

Six Principles for Biologically-Based Computational Models of Cortical Cognition

Randall C. O'Reilly

Department of Psychology
University of Colorado at Boulder
Campus Box 345
Boulder, CO 80309
oreilly@psych.colorado.edu

January 19, 1999

Appeared in 1998: *Trends in Cognitive Sciences*, 2, 455-462.

Abstract

This paper describes and motivates six principles for computational cognitive neuroscience models: biological realism, distributed representations, inhibitory competition, bidirectional activation propagation, error-driven task learning, and Hebbian model learning. Although these principles are supported by a number of cognitive, computational, and biological motivations, the prototypical neural network model (a feedforward backpropagation network) incorporates only two of them, and no widely used model incorporates all of them. This paper argues that these principles should be integrated into a coherent overall framework, and discusses some potential synergies and conflicts in doing so.

A number of important principles have been developed for computational neural network models of cortical learning and cognitive processing. However, relatively little work has been done to try to integrate these principles into a coherent overall framework. Integrating these principles allows one to demonstrate the consistency of different models, capitalize on synergies between different principles, organize and consolidate existing findings, and generate novel insights into the nature of cognition. This paper describes and motivates a provisional set of six principles (illustrated in Figure 1) that have proven individually useful in a number of existing models. Despite their proven utility, most models incorporate only a small number of these principles (e.g., the prototypical feedforward backpropagation network uses only two). Thus, this paper attempts to highlight the potential advantages and pitfalls of using a more inclusive set of principles.

Although a specific algorithm called *Leabra* has been developed to incorporate these principles (see Box 1), the focus of this paper is on the history and importance of the principles themselves, and the ways in which these principles interact with each other. As an important caveat, this discussion focuses on biologically-based principles that are particularly relevant for cognition, and does not include a number of functional and cognitive-level principles that could also be enumerated.

The proposed set of principles can be considered an extension of the GRAIN framework of McClelland [1]. GRAIN stands for graded, random, adaptive, interactive, (nonlinear) network. This framework was primarily motivated by (and applied to) issues surrounding the dynamics of activation flow through a neural network. By way of extension, the present framework emphasizes learning mechanisms and the architectural properties that support them. Two of the key principles in GRAIN, interactivity and competition, are among the six principles emphasized here. The other GRAIN principles (graded, nonlinear activations, graded activation propagation, and intrinsic variability) are assumed, but not emphasized in this framework because of their nearly ubiquitous acceptance within neural network models (but see [2] for an interesting application of these principles to controversies in cognitive development).

The Principles

The six principles can be grouped into three categories. The first principle, biological realism, is in a category by itself, providing a general overriding constraint on the framework. The next three principles, distributed representations, inhibitory competition, and bidirectional activation propagation (interactivity), are concerned with the architecture of the network and the general behavior of the neuron-like processing units within it. The final two principles, error-driven task learning and Hebbian model learning, govern the way that learning occurs in the network.

1. Biological Realism

Biological realism lies at the foundation of the entire enterprise of computational modeling in cognitive neuroscience. This approach seeks to understand how the *brain* (and specifically the cortex in the present case) gives rise to cognition, not how some abstraction of uncertain validity does so. Thus, wherever possible, computational models should be constrained and informed by biological properties of the cortex. Moreover, computational mechanisms that violate known biological properties should not be relied upon. This point has implications for error-driven learning, as discussed below.

Although the issue of biological realism is easy to state, it can be difficult to apply, because the known biology often does not provide sufficient constraints. Thus, biological realism often reduces to plausibility arguments, which depend on things like how simple and local the mechanism in question is, and that it is not inconsistent with known biology. Also, one can adopt a converging evidence approach, where multiple constraints from biology, computation, and cognition converge to support a given principle. This approach is emphasized here.

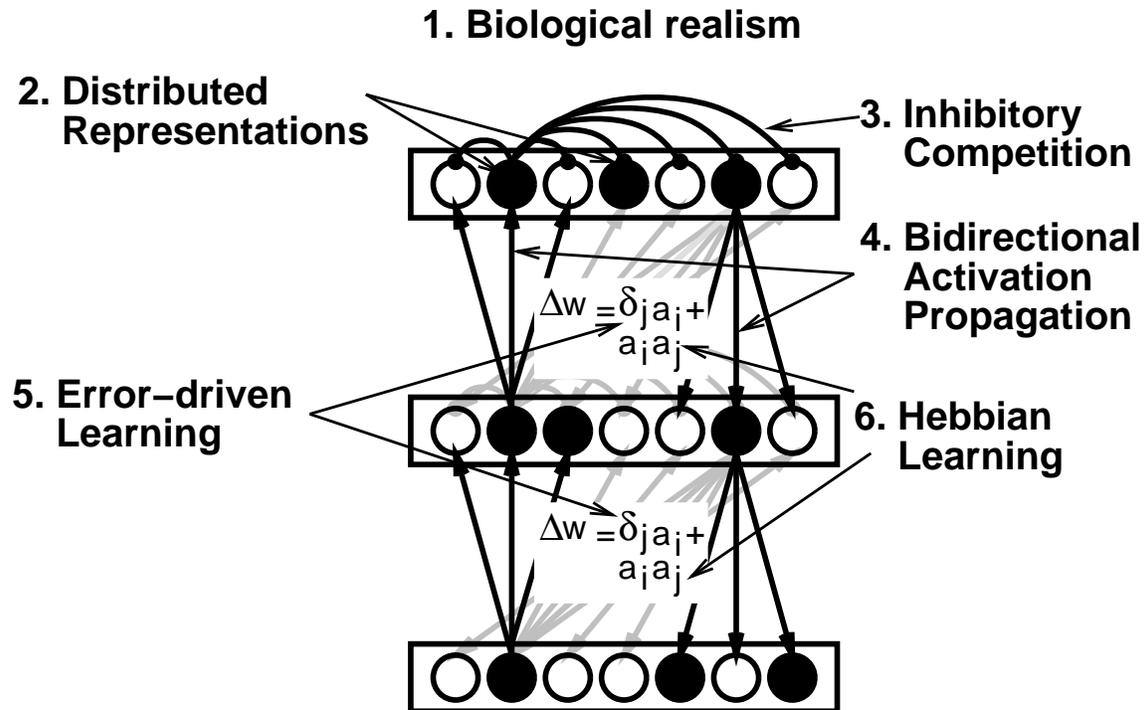


Figure 1: Illustration of the six core principles, and their instantiation in a neural network. Biological realism (1) is an overarching constraint. Distributed representations (2) have multiple units active, while inhibitory competition (3, implemented via inhibitory connectivity) ensures that relatively few such units are active. Bidirectional activation propagation (4, implemented by bidirectional connectivity) enables both bottom-up and top-down constraints to simultaneously shape the internal representation. Error-driven learning (5) shapes representations according to differences between expected outputs and actual ones (represented by the error term δ_j). Hebbian learning (6) shapes representations according to the co-occurrence (correlation) statistics of items in the environment (represented by the product of the sending and receiving unit activations).

Architectural Principles

2. Distributed Representations

The cortex is widely believed to use distributed representations to encode information. A distributed representation uses multiple active neuron-like processing units to encode information (as opposed to a single unit, localist representation), and the same unit can participate in multiple representations. Each unit in a distributed representation can be thought of as representing a single *feature*, with information being encoded by particular combinations of such features. Electrophysiological recordings demonstrate that distributed representations are widely used in the cortex (e.g., [3, 4, 5]). The functional benefits of distributed representations include greater efficiency, robustness, and accuracy, and the ability to represent similarity relationships [6]. The efficiency of distributed representations can be appreciated by analogy with letters. Just as different combinations of a small number of letters can represent a large number of words, so can different combinations of a small set of units represent a large amount of information. The robustness of distributed representations comes from the redundancy of having each item represented by many units. Distributed representations can more accurately represent graded values through *coarse coding*, where a value is encoded by the relative magnitudes of a number of broadly tuned units. Finally, similarity is represented by the shared units involved in the distributed representations of different items.

3. Inhibitory Competition

Inhibitory competition is important because it selects representations for processing and for subsequent refinement over learning. Inhibitory competition arises when mutual inhibition among a set of units (i.e., as mediated by inhibitory interneurons) prevents all but a subset of them from becoming active at a time. Roughly 20% of the neurons in the cortex are inhibitory interneurons [7], and it is clear that they control the explosion of activation that would otherwise result from all the positive interconnectivity among cortical pyramidal neurons (e.g., as happens in epilepsy). Inhibitory competition allows only the most strongly excited representations to prevail, with this *selection* process identifying the most appropriate representations for subsequent processing. Furthermore, most learning mechanisms (including those discussed later) are affected by this selection process such that only the selected representations are refined over time through learning, resulting in an effective differentiation and distribution of representations [8, 9, 10, 11].

Aside from the selection and refinement of representations, another benefit of inhibitory competition comes from the idea that, given the general structure of the environment, *sparse* distributed representations (i.e., having relatively few units active at a time) are particularly useful [12, 13]. For example, in visual processing, a given object can be defined along a set of feature dimensions (e.g., shape, size, color, texture), with a large number of different values along each dimension (i.e., many different possible shapes, sizes, colors, textures, etc). Assuming that the individual units in a distributed representation encode these feature values, a representation of a given object will only activate a small subset of units (i.e., the representations will be sparse). To further substantiate this argument, Olshausen & Field [14] showed that imposing a bias for developing sparse distributed representations can result in the development of realistic early visual representations (oriented edge detectors) of natural visual scenes. More generally, it seems as though the world can be usefully represented in terms of a large number of categories with a large number of exemplars per category (animals, furniture, trees, etc.). If we again assume that only a relatively few such exemplars are processed at a given time, a bias favoring sparse representations is appropriate.

4. Bidirectional Activation Propagation (Interactivity)

Bidirectional activation propagation is a critical principle for information flow through the network. Bidirectional activation propagation (also called *interactivity* or *recurrence*) is the communication of activation simultaneously in both bottom-up and top-down directions. This contrasts with *feedforward* activation propagation where information only goes in one direction (bottom-up). To enable information to flow in both directions simultaneously in a stable and effective manner, processing must proceed in gradual, iterative steps. Thus, a temporally-extended *settling* process with many iterative steps is required for the network to achieve an appropriate representation of a given input pattern. This is a central feature of GRAIN [1]. Bidirectional connectivity is ubiquitous in the cortex (e.g., [15, 16, 17]). An important benefit of bidirectional activation propagation is powerful *constraint satisfaction* processing [18, 19], where both lower-level (e.g., perceptual) and higher-level (e.g., conceptual) constraints can be simultaneously brought to bear in interpreting and processing inputs. The importance of interactivity for understanding cognitive processing was demonstrated in the word superiority model of McClelland & Rumelhart [9]. They showed that interactivity could explain the counterintuitive finding that higher-level word processing can influence lower-level letter perception. More recently, Vecera & O'Reilly [20] showed that bidirectional constraint satisfaction can model people's ability to resolve ambiguous visual inputs in favor of familiar versus novel objects [21]. They also showed that adding inhibitory competition to an interactive network significantly speeded the settling process, and greatly reduced the number of times the network settled into bad local minima.

Learning Principles

Learning is essential for shaping the representations of neural networks according to the structure of the environment. A key issue is what aspects of the environmental structure should be learned, with the understanding that not everything can or should be represented. The following two learning principles exploit two complementary aspects of environmental structure: task demands, and the extent to which different things co-occur. The first is referred to as *task learning* for obvious reasons, and the second is referred to as *model learning* because the objective is to develop an internal model of the environment irrespective of specific tasks. These two learning objectives can be achieved by two different forms of implementational mechanisms, *error-driven* and *Hebbian* learning, respectively.

5. Error-driven Task Learning

Error-driven learning (also called supervised learning) is important for shaping representations according to task demands by learning to minimize the difference (i.e., the error) between a desired outcome and what the network actually produced. This principle captures the idea that you learn what enables you to succeed at the necessary tasks of life, without bothering to represent aspects of the environment that are not relevant to these tasks. The widely used backpropagation learning algorithm [22] directly minimizes error through gradient descent, and has proven to be very powerful. Although task learning is clearly psychologically relevant, and a majority of psychological models have used this form of learning, its biological plausibility has been widely questioned because it requires the propagation of error signals in a manner inconsistent with known neurobiological properties (e.g., [23, 24]). Furthermore, it has not been clear where the necessary “teaching” signals could plausibly come from. However, it has recently been shown that biologically plausible bidirectional activation propagation (see principle 4) can be used to perform essentially the same error-driven learning as backpropagation [25], using any of a number of readily available teaching signals. The resulting algorithm generalizes the recirculation algorithm of Hinton & McClelland [26], and is thus called *GeneRec*.

The basic idea behind GeneRec is that instead of propagating an error signal, which is a difference between two terms, one can propagate the two terms separately as activation signals, and then take their difference locally at each unit. This works by having two phases of activations for computing the two terms. In the *expectation* phase, the bidirectionally-connected network settles based on an input activation pattern into a state that reflects the expected consequences or correlates of that input pattern. Then, in the *outcome* phase, the network experiences actual consequence(s) or correlate(s). The difference between outcome and expectation is the error signal, and the bidirectional connectivity propagates this error signal throughout the network via local activation signals. Thus, interactivity enables units everywhere in the network to receive (possibly indirectly via hidden layers) activation signals from the layer(s) where the expectation and outcome are represented. The remarkable thing is that the activation signals in an interactive network are naturally propagated (even through hidden layers) in just the right way to enable the correct error gradient to be simply and locally computed at each unit [25].

The GeneRec analysis also showed that Boltzmann machine learning and its deterministic versions [19, 27, 28, 29] can be seen as variants of this more biologically plausible version of the backpropagation algorithm. This means that all of the existing approaches to error-driven learning using activation-based signals converge on essentially the same basic mechanism, making it more plausible that this is the way the brain does error driven learning. Furthermore, the form of synaptic modification necessary to implement this algorithm is consistent with (though not directly validated by) known properties of biological synaptic modification mechanisms. Finally, there are many sources in the natural environment for the necessary outcome phase signals in the form of actual environmental outcomes that can be compared with internal expectations to provide error signals [30, 25]. Thus, one does not need to have an explicit “teacher” to perform

error-driven learning. Taken together, these developments make it difficult to continue to object to the use of error-driven learning on the grounds that it is not biologically plausible.

6. Hebbian Model Learning

Model learning (also called self-organizing or unsupervised learning) is important for forming internal representations of the general (statistical) structure of the environment, without respect to particular tasks. Many versions of this general idea exist, defined by what aspects of environmental structure are deemed important to represent, but it is generally agreed that something like correlational structure is important. Hebbian learning mechanisms represent this correlational structure, encoding the extent to which different things co-occur in the environment [31]. Biologically, Hebbian learning requires that the synaptic strength change as a function of the co-activation of the sending and receiving neurons. NMDA-mediated long-term potentiation (LTP) has this Hebbian property (e.g., [32]). Thus, Hebbian learning is almost universally regarded as being biologically plausible. At a functional level, the co-occurrence of items suggests that there might be a causal relationship between them. Furthermore, co-occurring items can be more efficiently represented together within a common representational structure. Mathematical analyses have shown that Hebbian learning performs something like principal components analysis [33], which extracts the principal dimensions of covariance within the environment. An interesting demonstration of the power of this kind of Hebbian model learning was recently provided in the form of a model that performs principal components analysis on the co-occurrence statistics of words within large texts, yielding surprisingly powerful representations of word meaning [34].

Interactions Among the Principles

The preceding discussion provided specific and compelling motivations for each of the individual principles. In this section, three examples of interactions (synergies and conflicts) among the six principles will be discussed. The first example comes from the GRAIN framework, and deals with the consequences of interactivity and noise. The second explores the interactions between distributed representations and competition, which can be at odds with each other. The last explores the interactions between error-driven and Hebbian learning.

Interactivity and Noise

The GRAIN framework has been used to explore the implications of some of the principles on the activation dynamics of a network [1]. For example, although interactivity (bidirectional activation propagation) is important for accounting for a range of different behavioral phenomena, it can also be problematic for others. Specifically, interactivity interfered with the ability of a network to exhibit independent contributions from context and stimulus strength in a stimulus identification situation [35, 36]. McClelland showed that the use of intrinsic variability (noise) can resolve this conflict, resulting in a model that captures a wider range of phenomena, including standard interactive phenomena (e.g., top-down effects) and the independent contributions of context and stimuli [1].

This example illustrates a theme that emerges repeatedly when attempting to integrate different principles (see Box 1 for another example): often, subsets of principles do not work as well as a more complete set of principles. Thus, instead of abandoning a principle (e.g., interactivity) when it appears to introduce a problem, one should consider how other principles might be adopted that would resolve the problem. The advantage of the integrative approach is that the resulting model then accounts for a much wider range of phenomena, and may provide important new insights into the nature of the originally problematic phenomenon. For example, the GRAIN model can explain the conditions under which you would *not* expect to

find an independent contribution of stimulus and context (see [1] for details and empirical validation).

Distributed Representations and Competition

Perhaps one of the most important challenges in integrating the six principles comes in combining distributed representations and competition, which tend to work at cross purposes. Distributed representations require multiple active units that cooperatively represent something, whereas competition causes fewer units to be active, and it can inhibit cooperativity. A reasonable compromise between these two principles is the sparse distributed representation as discussed previously. Although seemingly straightforward, achieving a sparse distributed representation is technically challenging, primarily because this case is difficult to analyze mathematically within a probabilistic learning framework. The problem is one of combinatorial explosion — one needs to take into account all the different possible combinations of active and inactive units to analyze a sparse distributed representation based on true inhibitory competition. Thus, sparse distributed representations fall in a complex intermediate zone between two easily analyzed frameworks [37]: (a) The winner-take-all (WTA) framework [10, 11, 38], where only one unit is allowed to be active at a time. Having a single active unit eliminates the combinatorial problems, but this also violates principle 2 by not allowing for distributed representations. (b) The independent units framework, where the units are considered to be (conditionally) independent of each other (e.g., a standard backpropagation network). This allows the combined probability of an activation pattern to be represented as a simple product of the individual unit probabilities (and for distributed representations), but it also violates principle 3 because there is no competition.

There have been a number of attempts to remedy the limitations of these two analytical frameworks, by introducing distributed representations within a basically WTA framework, or by introducing sparseness constraints within the independent units framework. However, the basic limitations of these frameworks are difficult to overcome. Basically, any use of WTA prevents the cooperativity and combinatoriality of true distributed representations, and the need to preserve independence among the units in the independent units framework prevents the introduction of any true activation-based competition. After discussing these approaches and their limitations, the more difficult to analyze approach of directly implementing sparse distributed representations using inhibitory competition will be discussed.

The following are extensions of the WTA framework. In the mixture-of-experts framework [39], a WTA competition takes place within a specialized “gating” network, that regulates the participation of a set of “expert” networks, which can themselves have distributed representations. A limitation of this approach is the coarse-grained level of the competition — whole groups of units compete, but individual units do not. Also, multiple experts cannot easily cooperate due to the WTA limitation. The model of Dayan & Zemel [40] uses a WTA assumption where units in the hidden layer compete to determine the value of a given unit in the output. However, this just transfers the WTA assumption from representing the input to representing the output, and a WTA assumption anywhere is likely to be problematic. The Dayan & Zemel model was intended as an improvement over the *noisy-or* model of Saund [41], which did not result in sufficient competition. Finally, the Kohonen network [8] uses a WTA to select a single winner, but then a neighborhood of units around that winner are also activated. Although useful for achieving topographic representations, this kind of fixed, imposed activation state does not enable the full combinatorial representational power that is an essential feature of true distributed representations.

Within the independent units framework, the main approach has been to introduce a sparseness constraint that does not actually involve direct activation-based competition. This usually involves adding an extra factor to the learning rule that favors sparse representations (e.g., [42, 14, 43, 44]), or adding a sparseness bias into the activation function itself (e.g., [37]). Thus, units are only competing over the long time-course of learning (or against their own negative bias), and not directly with one another to represent the current input pattern (i.e., selection). Furthermore, the dynamic thresholding behavior one achieves with activation-

based competition (which for example makes the system robust to changes in absolute levels of excitation) is not present in these approaches. This limitation is particularly evident in bidirectionally connected networks, where the need to control positive feedback requires the dynamic thresholding of true competition (O’Reilly, 1996, PhD Thesis, Carnegie Mellon University). Thus, integrating all of the principles places further demands on the competition mechanism.

It seems clear that the cortex implements inhibitory competition (and sparse distributed representations) via inhibitory interneurons. One way of understanding the effects of this inhibitory competition is in terms of a k -winners-take-all (KWTA) mechanism, that generalizes the WTA approach to k winners [45]. A KWTA mechanism can enforce true competition amongst the units, while allowing for a (sparse) distributed representation across the subset of k units. KWTA mechanisms have been analyzed for factors such as stability and convergence onto k units, and can be implemented with biologically-plausible lateral inhibition mechanisms [45, 46]. However, they have not been analytically treated within a probabilistic learning framework, due to the combinatorial explosion problems. Nevertheless, a simple form of KWTA that works well in bidirectionally connected networks has been shown to be useful for modeling a wide range of cognitive phenomena (see Box 1).

Learning Principles

Before discussing the interactions between error-driven task learning and Hebbian model learning, the distinction between the computational objectives of learning (i.e., task and model learning) and the implementational mechanisms (i.e., error-driven and Hebbian learning) needs to be clarified. Two points of potential confusion exist: (a) error-driven learning can be used to achieve a model-learning objective, and (b) some error-driven learning mechanisms resemble Hebbian mechanisms. The first point of confusion arises because one can train a network to reproduce the information in the environment using error-driven mechanisms, resulting in a task-independent model of the environment (i.e., via an auto-encoder [47, 48, 42] or a generative model [49, 14]). One can also learn an internal model based on error signals derived from the mismatch between different sensory representations of the same underlying event [50, 51, 52]¹ These examples raise the issue of why one should use Hebbian mechanisms to implement model learning, instead of using error-driven learning for both task and model learning. The subsequent discussion of the advantages of combining error-driven and Hebbian learning addresses this issue.

As for the second point of confusion, a version of the GeneRec algorithm analyzed in [25] is equivalent to the “contrastive Hebbian learning” (CHL) algorithm of Movellan [29], which uses the difference between two Hebbian terms. Also, other algorithms have been proposed that achieve quasi-error-driven learning with Hebbian-like mechanisms (e.g., [53]). However, despite these apparent similarities in the surface form of the learning rule, error-driven learning achieves a very different computational objective from simple Hebbian learning; only error-driven learning can achieve a fully general, powerful form of task learning (i.e., that is capable of learning arbitrary input-output mappings).

Thus, it seems clear that we should begin with the assumption that error-driven learning is essential for achieving task learning. From this error-driven perspective, it would then be of interest to know if further constraining the learning with Hebbian model learning would yield any benefits. A general framework for understanding why this might be the case was articulated by Geman, Bienenstock, & Doursat [54]. They argued that standard neural networks (e.g., generic backpropagation networks) are typically underconstrained by the learning task, and thus suffer from too much *variance* in solutions. This can have negative consequences for generalization to novel inputs, among other things. The solution is to add *biases* to networks that further constrain the learning by favoring particular forms of solutions (representations). To be beneficial,

¹This idea can also be viewed as an instance of the GeneRec expectation-outcome framework, where each modality creates an “expectation” about how the other modality will represent the event. The difference between this expectation and how the modality actually represents the event constitutes the error signal.

these biases obviously need to be appropriate — there are no generically optimal set of biases — but there may be a set of biases that are particularly appropriate for representing the real world. Indeed, encouraging sparse distributed representations can be seen as just such a bias that has been justified in terms of real world properties as discussed previously. It is likely, given the general importance of correlational information in the world (e.g., for suggesting causal relationships), that including a Hebbian bias towards representing co-occurrence statistics would be another such generally useful bias.

Although error-driven learning can be sensitive to correlational information, Hebbian learning is directly constrained to learn a correlation-based model because Hebbian weight changes directly reflect unit correlations. Thus, Hebbian model learning can provide a distinct and useful additional bias to further constrain error-driven task learning. This additional Hebbian bias can be thought of as a somewhat “smarter” version of the widely-used weight decay bias (e.g., [55]). Aside from the Leabra algorithm described in Box 1, there is at least one other example in the literature where error-driven (backpropagation) and Hebbian learning have been combined, with the expected beneficial results [56]. In addition to the synergy between these two forms of learning, combining both task-based and model-based learning enables one to account for phenomena associated specifically with these different types of learning. For example, it seems reasonable to assume that some kinds of learning occur as a result of mere exposure to stimuli (i.e., as would be expected by model learning, but not task learning). However, other kinds of learning (e.g., complex input-output mappings) clearly require task learning.

Outstanding Questions

- Are there cognitive phenomena or biological facts that appear to directly contradict the core principles?
- Is it possible that different parts of the cortex emphasize some principles over others? How might this influence functional specialization in the cortex?
- How many other important principles are missing from this list?
- Can complex sequential cognitive processing be shown to emerge from such basic principles as those discussed here, or does this require a whole new set of principles?
- How might error-driven and Hebbian learning co-exist and interact with reinforcement learning, which is likely to be taking place in sub-cortical structures, and possibly the cortex?

Acknowledgements

I thank Yuko Munakata and Jerry Rudy for their helpful comments. This work was supported in part by NIH program project grant MH47566.

Box 1: The Leabra Implementation

The six principles have been implemented in an algorithm called *Leabra*, which is briefly presented here. Leabra stands for “learning in an error-driven and associative, biologically-realistic algorithm” (where associative is another term for Hebbian learning). Leabra has been used in an upcoming textbook [57] to implement a wide range of cognitive neuroscience models. The scope of phenomena it is capable of modeling is commensurate with the breadth of the principles as discussed in the paper, and demonstrates their sufficiency and mutual compatibility.

Point Neuron Activation Function

Leabra uses a *point neuron* activation function that models the electrophysiological properties of real neurons, while simplifying their geometry to a single point. This is nearly as simple computationally as the standard sigmoidal activation function, but the more biologically-based implementation makes it considerably easier to model inhibitory competition, as described below. Further, it enables cognitive models to be more easily related to more physiologically detailed simulations, thereby facilitating bridge-building between biology and cognition.

The membrane potential V_m is updated as a function of ionic conductances g with reversal (driving) potentials E as follows:

$$\frac{dV_m(t)}{dt} = \tau \sum_c g_c(t) \overline{g}_c (E_c - V_m(t)) \quad (1)$$

with 3 channels (c) corresponding to: e excitatory input; l leak current; and i inhibitory input. The equilibrium potential can be written in a simplified form by setting the excitatory driving potential (E_e) to 1 and the leak and inhibitory driving potentials (E_l and E_i) of 0:

$$V_m^\infty = \frac{g_e \overline{g}_e}{g_e \overline{g}_e + g_l \overline{g}_l + g_i \overline{g}_i} \quad (2)$$

This shows that the neuron is computing a balance between excitation and the opposing forces of leak and inhibition. This form of the equation can be understood in terms of a Bayesian decision making framework [57]. Activation communicated to other cells (y) is a thresholded (Θ), sigmoidal function of the membrane potential with gain parameter γ :

$$y_j(t) = \frac{1}{\left(1 + \frac{1}{\gamma[V_m(t) - \Theta]_+}\right)} \quad (3)$$

This can be convolved with Gaussian noise, producing a less discontinuous function with a softer lower threshold.

k-Winners-Take-All Inhibition

Leabra uses a KWTA function to achieve sparse distributed representations. This function is achieved by setting a uniform level of inhibition for all units in the layer that prevents more than k units from getting over threshold. This inhibitory current is given by:

$$g_i = g_{k+1}^\ominus + q(g_k^\ominus - g_{k+1}^\ominus) \quad (4)$$

where q is typically .25, and the threshold-level inhibition terms are $g^\ominus = \frac{\sum_{c \neq i} g_c \overline{g}_c (E_c - \Theta)}{\Theta - E_i}$ for the units with the k th and $k + 1$ th highest excitatory inputs. Activation dynamics similar to those produced by this

function have been shown to result from simulated inhibitory interneurons that project both feedforward and feedback inhibition [57].

Error-Driven Learning

Error-driven learning is implemented in Leabra using a symmetric version of the biologically plausible GeneRec algorithm [25], that is functionally equivalent to the deterministic Boltzmann machine and contrastive Hebbian learning (CHL) [27, 29]. The network settles in two phases, an expectation (minus) phase and an outcome (plus) phase, and then computes a simple difference of a pre and postsynaptic activation product across these two phases:

$$\Delta w_{ij} = x_i^+ y_j^+ - x_i^- y_j^- \quad (5)$$

for sending unit x_i and receiving unit y_j in the two phases.

Hebbian Learning

The simplest form of Hebbian learning adjusts the weights in proportion to the product of the sending (x_i) and receiving (y_j) unit activations: $\Delta w_{ij} = x_i y_j$. The weight vector is dominated by the principal eigenvector of the pairwise correlation matrix of the input, but it also grows without bound. Leabra uses a variant of the Oja normalization [33]:

$$\Delta w_{ij} = x_i y_j - y_j w_{ij} = y_j (x_i - w_{ij}) \quad (6)$$

which can also be seen as computing the expected value of the sending unit activity conditional on the receiver's activity (if treated like a binary variable active with probability y_j): $w_{ij} \approx \langle x_i | y_j \rangle_p$. This is essentially the same rule used in standard competitive learning or mixtures-of-Gaussians [10, 38].

Error-driven and Hebbian learning are linearly combined at each synapse in the network, using a normalized mixing constant. To keep the error-driven component within the same 0-1 range of the Hebbian term, soft weight bounding with exponential approach to these extremes is used on this component. Finally, a sigmoidal contrast enhancement function on the weights can be used to facilitate learning in environments with underlying binary features (i.e., imposing a bias towards binary weights). See [57] for details.

Principal Results

In [57], Leabra is used to replicate a large number of published models that were originally implemented using a variety of different algorithms from backpropagation to Hebbian self-organizing learning. Leabra also illustrates many of the issues discussed in this paper regarding the interactions among the different principles. For example, just adding interactivity to an otherwise generic error-driven network (e.g., a GeneRec network) significantly impairs generalization performance. However, also adding Hebbian learning and inhibitory competition (in Leabra) restores good generalization performance within an interactive network (O'Reilly, 1996, PhD Thesis, Carnegie Mellon University). The conclusion is similar to that of the GRAIN exploration of interactivity and noise — interactivity itself may cause problems, but these can be remedied with additional principles.

In addition to replicating existing models, Leabra also provides better models of several phenomena. One salient example of this is in the case of the U-shaped past-tense overregularization phenomenon, which has proven difficult to capture using purely error-driven backpropagation networks without also manipulating the training environment in a potentially implausible fashion [58, 59, 60]. By adding Hebbian learning and

inhibitory competition, Leabra introduces biases that produce a much more pronounced U-shaped effect (including a longer period of early competence) [57]. This can be contrasted with the essentially monotonic decrease in overregularizations that, in retrospect, is exactly what would be expected from a purely error-driven algorithm (see [57] and Hoeffner, 1997, PhD Thesis, Carnegie Mellon University for a more detailed critique of existing models).

*

References

- [1] J. L. McClelland. The GRAIN model: A framework for modeling the dynamics of information processing. In D. E. Meyer and S. Kornblum, editors, *Attention and Performance XIV: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience*, pages 655–688. Lawrence Erlbaum Associates, Hillsdale, NJ, 1993.
- [2] Yuko Munakata, James L. McClelland, M. J. Johnson, and Robert S. Siegler. Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review*, 104:686–713, 1997.
- [3] R. Desimone and L. G. Ungerleider. Neural mechanisms of visual processing in monkeys. In F. Boller and J. Grafman, editors, *Handbook of Neurophysiology, Vol. 2*, chapter 14, pages 267–299. Elsevier, Amsterdam, 1989.
- [4] A. P. Georgopoulos. Neurophysiology and reaching. In M. Jeannerod, editor, *Attention and Performance, vol. 13*, pages 227–263. Erlbaum, Hillsdale, N.J., 1990.
- [5] S. C. Rao, G. Rainer, and E. K. Miller. Integration of what and where in the primate prefrontal cortex. *Science*, 276:821, 1997.
- [6] G. E. Hinton, J. L. McClelland, and D. E. Rumelhart. Distributed representations. In D. E. Rumelhart, J. L. McClelland, and PDP Research Group, editors, *Parallel Distributed Processing. Volume 1: Foundations*, chapter 3, pages 77–109. MIT Press, Cambridge, MA, 1986.
- [7] P. L. A. Gabbot and P. Somogyi. Quantitative distribution of GABA-immunoreactive neurons in the visual cortex (area 17) of the cat. *Experimental Brain Research*, 61:323–331, 1986.
- [8] T. Kohonen. *Self-Organization and Associative Memory*. Springer-Verlag, Berlin, 1984.
- [9] J. L. McClelland and D. E. Rumelhart. An interactive activation model of context effects in letter perception: Part 1. an account of basic findings. *Psychological Review*, 88(5):375–407, 1981.
- [10] D. E. Rumelhart and D. Zipser. Feature discovery by competitive learning. In D. E. Rumelhart, J. L. McClelland, and PDP Research Group, editors, *Parallel Distributed Processing. Volume 1: Foundations*, chapter 5, pages 151–193. MIT Press, Cambridge, MA, 1986.
- [11] S. Grossberg. Adaptive pattern classification and universal recoding I: Parallel development and coding of neural feature detectors. *Biological Cybernetics*, 23:121–134, 1976.
- [12] H. B. Barlow. Unsupervised learning. *Neural Computation*, 1:295–311, 1989.
- [13] David J. Field. What is the goal of sensory coding? *Neural Computation*, 6(4):559–601, 1994.
- [14] B. A. Olshausen and D. J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381:607, 1996.
- [15] D. J. Felleman and D. C. Van Essen. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1:1–47, 1991.

- [16] J. B. Levitt, D. A. Lewis, T. Yoshioka, and J. S. Lund. Topography of pyramidal neuron intrinsic connections in macaque monkey prefrontal cortex (areas 9 & 46). *Journal of Comparative Neurology*, 338:360–376, 1993.
- [17] E. L. White. *Cortical Circuits: Synaptic Organization of the Cerebral Cortex, Structure, Function, and Theory*. Birkhäuser, Boston, 1989.
- [18] P. Smolensky. Information processing in dynamical systems: Foundations of harmony theory. In D. E. Rumelhart, J. L. McClelland, and PDP Research Group, editors, *Parallel Distributed Processing. Volume 1: Foundations*, chapter 5, pages 282–317. MIT Press, Cambridge, MA, 1986.
- [19] David H. Ackley, G. E. Hinton, and T. J. Sejnowski. A learning algorithm for Boltzmann machines. *Cognitive Science*, 9:147–169, 1985.
- [20] Shaun P. Vecera and Randall C. O’Reilly. Figure-ground organization and object recognition processes: An interactive account. *Journal of Experimental Psychology: Human Perception and Performance*, 24:441–462, 1998.
- [21] M. A. Peterson. Object recognition processes can and do operate before figure ground organization. *Current Directions in Psychological Science*, 3:105–111, 1994.
- [22] D. E. Rumelhart, G. E. Hinton, and R. J. Williams. Learning internal representations by error propagation. In D. E. Rumelhart, J. L. McClelland, and PDP Research Group, editors, *Parallel Distributed Processing. Volume 1: Foundations*, chapter 8, pages 318–362. MIT Press, Cambridge, MA, 1986.
- [23] Francis H. C. Crick. The recent excitement about neural networks. *Nature*, 337:129–132, 1989.
- [24] D. Zipser and R. A. Andersen. A back propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, 331:679–684, 1988.
- [25] R. C. O’Reilly. Biologically plausible error-driven learning using local activation differences: The generalized recirculation algorithm. *Neural Computation*, 8(5):895–938, 1996.
- [26] G. E. Hinton and J. L. McClelland. Learning representations by recirculation. In D. Z. Anderson, editor, *Neural Information Processing Systems, 1987*, pages 358–366, New York, 1988. American Institute of Physics.
- [27] G. E. Hinton. Deterministic Boltzmann learning performs steepest descent in weight-space. *Neural Computation*, 1:143–150, 1989.
- [28] C. Peterson and J. R. Anderson. A mean field theory learning algorithm for neural networks. *Complex Systems*, 1:995–1019, 1987.
- [29] J. R. Movellan. Contrastive Hebbian learning in the continuous Hopfield model. In D. S. Touretzky, G. E. Hinton, and T. J. Sejnowski, editors, *Proceedings of the 1989 Connectionist Models Summer School*, pages 10–17, San Mateo, CA, 1990. Morgan Kaufman.
- [30] J. L. McClelland. The interaction of nature and nurture in development: A parallel distributed processing perspective. In P. Bertelson, P. Eelen, and G. D’Ydewalle, editors, *Current Advances in Psychological Science: Ongoing Research*, pages 57–88. Erlbaum, Hillsdale, NJ, 1994.
- [31] D. O. Hebb. *The Organization of Behavior*. Wiley, New York, 1949.

- [32] G. L. Collingridge and T. V. P. Bliss. NMDA receptors - their role in long-term potentiation. *Trends in Neurosciences*, 10:288–293, 1987.
- [33] E. Oja. A simplified neuron model as a principal component analyzer. *Journal of Mathematical Biology*, 15:267–273, 1982.
- [34] Thomas K. Landauer and Susan T. Dumais. a solution to plato's problem: the latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104:211–240, 1997.
- [35] D. W. Massaro. Testing between the TRACE model and the fuzzy logical model of speech perception. *Cognitive Psychology*, 21:398–421, 1989.
- [36] J. Morton. The interaction of information in word recognition. *Psychological Review*, 76:165–178, 1969.
- [37] G. E. Hinton and Z. Ghahramani. Generative models for discovering sparse distributed representations. *Philosophical Transactions of the Royal Society (London) B*, 352:1177–1190, 1997.
- [38] S. J. Nowlan. Maximum likelihood competitive learning. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems*, 2, pages 574–582. Morgan Kaufman, San Mateo, CA, 1990.
- [39] R. A. Jacobs, M. I. Jordan, S. J. Nowlan, and G. E. Hinton. Adaptive mixtures of local experts. *Neural Computation*, 3:79–87, 1991.
- [40] P. Dayan and R. S. Zemel. Competition and multiple cause models. *Neural Computation*, 7:565–579, 1995.
- [41] Eric Saund. A multiple cause mixture model for unsupervised learning. *Neural Computation*, 7:51–71, 1995.
- [42] R. S. Zemel. *A Minimum Description Length Framework for Unsupervised Learning*. PhD thesis, University of Toronto, Canada, 1993.
- [43] E. L. Bienenstock, L. N. Cooper, and P. W. Munro. Theory for the development of neuron selectivity: Orientation specificity and binocular interaction in visual cortex. *Journal of Neuroscience*, 2(2):32–48, 1982.
- [44] P. Földiák. Forming sparse representations by local anti-Hebbian learning. *Biological Cybernetics*, 64(2):165–170, 1990.
- [45] E. Majani, R. Erlarson, and Y. Abu-Mostafa. The induction of multiscale temporal structure. In D. S. Touretzky, editor, *Advances in Neural Information Processing Systems*, 1, pages 634–642, San Mateo, CA, 1989. Morgan Kaufmann.
- [46] T. Fukai and S. Tanaka. A simple neural network exhibiting selective activation of neuronal ensembles: from winner-take-all to winners-share- all. *Neural Computation*, 9:77–97, 1997.
- [47] Geoffrey E. Hinton. Connectionist learning procedures. *Artificial Intelligence*, 40:185–234, 1989.
- [48] P. Baldi and K. Hornik. Neural networks and principal components analysis: Learning from examples without local minima. *Neural Networks*, 2:53–58, 1989.

- [49] P. Dayan, G. E. Hinton, R. N. Neal, and R. S. Zemel. The helmholtz machine. *Neural Computation*, 7:889–904, 1995.
- [50] S. Becker. Mutual information maximization: Models of cortical self-organization. *Network : Computation in Neural Systems*, 7:7–31, 1996.
- [51] V. R. de Sa and D. H. Ballard. Category learning through multimodality sensing. *Neural Computation*, 10:1097–1118, 1998.
- [52] J. Kay, D. Floreano, and W. Phillips. Contextually guided unsupervised learning using local multivariate binary processors. *Neural Networks*, 11:117–140, 1998.
- [53] P. J. B. Hancock, L. S. Smith, and W. A. Phillips. A biologically supported error-correcting learning rule. *Neural Computation*, 3:201–212, 1991.
- [54] Stuart Geman, E. L. Bienenstock, and Rene Doursat. Neural networks and the bias/variance dilemma. *Neural Computation*, 4:1–58, 1992.
- [55] A. S. Weigend, D. E. Rumelhart, and B. A. Huberman. Generalization by weight-elimination with application to forecasting. In Richard P. Lippmann, John E. Moody, and David S. Touretzky, editors, *Advances In Neural Information Processing Systems*, 3, pages 875–882. Morgan Kaufman, San Mateo, CA, 1991.
- [56] Dong-Gyu Jeong and Soo-Young Lee. Merging back-propagation and hebbian learning rules for robust classifications. *Neural Networks*, 9:1213–1222, 1996.
- [57] R. C. O’Reilly and Y. Munakata. *Cognitive Neuroscience: A Computational Exploration*. MIT Press, Cambridge, MA, in press.
- [58] David E. Rumelhart and James L. McClelland. On learning the past tenses of English verbs. In J. L. McClelland, D. E. Rumelhart, and PDP Research Group, editors, *Parallel Distributed Processing. Volume 2: Psychological and Biological Models*, pages 216–271. MIT Press, Cambridge, MA, 1986.
- [59] Kim Plunkett and Virginia A. Marchman. From rote learning to system building: Acquiring verb morphology in children and connectionist nets. *Cognition*, 48(1):21–69, 1993.
- [60] Gary F. Marcus, Steven Pinker, Michael Ullman, Michelle Hollander, John T. Rosen, and Fei Xu. Over-regularization in language acquisition. *Monographs of the Society for Research in Child Development*, 57(4):1–165, 1992.