.C. (1994). Neural Comput. *6*, 100–126.

Miller, K.D., and Stryker, M.P. (1990). In Connectionist Modeling and Brain Function: The Developing Interface, S.J. Hanson and C.R. Olson, eds. (Cambridge, Massachusetts: MIT Press/Bradford), pp. 255–350.

Peinado, A., Yuste, R., and Katz, L.C. (1993). Neuron *10*, 103–114.

Scanziani, M., Malenka, R.C., and Nicoll, R.A. (1996). Nature *380*, 446–450.

Schuman, E.M., and Madison, D.V. (1994). Science *263*, 532–536.

Shatz, C.J. (1990). Neuron *5*, 745–756.

Turrigiano, G.G., Abbott, L.F., and Marder, E. (1994). Science *264*, 974–977.

Wong, R.O. and Oakley, D.M. (1996). Neuron *16*, 1087–1095.

Xiong, M., Pallas, S.L., Lim, S., and Finlay, B.L. (1994). J. Comp. Neurol. *344*, 581–597.