# Optimizing search strategies for supervised learning by response modulation

741.16


Christian D. Swinehart & L.F. Abbott

Department of Biology & Volen Center for Complex Systems | Brandeis University, Waltham MA, USA


## Function Approximation Learning

*Feedback-based Model*

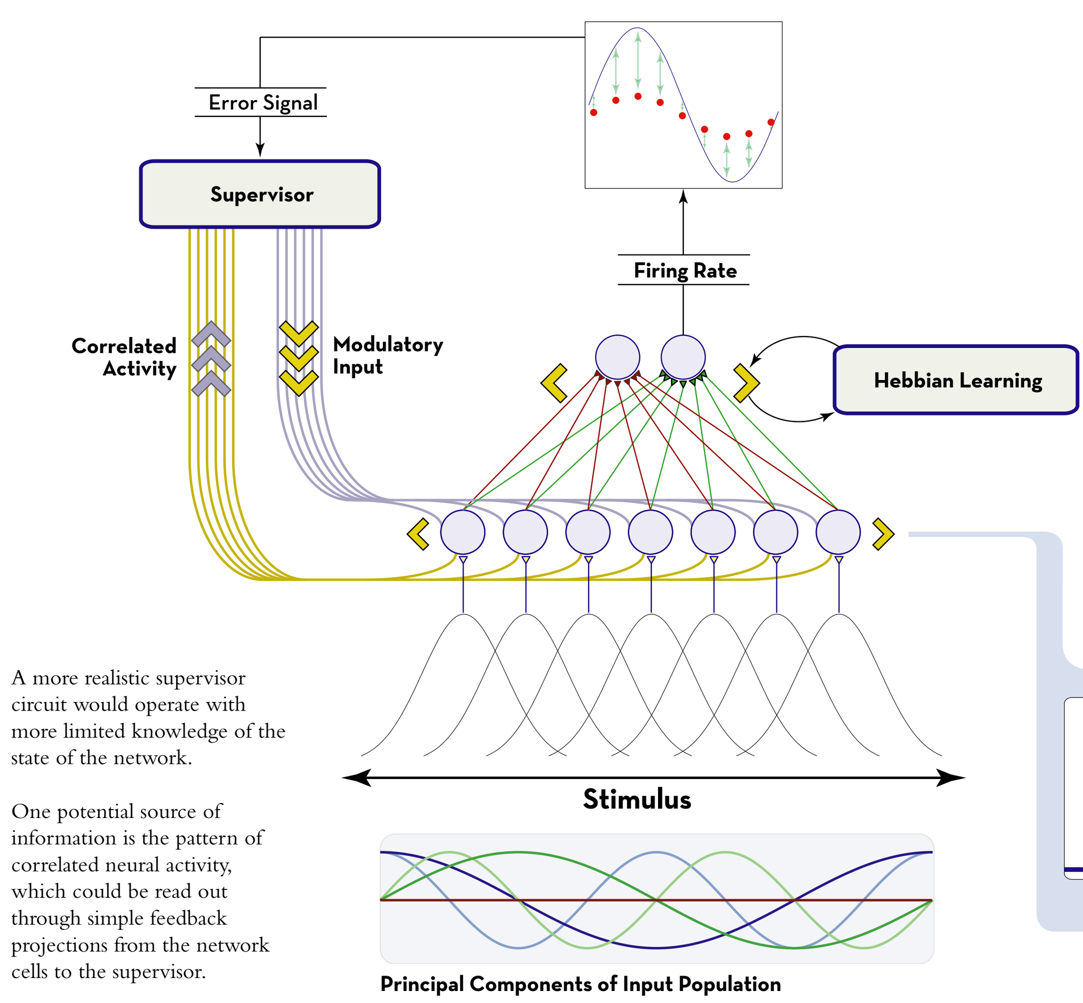


A major drawback of the traditional, synaptic approach to supervised learning is the lack of a plausible mechanism by which error signals generated by the supervisor could directly control synaptic plasticity.

In this study, the supervisor modifies intrinsic neuronal response properties, rather than synaptic strengths, to guide learning.

Recent work (Chance et al., 2002) has shown that the response properties of individual neurons can be modified through changes in the level and balance of their excitatory and inhibitory inputs. This provides a pathway by which supervisory error signals can reach elements of a network to guide its modification during learning.

Synaptic plasticity occurs in parallel via unsupervised, Hebbian learning. Thus the supervisor can only influence synaptic modification indirectly, by affecting the excitability of the cells in a manner which will lead to Hebbian strengthening or weakening of connections.

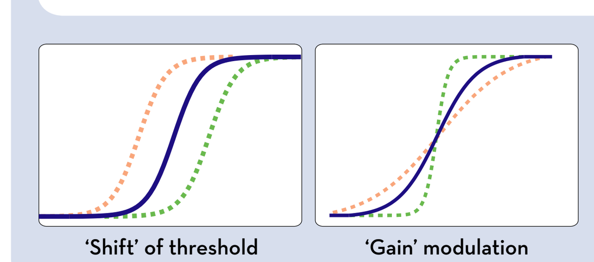


Intrinsic Properties Modification

A more realistic supervisor circuit would operate with more limited knowledge of the state of the network.

One potential source of information is the pattern of correlated neural activity, which could be read out through simple feedback projections from the network cells to the supervisor.

## The Chicken & Egg Problem of Plasticity

*Hebbian plasticity can't create something from nothing...*

### Hebbian Learning Rule

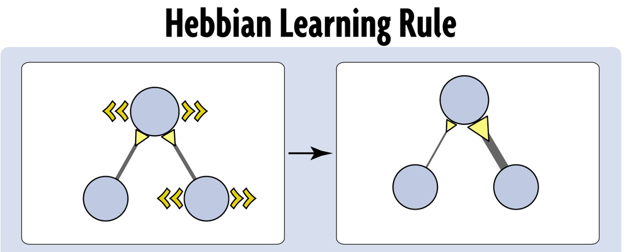

Hebbian Learning modifies connections using purely local information: the pre- and post-synaptic firing rates.

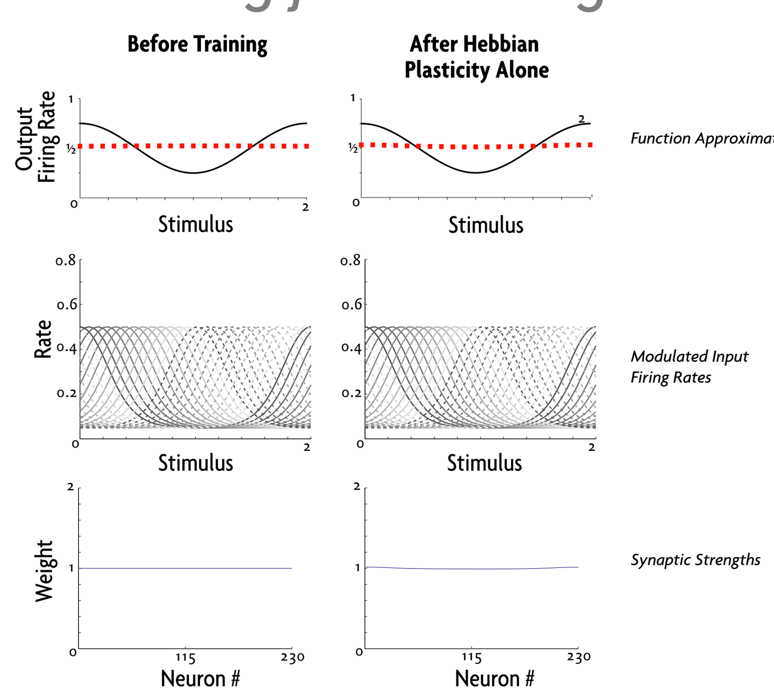


As a correlation-based mechanism, Hebbian plasticity modifies connections as a function of preexisting patterns of activity within the network.

However, when network behavior is homogeneous (as is the case before training), there is no structured activity for Hebb to amplify.

Thus a correlation rule alone is not sufficient to generate a desired behavior appropriate to a given task.

*...while Response Modulation can influence behavior, but cannot make permanent, synaptic changes to the network*

### Response Modulation

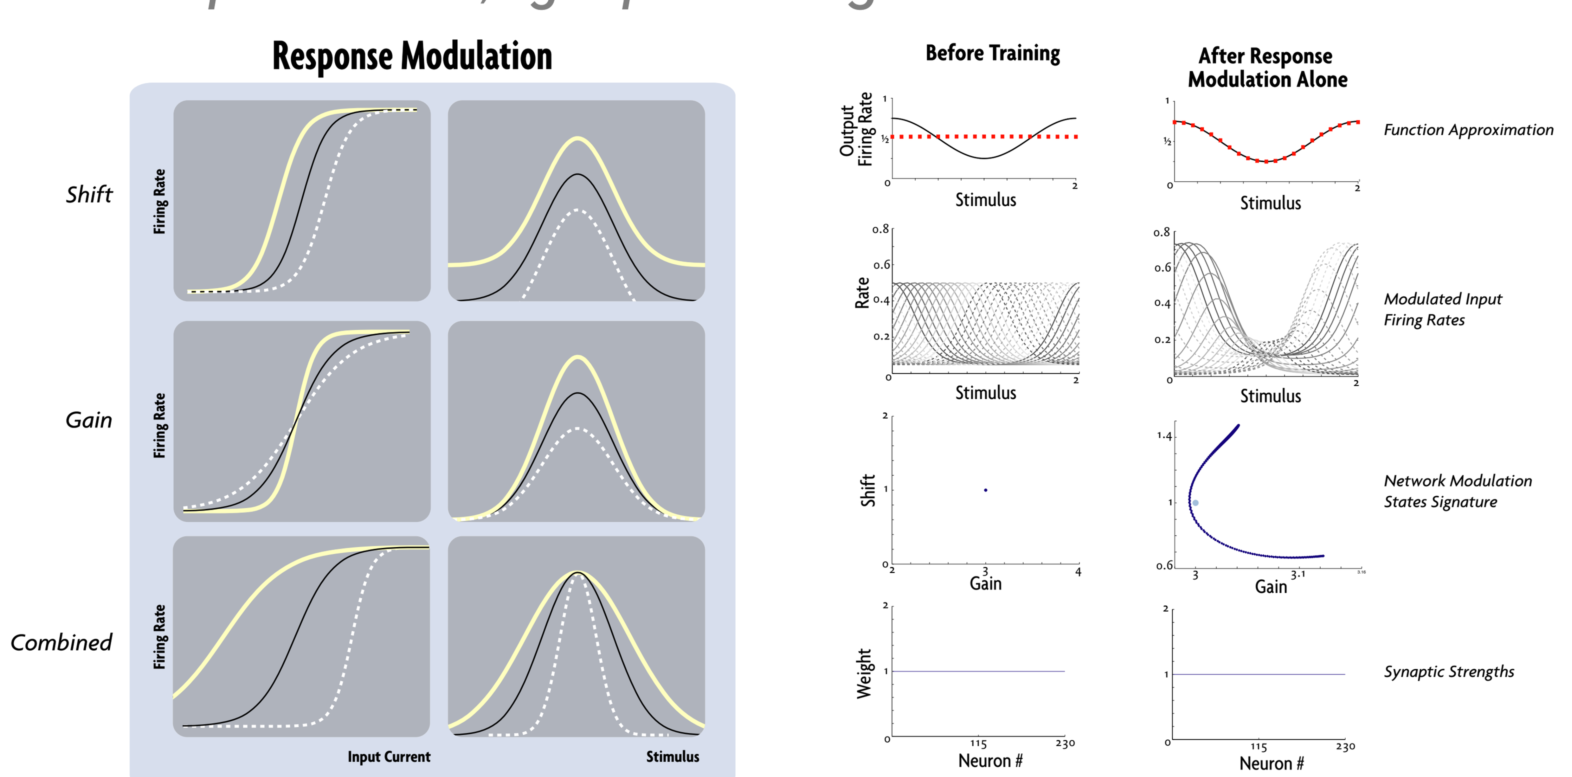


## Learning in Simple Networks

*Response Modulation-based Hebbian Learning*

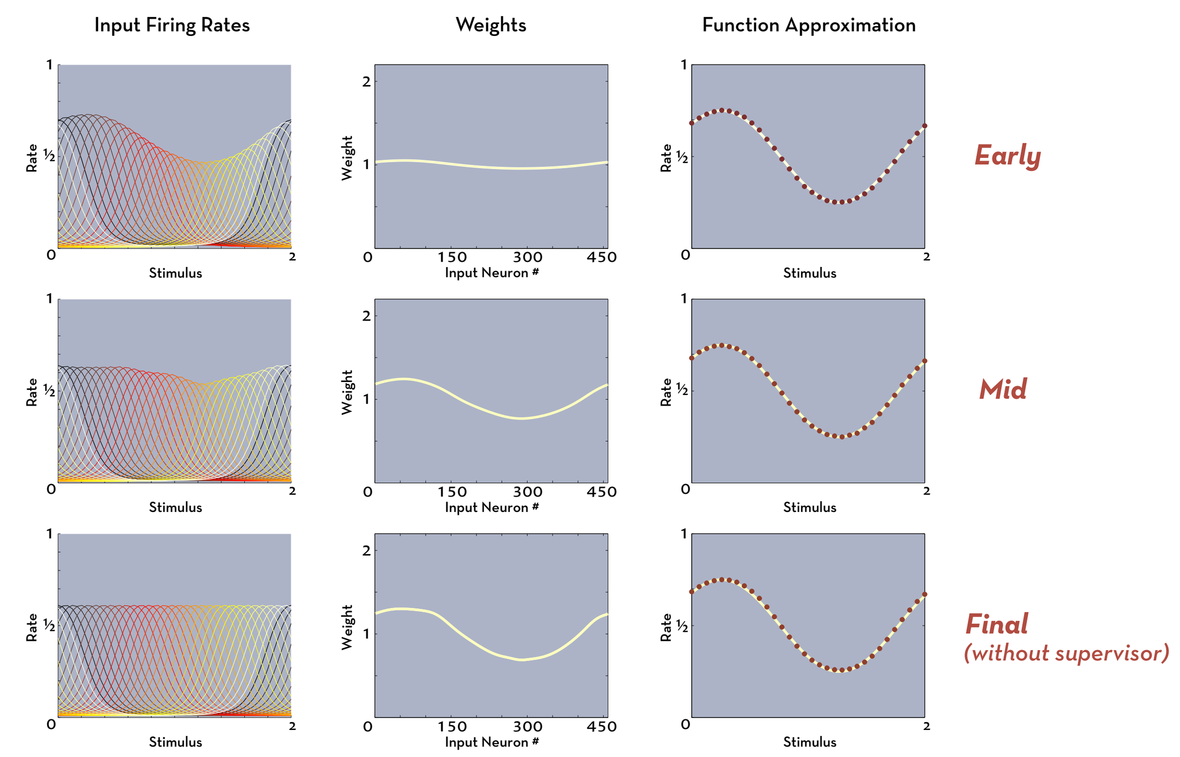


## Nature of the Supervisor

*Profiting from redundancy in the input population*

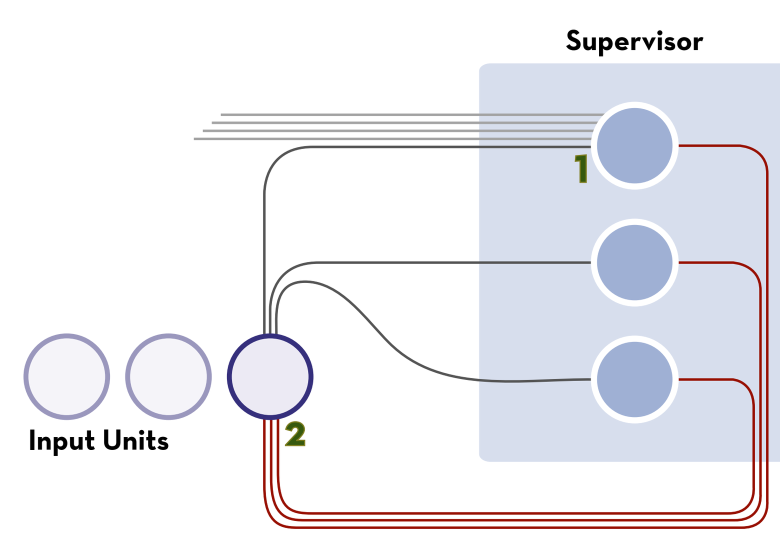


Rules governing plasticity

**1** $$\Delta w_{ij} = \eta_s y_i \left(x_j - \sum_{k=1}^{i} y_k w_{kj}\right)$$

*The Sanger Rule*

**2** $$w_{ji} \rightarrow w_{ji} + \eta_o x_j \left(y_i - x_j w_{ji}\right)$$

*The Oja Rule*

### An Aside...

***Principal Component Analysis***

- A method of accounting for variance in a set of data and discovering relationships between variables
- Correlations in data can be used to extract new axes
- The original data can be projected onto these new axes which better represent the underlying structure

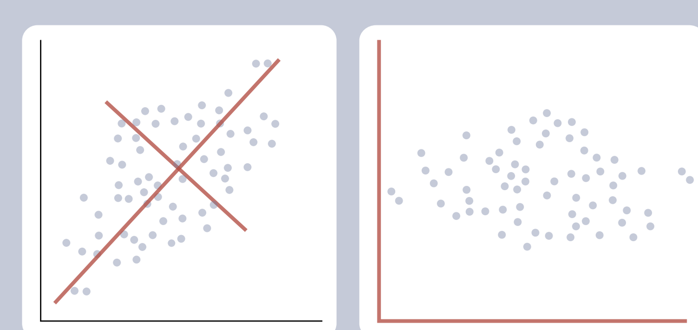

*Learning the strengths of the feedback connections*

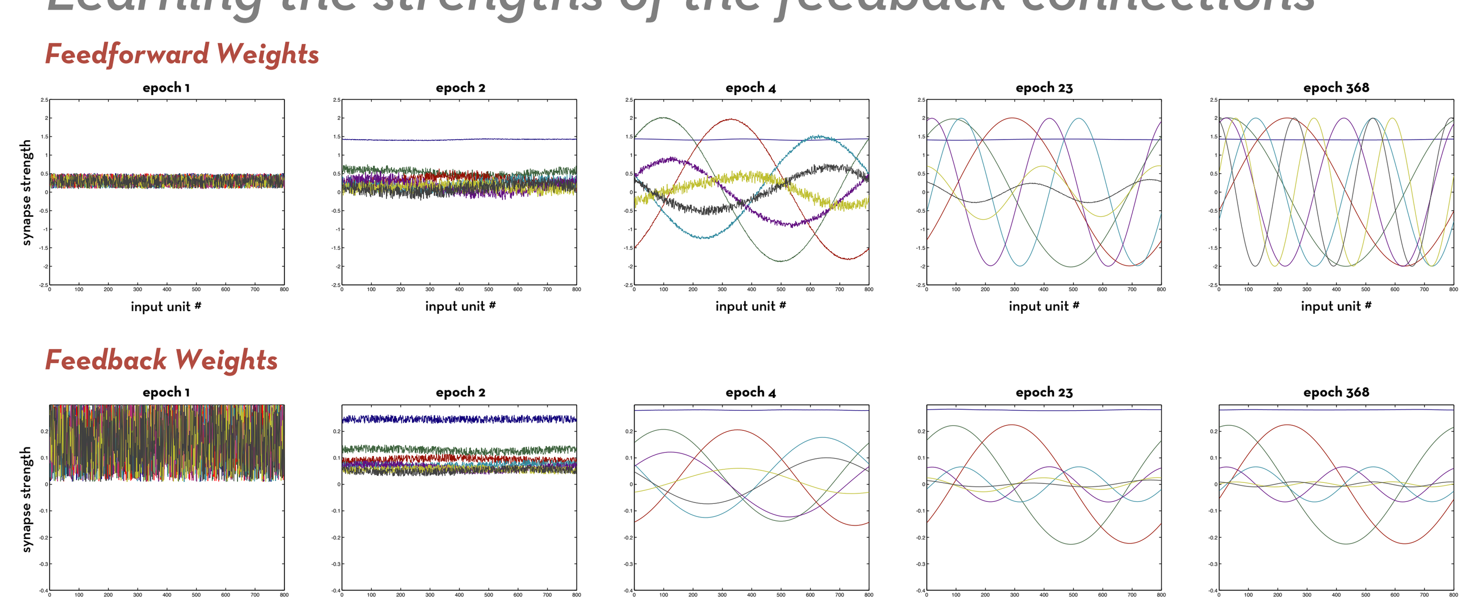


## A PCA-derived Strategy

*Supervision by driving feedback cells*

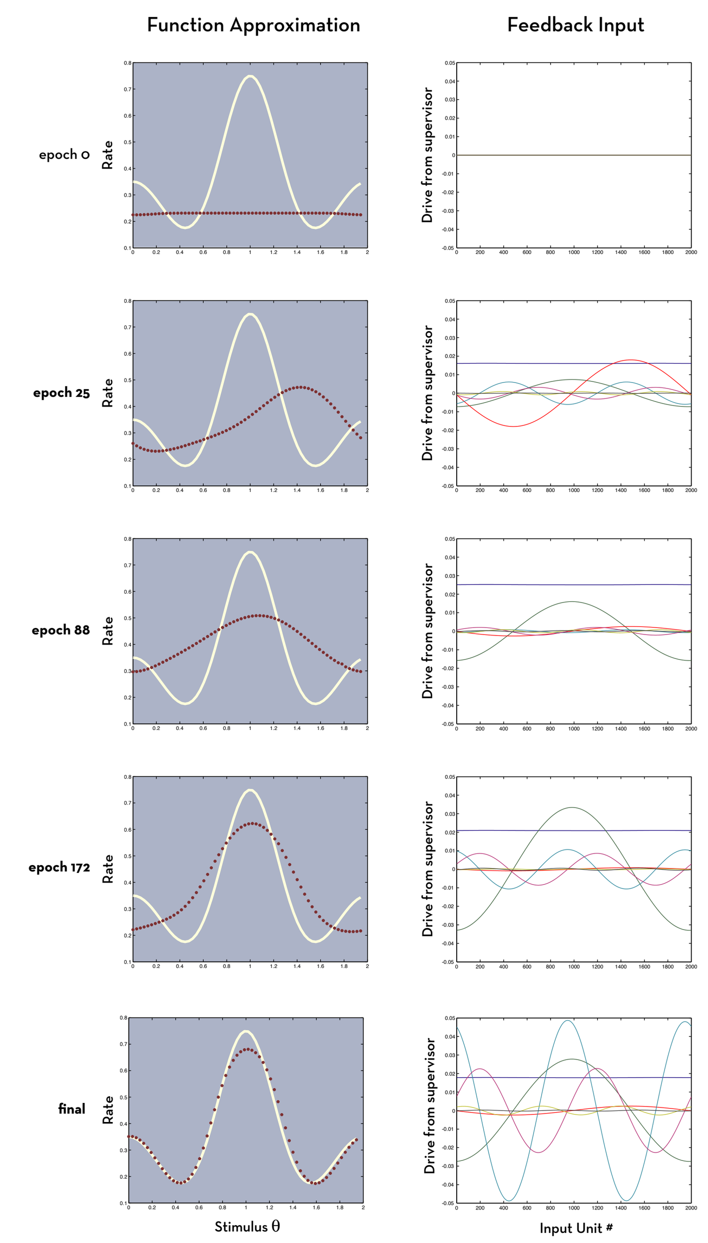


Previously, learning consisted of the supervisor modulating the shifts and gains of cells in the input population by independently modifying the balance of excitation and inhibition for each individual cell.

This required making a decision of how to minimize error on a cell-by-cell basis, making for a search space with 2N dimensions—a situation that will likely scale quite poorly beyond even moderately-sized networks.

In contrast, the PCA approach simplifies the problem for the supervisor circuit to setting the firing rates of seven cells. Provided that it works, this represents a substantial reduction in the dimensionality of the search space and shrinks the problem to one solvable by very simple search algorithms that could easily be realized neurally.

The plots at left demonstrate that this is in fact a workable strategy. By doing a random walk over this 7-dimensional space (guided by a leaky integrator driven by accumulated error), a successful approximation of the target function can be generated.

This is possible (despite the few degrees of freedom available to the supervisor circuit) because of the relationship between the innervation pattern of the feedback projections and the response properties of the input units. Since they were derived from the correlations in the activity of the input population, the negative effects of this reduction in control is minimized, while the benefits (in terms of search-space reduction) can be exploited.

## Alternative Input Statistics

*Learning under alternative input distributions*

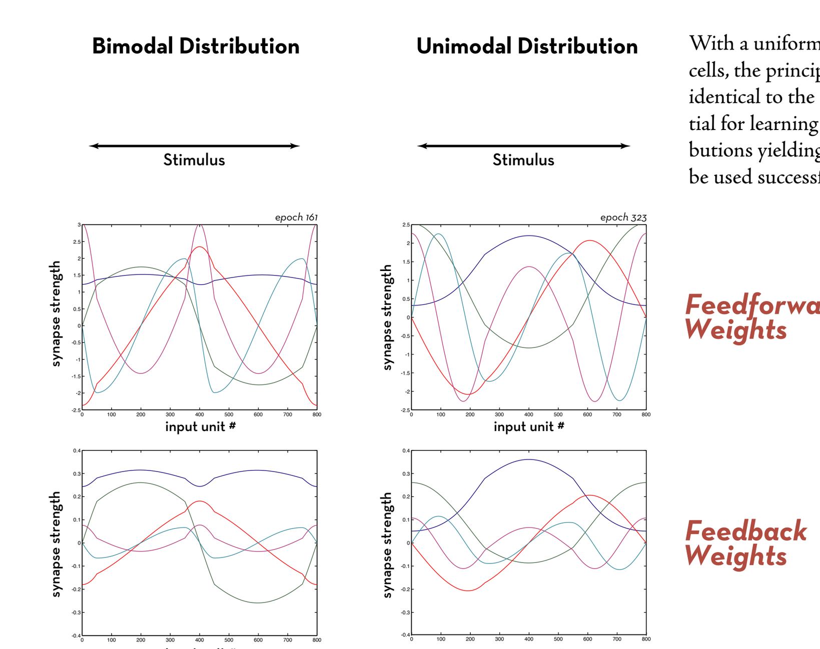


With a uniform, periodic distribution of input cells, the principal components extracted are identical to the Fourier modes. This is not essential for learning to work however. Other distributions yielding other basis functions can also be used successfully.

*Modulation compensates for input bias*

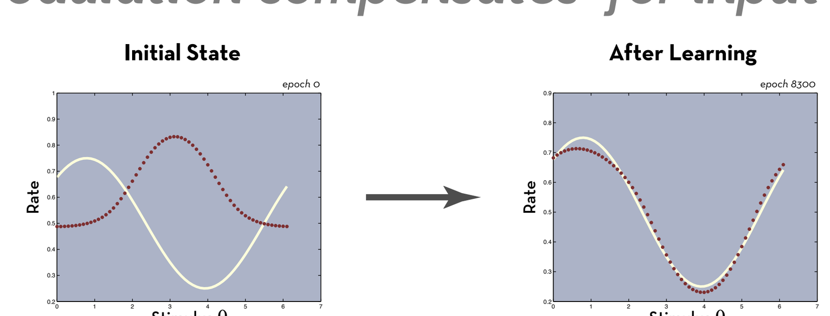


This network's input population has a central bias with twice the input population density between ¾π and 1¼π of the surrounding region.

As a result the output cell receives greater drive in this central region leading to a peaked response function in its untrained state.

Despite the population bias toward high firing rates in the center and low rates at the edges, modulating the firing rates of the supervisor cells can compensate for this.

Again, this would not be possible were the connectivity patterns not matched to the population statistics

## Assessing Improvement

*Quantifying the Dimensional Reduction*

**Performance vs Population Size**

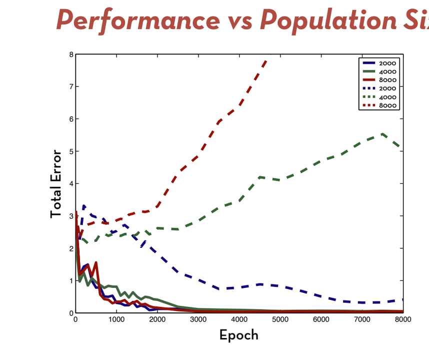


On the basis of the dimensional reduction in the search space, one would expect faster convergence to a solution in the PCA case than under the individual modulation scheme. This prediction is borne out in the comparison data, showing that the PCA supervisor is both faster and scales to much larger population sizes without a significant change in speed or accuracy.

Additionally, increasing the number of components the supervisor draws on improves accuracy, but only a relatively small number is required to account for most of the variability in a given target function, with improvements plateauing after ~5 components are used.

**Performance vs Number of Components**

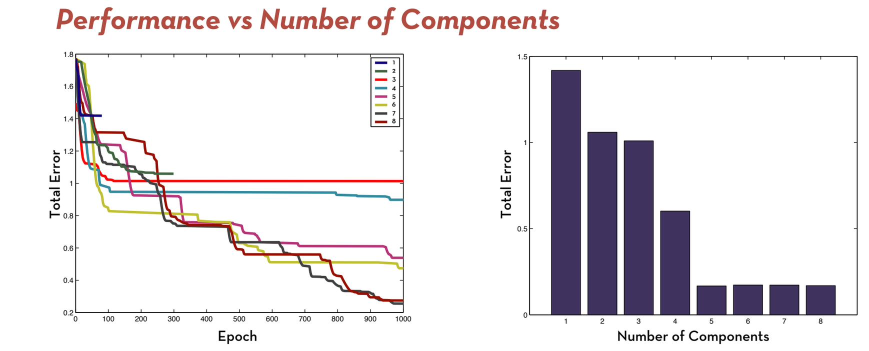


*Implications for Hebbian learning*

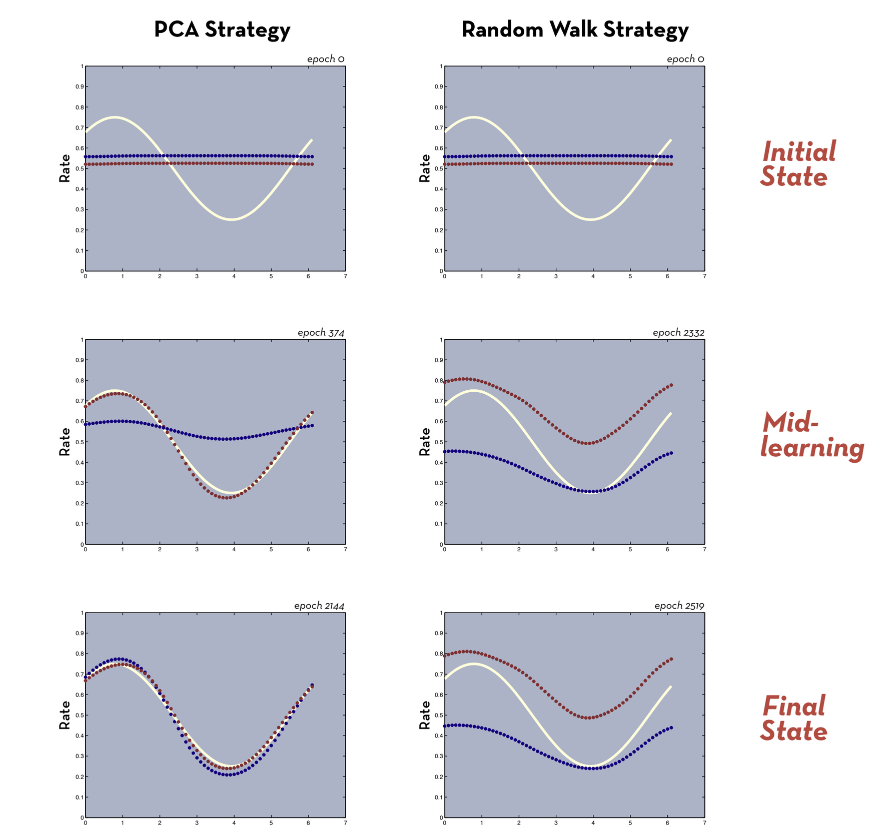


*Dynamics of PCA-based learning*

**PCA Strategy**

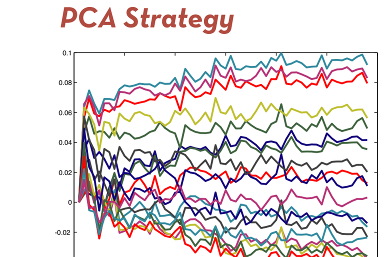

**Random Walk Strategy**

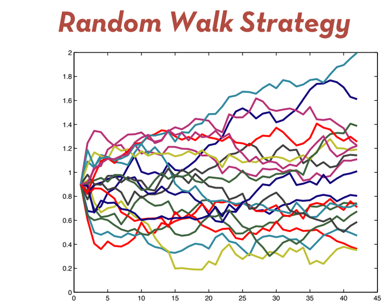

The reason PCA-based learning works for this task is that nearby input units behave similarly in response to stimulus values. As a result the supervisor can exploit this by modulating them similarly rather than wasting resources treating them as though they were independent.

This local similarity of modulation can be seen in the traces at left. Input cells with similar preferences follow similar trajectories as the supervisor finds their optimal values.

The pure random walk strategy treats each unit as unique, thus no similar pattern can be observed in these traces. Instead each unit does its own walk without any relationship to its neighbors.

### References


Chance, FS, Abbott, LF & Reyes, AD (2002) Gain Modulation Through Background Synaptic Input. Neuron **35**:773–782.

Chauvin, Y, & Rumelhart, DE, eds. (1995) Back Propagation: Theory, Architectures, and Applications Hillsdale, nj: Erlbaum.

Doiron, B, Longtin, A, Berman, N, & Maler, L (2001). Subtractive and divisive inhibition: effect of voltage-dependent inhibitory conductances and noise. Neural Comput. **13**:227–248.

Lukashin AV, Wilcox GL, Georgopoulos AP (1994) Overlapping neural networks for multiple motor engrams. Proc Natl Acad Sci USA **9**:8651–8654.

O'Reilly, RC (1996) Biologically plausible error-driven learning using local activation differences: The generalised recirculation algorithm. Neural Computation **8**:895–938.

Prescott, SA & De Koninck Y (2003) Gain control of firing rate by shunting inhibition: Roles of synaptic noise and dendritic saturation. Proc Natl Acad Sci USA **100**:2076–2081.

Poggio, T (1990) A theory of how the brain might work. Cold Spring Harbor Symposium on Quantitative Biology **55**:899–910.

Schultz W, Dayan P, Montague PR (1997) A neural substrate of prediction and reward. Science **275**:1593–1599.

Widrow, B, & Stearns, SD (1985) Adaptive Signal Processing. Englewood Cliffs, nj: Prentice-Hall.

Xie, X & Seung, S (2003) Learning in neural networks by reinforcement of irregular spiking. (unpublished).