

Hebbian learning and development

Yuko Munakata and Jason Pfaffly

Department of Psychology, University of Colorado Boulder, USA

Abstract

Hebbian learning is a biologically plausible and ecologically valid learning mechanism. In Hebbian learning, 'units that fire together, wire together'. Such learning may occur at the neural level in terms of long-term potentiation (LTP) and long-term depression (LTD). Many features of Hebbian learning are relevant to developmental theorizing, including its self-organizing nature and its ability to extract statistical regularities from the environment. Hebbian learning mechanisms may also play an important role in critical periods during development, and in a number of other developmental phenomena.

Introduction

Hebbian learning algorithms are highly relevant for investigating development. As elaborated below, such algorithms are biologically plausible and ecologically valid. Further, Hebbian learning algorithms can account for a wide range of behaviors and changes during development. These include critical periods, development of cortical receptive fields, learning of statistical regularities in the environment, development of object knowledge, learning of phonemes, and development of flexible behaviors (Blair, Intrator, Shouval & Cooper, 1998; McClelland, Thomas, McCandliss & Fiez, 1999; Morton & Munakata, 2002a; O'Reilly & Johnson, 2002).

Note that Hebbian learning algorithms do not represent a *new* alternative to backpropagation. Hebbian learning algorithms existed prior to the advent of backprop, and were ultimately supplanted by backprop in many domains due to backprop's greater power to master complex tasks. However, the basic mechanisms of Hebbian learning are important in their own right. Moreover, Hebbian-like mechanisms play a role in more powerful learning algorithms that address the biological implausibility of backprop.

Overview of Hebbian learning

Biological basis of Hebbian learning

Donald Hebb was the first to suggest that the 'efficiency' of a given neuron, in contributing to the firing of another, could increase as that cell is repeatedly involved

in the activation of the second (Hebb, 1949). Thus, the basic tenet of Hebbian learning in neural networks is that 'units that fire together, wire together'. This idea is consistent with the neural mechanisms of long-term potentiation (LTP) (Bliss & Lomo, 1973) and long-term depression (LTD) (e.g. Artola, Brocher & Singer, 1990). LTP and LTD influence the extent to which activity in a sending neuron leads to depolarization of a receiving neuron, by influencing the efficacy of synapses, or junctions between neurons. LTP is a long-lasting potentiation (increase) of synaptic efficacy, while LTD is a long-lasting depression (weakening) of synaptic efficacy.

LTP and LTD likely rely on calcium ions entering postsynaptic dendrites through NMDA channels. LTP may be understood in terms of simultaneous presynaptic and postsynaptic activation. Presynaptic activation causes the release of the excitatory neurotransmitter glutamate, which can then bind with and open postsynaptic NMDA receptors. Postsynaptic activation causes magnesium ions to move out of the opening of NMDA receptor channels, which they would otherwise block. Calcium ions can then enter the postsynaptic neuron, where they are involved in a series of molecular changes that ultimately increase overall synaptic efficacy. Both presynaptic and postsynaptic activation are thus required for this process. Neurons that fire together could wire together through these increases in synaptic efficacy. LTD may occur when a smaller concentration of calcium enters the postsynaptic neuron (Artola *et al.*, 1990; Bear & Malenka, 1994; Lisman, 1994). This process might reflect a lack of coordinated activity between presynaptic and postsynaptic neurons. Neurons that do not fire together would not wire together due to these decreases in synaptic efficacy.

Address for correspondence: Yuko Munakata, Department of Psychology, University of Colorado Boulder, 345 UCB, Boulder, CO 80309-0345, USA; e-mail: munakata@psych.colorado.edu

Ecological validity of Hebbian learning

In addition to being biologically plausible, the mechanisms underlying Hebbian learning are ecologically valid. These mechanisms rely only on inputs coming into the system to produce patterns of activity. Learning proceeds on the basis of these resulting patterns of activity. Thus, such models can learn without needing to produce particular responses, without any explicit tasks or directed attempts to learn, and without requiring teaching signals.

Network implementation of Hebbian learning

Hebbian learning is implemented in neural network models through changes in the strength of connection weights between units. The strength of a connection weight determines the efficacy of a sending unit in activating a receiving unit, and thus captures this aspect of synaptic efficacy. In Hebbian learning, weights change as a function of units' activity levels. Through such learning, weights come to reflect statistical regularities in the environment, with networks self-organizing so that different units learn to represent different environmental regularities. The mathematical basis for such learning is described below.

A basic Hebbian learning rule takes the following form:

$$\Delta w_{ij} = \varepsilon a_i a_j \quad (1)$$

where Δw_{ij} denotes the change in weight from unit i to unit j , a_i and a_j denote the activation levels of units i and j respectively, and ε denotes the learning rate – how quickly the weights change in response to unit activations.¹ The activation level of a given unit typically varies between 0 and 1, and is computed as a nonlinear function of the activations of other units and the strength of their connections to the unit. Equation 1 captures the basic notion of Hebbian learning that units that fire together become strongly connected, because when a_i and a_j are simultaneously large, the weight increase will also be large. The weight value at any given time t can then be computed as follows:

$$w_{ij}(t) = w_{ij}(t-1) + \Delta w_{ij} \quad (2)$$

That is, the weight from unit i to unit j at any given time t is the weight value from the previous time step plus the change in weight resulting from the activity of units i and j .

¹ The Δ and a terms are functions of time, but the time terms are omitted to simplify the presentation.

Although Equation 1 captures some aspects of the notion of Hebbian learning, there is a key problem with this basic form of the equation: Weights can increase, but they cannot decrease. This is problematic for two reasons. First, over successive time steps, learning according to this rule will drive the weights within a network to become infinitely large. Such learning would be computationally problematic and biologically implausible. Second, because the weights cannot decrease according to the basic equation, it fails to capture the known biological mechanism of LTD. Both problems can be addressed by normalizing the weight updates with an equation like the following (see also Oja, 1982):

$$\Delta w_{ij} = \varepsilon a_j (a_i - w_{ij}) \quad (3)$$

This equation subtracts the weight value from the sending unit activation, which allows Δw_{ij} to be negative. Weights can thus decrease, addressing the problems of infinite growth and a missing LTD mechanism in Equation 1. Equation 3 leads to a form of soft weight bounding, where increases to a weight become less likely as the weight becomes quite large, and decreases to the weight become less likely as the weight becomes quite small. This same pattern is observed in LTP/LTD.

In addition to providing a computationally useful and biologically plausible form of learning, Hebbian learning rules lead units to represent correlations in the environment. For example, a form of Hebbian learning like that in Equation 3 performs principal components analysis on the correlations between inputs (Oja, 1982). Equation 3 also leads connection weights to represent the conditional probability that a sending unit was active given that the receiving unit was active (O'Reilly & Munakata, 2000; Rumelhart & Zipser, 1986), or:

$$w_{ij} = P(a_i | a_j) \quad (4)$$

Thus, whenever a given receiving unit (j) is active, if a sending unit (i) also tends to be active (i.e. the conditional probability of a_i given a_j is high – near 1), the weight will similarly be high. In contrast, whenever a given receiving unit is active, if a sending unit tends not to be active (i.e. the conditional probability of a_i given a_j is low – near 0), the weight will similarly be low. In this way, Hebbian learning yields weights that reflect conditional probabilities of activities, and in turn, yields units that represent correlations in the environment.

The self-organizing nature of Hebbian networks, supported through differences in starting weights and inhibitory competition among units, allows different units to represent different correlations in the environment. A given unit might start with a set of connections that

give it a slight advantage in responding to a certain kind of correlation in the environment (e.g. bars of light in a vertical orientation). Hebbian learning would increase the connections to this unit from relevant sending units, so the unit would become increasingly responsive to this correlation. And Hebbian learning would decrease the connections to this unit from irrelevant sending units, so the unit would become less responsive to other correlations in the environment (e.g. bars of light in a horizontal orientation). Inhibitory competition among the units would lead other units to be less responsive to the correlation represented by the first unit. These other units could then tune their weights to respond to other correlations in the environment. In this way, Hebbian models can self-organize so that units specialize in representing distinct statistical regularities in the environment.

Many different types of networks use Hebbian-style learning (e.g. Bienenstock, Cooper & Munro, 1982; Kohonen, 1984; Miller, 1994). In the Kohonen network (Kohonen, 1984), a network is presented with a particular stimulus, and the 'winner' unit (the one with the most input) is activated, along with its neighboring units, with activation decreasing with distance from the winner. This kind of activation pattern simulates the clusters of activity that result in systems when connections are locally excitatory but globally inhibitory (Konen & Von Der Malsburg, 1994). With Hebbian learning in such a network, weights change as a function of the simultaneous activity of units; because neighboring units become activated together, their weights change in similar ways, such that units learn to respond to similar inputs as their neighbors. This yields topographic maps like those in the brain, where groups of neighboring neurons respond similarly to particular stimuli (e.g. Blasdel & Salama, 1986; Livingstone & Hubel, 1988; Merzenich & Kaas, 1980). As elaborated below, Kohonen networks have also been used to explore critical period effects in language perception.

Features of Hebbian learning relevant to development

Many features of Hebbian learning are relevant to developmental theory and data, such as the ability to extract statistical regularities from the environment. Infants are very good at learning such regularities from relatively brief exposures to auditory or visual stimuli (Aslin, Saffran & Newport, 1998; Kirkham, Slemmer & Johnson, 2002; Maye, Werker & Gerken, 2002; Saffran, Johnson, Aslin & Newport, 1999). These statistical learning abilities are general across adults (Fiser &

Aslin, 2001; Saffran *et al.*, 1999) and other species (Hauser, Newport & Aslin, 2001). Hebbian learning can support such statistical learning. For example, when networks are presented with images of natural scenes, Hebbian learning allows them to develop representations that capture important statistical correlations present in the images, namely, the regular presence of edges in the images (e.g. Blair *et al.*, 1998; O'Reilly & Munakata, 2000). The resulting representations include many important properties of cortical receptive fields, such as edge detectors varying in size, position and orientation. Hebbian learning also allows models to extract statistical regularities of linguistic inputs, for example, the extent to which words co-occur in the environment (O'Reilly & Munakata, 2000). This co-occurrence information is remarkably useful for capturing semantic information about words (Landauer & Dumais, 1997). Note that statistical learning is not unique to Hebbian models. It is a general ability explored across a range of neural network models (e.g. Altmann & Dienes, 1999; Christiansen & Curtin, 1999; Lee & Seung, 1999; Rogers & McClelland, *in press*; Seidenberg & Elman, 1999).

However, Hebbian learning may be particularly relevant for certain kinds of statistical learning (as well as other kinds of learning) that appear to occur incidentally, without any explicit task or directed attempt to learn. Thus, another relevant feature of Hebbian learning for development is its automatic, self-organizing nature. This kind of learning can proceed simply in response to an array of inputs from the environment, without any consideration of what outputs should be produced in response to those inputs. In contrast, in error-driven algorithms such as backprop, learning requires a network's responses to be compared to target responses. Such error-driven algorithms can be reconceptualized in terms of networks producing outputs that are predictions about what will occur next in the environment, with subsequent inputs providing target information (McClelland, 1994; Munakata, McClelland, Johnson & Siegler, 1997); this obviates the need for producing outputs and comparing them to an explicit training signal. This reconceptualization of error-driven algorithms may broaden their relevance to a range of developmental domains. Nonetheless, Hebbian learning may provide a closer match to the kinds of learning shown by infants in certain cases, such as statistical learning tasks. For example, rather than infants forming predictions and learning from their errors based on brief exposures to arbitrary and rapid streams of phonemes or visual sequences, it may be more plausible that Hebbian type learning mechanisms automatically extract statistical regularities from such stimuli.

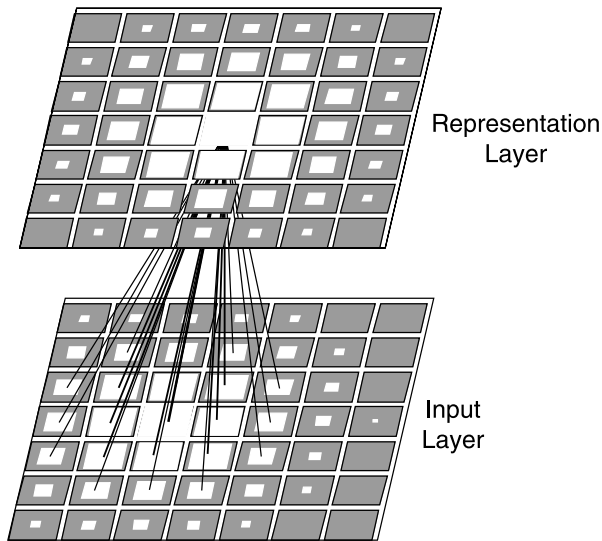


Figure 1 Model used to explore critical periods in phoneme perception (adapted from McClelland et al., 1999).

Application: critical periods

Why are there critical periods (or sensitive periods) in development, when a system seems to be much more responsive to environmental input than at later ages? For example, people are generally much better at learning language (and in particular, syntax) before puberty than after (Johnson & Newport, 1989; Newport, 1990). Also, infants learn the phonemes they are exposed to within the first months of life, while older children and adults have more difficulty with this task (Kuhl, Williams, Lacerda, Stevens & Lindblom, 1992; Werker & Tees, 1983). Such critical periods are not limited to language, nor to humans. For example, they are also observed in the effects of visual deprivation in cats (Hubel & Weisel, 1970), macaques (Horton & Hocking, 1997) and humans (Maurer & Lewis, 2001), and in the imprinting behavior of chicks (Bolhuis & Bateson, 1990).

Hebbian mechanisms may play an important role in such critical periods. As connections change to strengthen certain neural responses to stimuli, this can reduce the possibility of other responses, which can impair learning.

Here, we focus on a model demonstrating how such Hebbian mechanisms can lead 'young' networks to develop appropriate representations for phonemes in the environment, while 'old' networks fail to do so (McClelland et al., 1999; see also Guenther & Gjaja, 1996). This model (Figure 1) uses a variant of the Kohonen network (Kohonen, 1982) and the normalized Hebbian learning

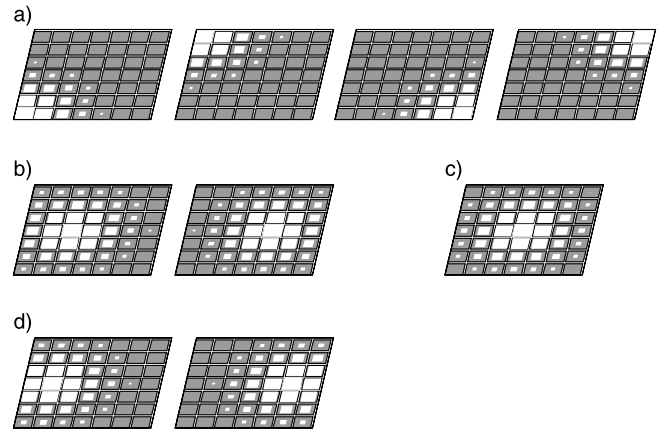


Figure 2 Sample input patterns presented to the phoneme model (adapted from McClelland et al., 1999): (a) distinct phonemes, (b) overlapping phonemes, (c) a single central phoneme and (d) exaggerated phonemes.

algorithm (Equation 3) described above. Networks were presented with phonemes represented as patterns of activity on the input layer. Some patterns were well separated (Figure 2a), corresponding to easily distinguished phonemes. Other patterns were overlapping (Figure 2b), corresponding to difficult to distinguish phonemes like /l/ and /r/. And other patterns were blended versions of overlapping patterns (Figure 2c), corresponding to single phonemes like the Japanese phoneme that includes /l/ and /r/ sounds. One group of networks was 'raised' in an English-language environment, with the well-separated and overlapping patterns. A second group of networks was raised in a Japanese-language environment, with the well-separated and blended patterns, and was later placed in the English-language environment. The networks demonstrated a critical period in learning to distinguish overlapping input patterns (like /r/ and /l/). Young models represented such inputs as distinct patterns on the representation layer, while older models instead represented such inputs as a single pattern on the representation layer. This corresponds to young English learners hearing /r/ and /l/ as distinct sounds, while older English learners (in particular, those who learned Japanese as their first language) cannot hear the difference between /r/ and /l/ sounds.

How did young networks learn to appropriately represent the phonemes in their environments, and why did older networks fail to do so? Different inputs tended to activate different combinations of perceptual units. Hebbian learning strengthened connections between input units and the perceptual units they activated, leading networks to appropriately represent the phonemes in their 'native' environments. For networks raised with

separate /l/ and /r/ inputs, this led to separate perceptual representations for these two inputs. For networks raised with a single blended sound, Hebbian learning led to strong connections from relevant input units to the single perceptual representation. When these networks were then exposed to separate, but overlapping, /l/ and /r/ sounds, these inputs tended to activate the single perceptual representation. Hebbian learning served to maintain this tendency to represent the distinct inputs as a single sound, by strengthening the connections between the input units representing the distinct sounds and the shared perceptual representation units. As a result, even with extended training on the separate /l/ and /r/ sounds, the networks that first learned the single blended sound were unable to activate distinct perceptual representations for the separate input sounds.

Note that the critical periods in the model were not determined by any maturational changes. That is, the fundamental mechanism of plasticity (i.e. Hebbian learning) was intact as networks 'aged'. Older models failed to learn distinct phonemes because they learned based on inappropriate (blended) perceptual representations, not because their learning mechanism was impaired.

This Hebbian framework for critical periods thus suggested a way to improve learning of overlapping phonemes after the critical period (McCandliss, Fiez, Protopapas, Conway & McClelland, 2002; see also Tallal, Miller, Bedi, Byma, Wang, Nagaraja, Schreiner, Jenkins & Merzenich, 1996). Exaggerated /l/ and /r/ sounds (Figure 2d) might elicit distinct perceptual representations. Hebbian learning could then strengthen the connections supporting such distinct representations, increasing the likelihood that /l/ and /r/ inputs would activate distinct perceptual representations. The degree of exaggeration of the /l/ and /r/ sounds could then be reduced gradually, maintaining distinct representations for the two inputs until they fell within the natural range. In contrast, simple repeated exposure to overlapping stimuli should not improve performance, because Hebbian learning would only strengthen the tendency to hear those phonemes as a single sound. These predictions from the Hebbian model about how to improve performance have been confirmed in adults who learned Japanese as their first language (McCandliss *et al.*, 2002).

Hebbian mechanisms can account for critical periods in the imprinting behavior of chicks in a similar way (O'Reilly & Johnson, 2002). Chicks tend to imprint on the first object they see after hatching, preferring this object over other objects (Bolhuis & Bateson, 1990). A Hebbian model simulated this critical period, based on units becoming recruited and tuned to an imprinting stimulus through repeated exposure to it. After sufficient

exposure, a majority of units became selective for the imprinting stimulus, and no amount of exposure to a different stimulus could shift this preference (O'Reilly & Johnson, 2002). The model also simulated the sensitivity of critical periods to factors such as the length of a chick's exposure to the first object experienced and the length of subsequent exposure to a second object. The longer the network's exposure to an object, the more units that became tuned to the object based on Hebbian learning, and the more difficult it became for those units to represent other objects.

Although Hebbian models have thus simulated a variety of aspects of critical periods, Hebbian mechanisms alone may not be sufficient to explain the full range of behaviors observed in critical periods. For example, error feedback can significantly improve performance, which Hebbian mechanisms cannot readily simulate. Japanese speakers improved their discrimination of /r/ and /l/ sounds when they were provided with error feedback, with training on either exaggerated or typical /l/ and /r/ stimuli (McCandliss *et al.*, 2002). Such findings suggest the importance of error-driven learning mechanisms. In fact, error-driven learning models have been used to investigate critical period effects, in word learning, reading and grammar (Ellis & Lambon-Ralph, 2000; Elman, 1993; Marchman, 1994; Rohde & Plaut, 1999; Zevin & Seidenberg, 2002). In addition, although the current model demonstrates that critical periods may arise without maturational changes, maturational changes may contribute and their effects on learning have been explored in models (e.g. Shrager & Johnson, 1996). However, it is important to note that apparently maturational changes can be triggered by learning mechanisms (Munakata & McClelland, 2003). For example, the NMDA receptor, crucial for changes in synaptic efficacy, decreases in responsiveness across development (Carmignoto & Vicini, 1992). Such changes in responsiveness, however, are experience-dependent and can be reversed if experience is withheld (Quinlan, Olstein & Bear, 1999).

Interestingly, the same Hebbian mechanisms that lead to critical periods could lead a system to get 'stuck' on a much smaller time scale. In particular, infants and children often *perseverate*, repeating prior behaviors when they are no longer appropriate (Aguiar & Baillargeon, 2000; Piaget, 1954; Zelazo, Frye & Rapus, 1996; Morton & Munakata, 2002b). For example, after infants search for a toy that is presented and then hidden, they tend to perseverate in searching for the toy in that location, even after watching as the toy is hidden in a new location (Diamond, 1985; Piaget, 1954). Similarly, after 3-year-olds sort cards according to one rule (e.g. by their shape), they tend to perseverate with this rule even after

they are told they should sort by a new rule (e.g. according to color; Zelazo *et al.*, 1996). Neural network models have demonstrated how such perseveration could arise from Hebbian learning mechanisms, because they tend to strengthen whatever response is made to a stimulus (Munakata, 1998; Morton & Munakata, 2002a; Stedron, Munakata & Sahni, 2002). Relative to the learning that occurs during critical periods, these perseverative tendencies might be easier to overcome because they result from a relatively small re-weighting of representations rather than a large-scale structuring of representations. Again, Hebbian mechanisms are not the only way to account for such behaviors (e.g. Dehaene & Changeux, 1989). Hebbian mechanisms may represent one of the most basic ways to account for such behaviors, given that feedback can seem to have no effect on children's perseveration (Diamond, 1983; Smith, Thelen, Titzer & McLin, 1999). However, error-driven mechanisms may also be required to account for cases where children respond to feedback by perseverating less (Yerys & Munakata, 2002).

Conclusions

In sum, Hebbian learning algorithms are biologically plausible, ecologically valid and powerful enough to account for a range of behaviors across development. At the same time, a complete account will likely require a more powerful error-driven learning mechanism as well. The most well-known option, backpropagation, is not ideal (Schlesinger & Parisi, this issue). However, it is important not to throw out the error-driven learning baby with the backpropagation bathwater. More biologically plausible error-driven learning mechanisms have been developed since backprop (e.g. Hinton & McClelland, 1988; O'Reilly, 1996). These algorithms estimate error signals from locally available activation values, with resulting weight changes computed in part based on the simultaneous activations of units, as in Hebbian learning. And the ecological validity of teaching signals in error-driven learning algorithms has been addressed to some degree (e.g. McClelland, 1994). Thus, Hebbian and error-driven algorithms may be integrated into a unified framework (O'Reilly, 1998, 2001). This type of unified framework can be more general than either Hebbian or error-driving learning algorithms alone, in simulating a range of behaviors across perception, attention, memory, language and higher-level cognition (O'Reilly & Munakata, 2000). Further integration of learning algorithms, capitalizing on the strengths of each in a biologically plausible manner, will likely play a key role in future progress in the modeling of development.

Acknowledgements

Preparation of this article was supported by a grant from NICHD (HD37163). We thank Randy O'Reilly for useful comments.

References

- Aguiar, A., & Baillargeon, R. (2000). Perseveration and problem solving in infancy. In H. Reese (Ed.), *Advances in child development and behavior*, Vol. 27 (pp. 135–180). New York: Academic Press.
- Altmann, G.T.M., & Dienes, Z. (1999). Rule learning by seven-month-old infants and neural networks. *Science*, **284**, 875.
- Artola, A., Brocher, S., & Singer, W. (1990). Different voltage-dependent thresholds for inducing long-term depression and long-term potentiation in slices of rat visual cortex. *Nature*, **347**, 69–72.
- Aslin, R.N., Saffran, J.R., & Newport, E.L. (1998). Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, **9**, 321–324.
- Bear, M.F., & Malenka, R.C. (1994). Synaptic plasticity: LTP and LTD. *Current Opinion in Neurobiology*, **4**, 389–399.
- 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** (2), 32–48.
- Blair, B.S., Intrator, N., Shouval, H., & Cooper, L.N. (1998). Receptive field formation in natural scene environments: comparison of single-cell learning rules. *Neural Computation*, **10**, 1797–1813.
- Blasdel, G.G., & Salama, G. (1986). Voltage-sensitive dyes reveal a modular organization in monkey striate cortex. *Nature*, **321**, 579–585.
- Bliss, T.V.P., & Lomo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology (London)*, **232**, 331–356.
- Bolhuis, J.J., & Bateson, P.P.G. (1990). The importance of being first: a primacy effect in filial imprinting. *Animal Behavior*, **40**, 472–483.
- Carmignoto, G., & Vicini, S. (1992). Activity-dependent decrease in NMDA receptor responses during development of the visual-cortex. *Science*, **258**, 1007–1011.
- Christiansen, M.H., & Curtin, S. (1999). Transfer of learning: rule acquisition or statistical learning? *Trends in Cognitive Sciences*, **3**, 289–290.
- Dehaene, S., & Changeux, J.-P. (1989). A simple model of prefrontal cortex function in delayed-response tasks. *Journal of Cognitive Neuroscience*, **1** (3), 244–261.
- Diamond, A. (1983). Behavior changes between 6–12 months of age: what can they tell us about how the mind of the infant is changing? *Dissertation Abstracts International*, **44** (01B), 337 (University Microfilms No. AAD8311882).
- Diamond, A. (1985). Development of the ability to use recall

- to guide action, as indicated by infants' performance on AB. *Child Development*, **56**, 868–883.
- Ellis, A.W., & Lambon-Ralph, M. (2000). Age of acquisition effects in adult lexical processing reflect loss of plasticity in maturing systems: insights from connectionist networks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **26**, 1103–1123.
- Elman, J.L. (1993). Learning and development in neural networks: the importance of starting small. *Cognition*, **48** (1), 71–99.
- Fiser, J., & Aslin, R.N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, **12**, 499–504.
- Guenther, F.H., & Gjaja, M.N. (1996). The perceptual magnet effect as an emergent property of neural map formation. *Journal of the Acoustical Society of America*, **100**, 1111–1121.
- Hauser, M.D., Newport, E.L., & Aslin, R.N. (2001). Segmentation of the speech stream in a nonhuman primate: statistical learning in cotton top tamarins. *Cognition*, **78**, B54–B64.
- Hebb, D.O. (1949). *The organization of behavior*. New York: Wiley.
- Hinton, G.E., & McClelland, J.L. (1988). Learning representations by recirculation. In D.Z. Anderson (Ed.), *Neural information processing systems, 1987* (pp. 358–366). New York: American Institute of Physics.
- Horton, J.C., & Hocking, D.R. (1997). Timing of the critical period for plasticity of ocular dominance columns in macaque striate cortex. *Journal of Neuroscience*, **17**, 3684–3709.
- Hubel, D.H., & Weisel, T. (1970). The period of susceptibility to the physiological effects of unilateral eye closure in kittens. *Journal of Physiology (London)*, **206**, 419–436.
- Johnson, J., & Newport, E. (1989). Critical period effects in second-language learning: the influence of maturational state on the acquisition of English as a second language. *Cognitive Psychology*, **21**, 60–99.
- Kirkham, N.Z., Slemmer, J.A., & Johnson, S.P. (2002). Visual statistical learning in infancy: evidence for a domain general learning mechanism. *Cognition*, **83**, B35–B42.
- Kohonen, T. (1982). Clustering, taxonomy, and topological maps of patterns. In M. Lang (Ed.), *Proceedings of the 6th International Conference on Pattern Recognition* (pp. 114–128). Silver Spring, MD: IEEE Computer Society Press.
- Kohonen, T. (1984). *Self-organization and associative memory*. Berlin: Springer-Verlag.
- Konen, W.K., & Von Der Malsburg, C. (1994). A fast dynamic link matching algorithm for invariant pattern recognition. *Neural Networks*, **7**, 1019–1030.
- Kuhl, P.K., Williams, K., Lacerda, F., Stevens, K., & Lindblom, B. (1992). Linguistic experience alters phonetic perception in infants by 6 months of age. *Science*, **255**, 606–608.
- Landauer, T.K., & Dumais, S.T. (1997). A solution to Plato's problem: the latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, **104**, 211–240.
- Lee, D.D., & Seung, H.S. (1999). Learning the parts of objects by non-negative factorization. *Nature*, **401**, 788.
- Lisman, J. (1994). The CaM Kinase II hypothesis for the storage of synaptic memory. *Trends in Neurosciences*, **17**, 406.
- Livingstone, M., & Hubel, D. (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science*, **24**, 740–749.
- McCandliss, B.D., Fiez, J.A., Protopapas, A., Conway, M., & McClelland, J.L. (2002). Success and failure in teaching the [r]–[l] contrast to Japanese adults: tests of a Hebbian model of plasticity and stabilization in spoken language perception. *Cognitive, Affective, and Behavioral Neuroscience*, **2**, 89–108.
- McClelland, J.L. (1994). The interaction of nature and nurture in development: a parallel distributed processing perspective. In P. Bertelson, P. Eelen & G. D'Ydewalle (Eds.), *Current advances in psychological science: Ongoing research* (pp. 57–88). Hillsdale, NJ: Lawrence Erlbaum Associates.
- McClelland, J.L., Thomas, A., McCandliss, B.D., & Fiez, J.A. (1999). Understanding failures of learning: Hebbian learning, competition for representational space, and some preliminary experimental data. In J. Reggia, E. Ruppini & D. Glanzman (Eds.), *Brain, behavioral, and cognitive disorders: The neurocomputational perspective* (pp. 75–80). Oxford: Elsevier.
- Marchman, V.A. (1994). Constraints on plasticity in a connectionist model of the English past tense. *Journal of Cognitive Neuroscience*, **5**, 215–234.
- Maurer, D., & Lewis, T.L. (2001). Visual acuity: the role of visual input in inducing postnatal change. *Clinical Neuroscience Research*, **1**, 239–247.
- Maye, J., Werker, J.F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition*, **82**, B101–B111.
- Merzenich, M.M., & Kaas, J.H. (1980). *Principles of organization of sensory-perceptual systems of mammals*. New York: Academic Press.
- Miller, K.D. (1994). A model for the development of simple cell receptive fields and the ordered arrangement of orientation columns through activity-dependent competition between ON- and OFF-center inputs. *Journal of Neuroscience*, **14**, 409–441.
- Morton, J.B., & Munakata, Y. (2002a). Active versus latent representations: a neural network model of perseveration and dissociation in early childhood. *Developmental Psychobiology*, **40**, 255–265.
- Morton, J.B., & Munakata, Y. (2002b). Are you listening? Exploring a knowledge-action dissociation in a speech interpretation task. *Developmental Science*, **5**, 435–440.
- Munakata, Y. (1998). Infant perseveration and implications for object permanence theories: a PDP model of the AB task. *Developmental Science*, **1**, 161–184.
- Munakata, Y., & McClelland, J.L. (2003). Connectionist models of development. *Developmental Science*, **6** (4), 413–429.
- Munakata, Y., McClelland, J.L., Johnson, M.H., & Siegler, R. (1997). Rethinking infant knowledge: toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, **104**, 686–713.
- Newport, E.L. (1990). Maturation constraints on language learning. *Cognitive Science*, **14**, 11–28.

- Oja, E. (1982). A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, **15**, 267–273.
- O'Reilly, R.C. (1996). Biologically plausible error-driven learning using local activation differences: the generalized recirculation algorithm. *Neural Computation*, **8** (5), 895–938.
- O'Reilly, R.C. (1998). Six principles for biologically-based computational models of cortical cognition. *Trends in Cognitive Sciences*, **2** (11), 455–462.
- O'Reilly, R.C. (2001). Generalization in interactive networks: the benefits of inhibitory competition and Hebbian learning. *Neural Computation*, **13**, 1199–1242.
- O'Reilly, R.C., & Johnson, M.H. (2002). Object recognition and sensitive periods: a computational analysis of visual imprinting. In M.H. Johnson, R.O. Gilmore & Y. Munakata (Eds.), *Brain development and cognition: A reader* (2nd edn., pp. 392–413). Oxford: Blackwell.
- O'Reilly, R.C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, MA: MIT Press.
- Piaget, J. (1954). *The construction of reality in the child*. New York: Basic Books.
- Quinlan, E.M., Olstein, D.H., & Bear, M.F. (1999). Bidirectional, experience-dependent regulation of n-methyl-d-aspartate receptor subunit composition in the rat visual cortex during postnatal development. *Neurobiology*, **96**, 12876–12880.
- Rogers, T.T., & McClelland, J.L. (in press). A parallel distributed processing approach to semantic cognition: applications to conceptual development. In D. Rakison & L. Gershkoff-Stowe (Eds.), *Building object categories in developmental time: Proceedings of the Carnegie Symposium on Cognition*, Vol. 32. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Rohde, D., & Plaut, D. (1999). Language acquisition in the absence of explicit negative evidence: how important is starting small? *Cognition*, **72**, 67–109.
- Rumelhart, D.E., & Zipser, D. (1986). Feature discovery by competitive learning. In D.E. Rumelhart, J.L. McClelland & PDP Research Group (Eds.), *Parallel distributed processing. volume 1: Foundations* (Chap. 5, pp. 151–193). Cambridge, MA: MIT Press.
- Saffran, J.R., Johnson, E.K., Aslin, R.N., & Newport, E.L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, **70**, 27–52.
- Seidenberg, M.S., & Elman, J.L. (1999). Do infants learn grammar with algebra or statistics? *Science*, **284**, 435–436.
- Shrager, J., & Johnson, M.H. (1996). Dynamic plasticity influences the emergence of function in a simple cortical array. *Neural Networks*, **9**, 1119.
- Smith, L.B., Thelen, E., Titzer, B., & McLin, D. (1999). Knowing in the context of acting: the task dynamics of the A-not-B error. *Psychological Review*, **106**, 235–260.
- Stedron, J., Munakata, Y., & Sahni, S.D. (2002). In plain sight: simulating the role of memory in perseveration with visible solutions. Manuscript in preparation.
- Tallal, P., Miller, S.L., Bedi, G., Byma, G., Wang, X., Nagaraja, S.S., Schreiner, C., Jenkins, W.M., & Merzenich, M.M. (1996). Language comprehension in language-learning impaired children improved with acoustically modified speech. *Science*, **271**, 81–84.
- Werker, J.F., & Tees, R.C. (1983). Developmental changes in the perception of non-native speech sounds. *Canadian Journal of Psychology*, **37**, 278–386.
- Yerys, B.E., & Munakata, Y. (2002). The effect of feedback on children's flexibility: rethinking perseveration in a card sorting task. Manuscript in preparation.
- Zelazo, P.D., Frye, D., & Rapus, T. (1996). An age-related dissociation between knowing rules and using them. *Cognitive Development*, **11**, 37–63.
- Zevin, J., & Seidenberg, M. (2002). Age of acquisition effects in reading and other tasks. *Journal of Memory & Language*, **47**, 1–29.