Neuronal Supervision of Hebbian Synaptic Plasticity

256.16

Christian D. Swinehart & L.F. Abbott

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

Function Approximation Learning

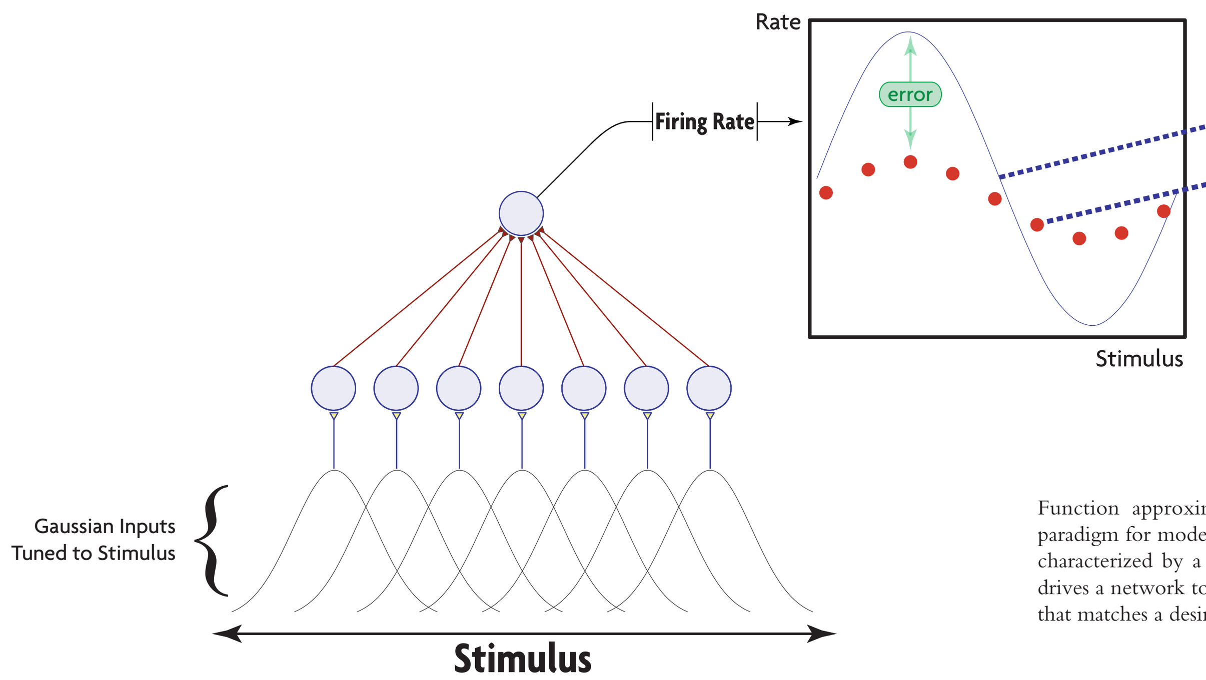

Function approximation is a standard neural network paradigm for modeling learning processes. In it, a stimulus characterized by a single parameter (the 'stimulus value') drives a network to generate a firing rate in an output unit that matches a desired target function of that parameter.

Traditional Supervised Learning

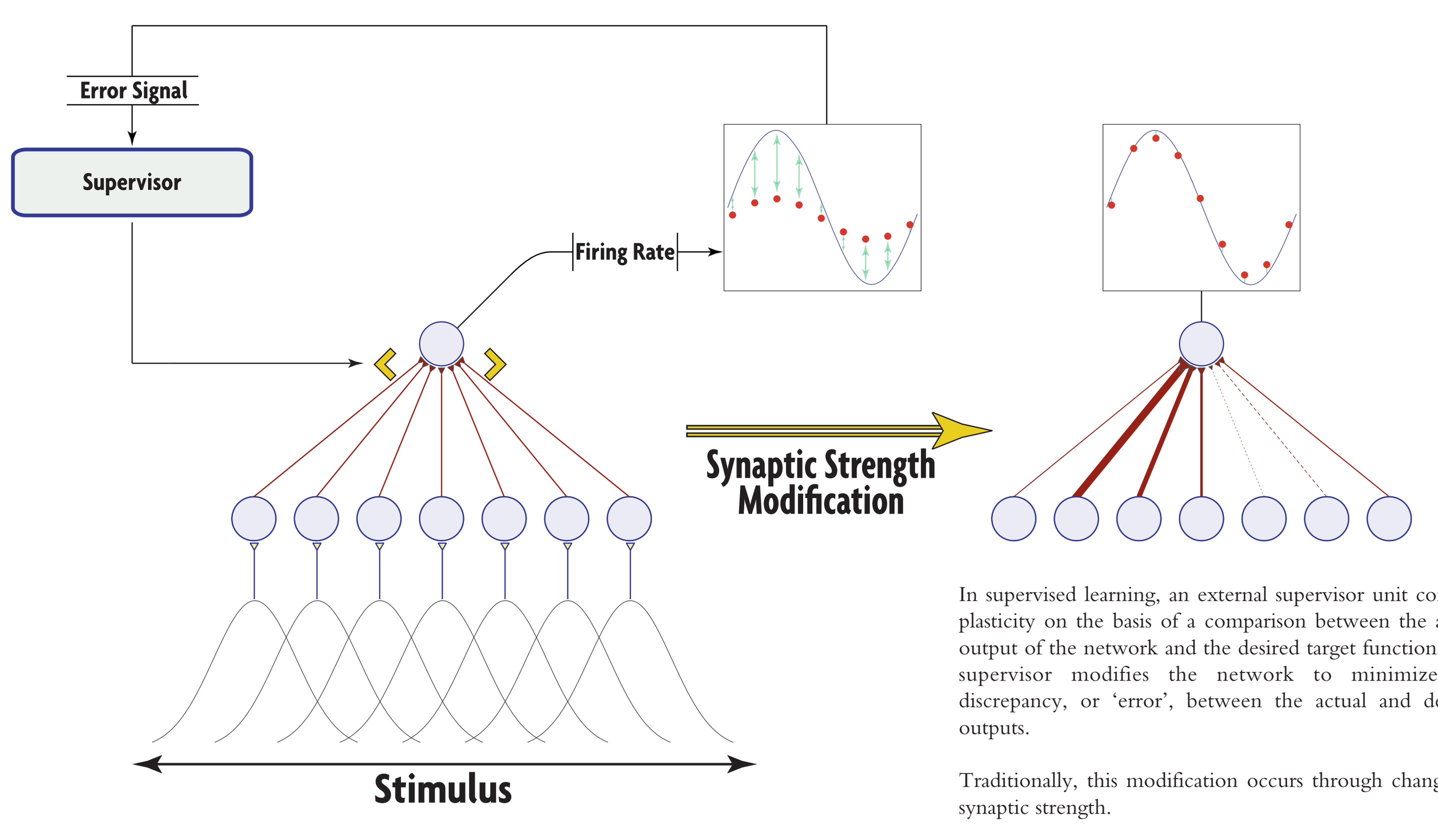

In supervised learning, an external supervisor unit controls plasticity on the basis of a comparison between the actual output of the network and the desired target function. The supervisor modifies the network to minimize the discrepancy, or 'error', between the actual and desired outputs.

Traditionally, this modification occurs through changes in synaptic strength.

Response Modulation-guided Learning

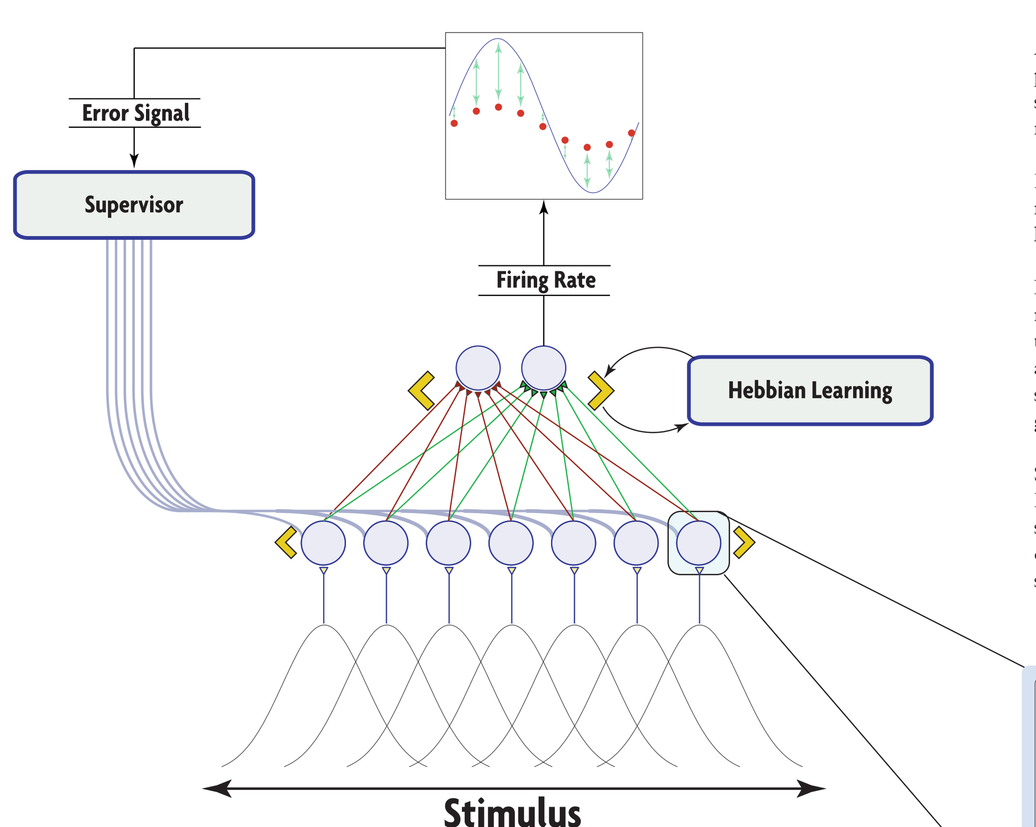

A major drawback of the traditional approach is the lack of a plausible mechanism by which error signals generated by the supervisor could reach individual synapses to guide their modification.

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 supervisor 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

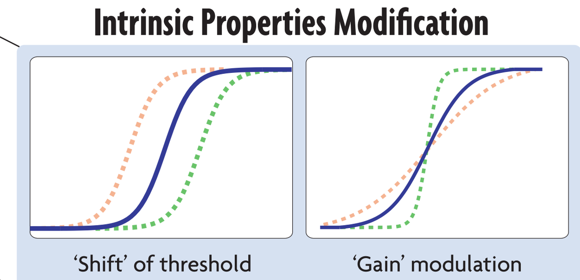

The Chicken & Egg Problem of Plasticity

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

Hebbian Learning Rule

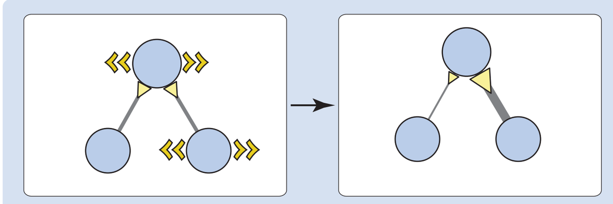

Hebbian Learning modifies connections using purely local information: the pre- and post-synaptic firing rates.

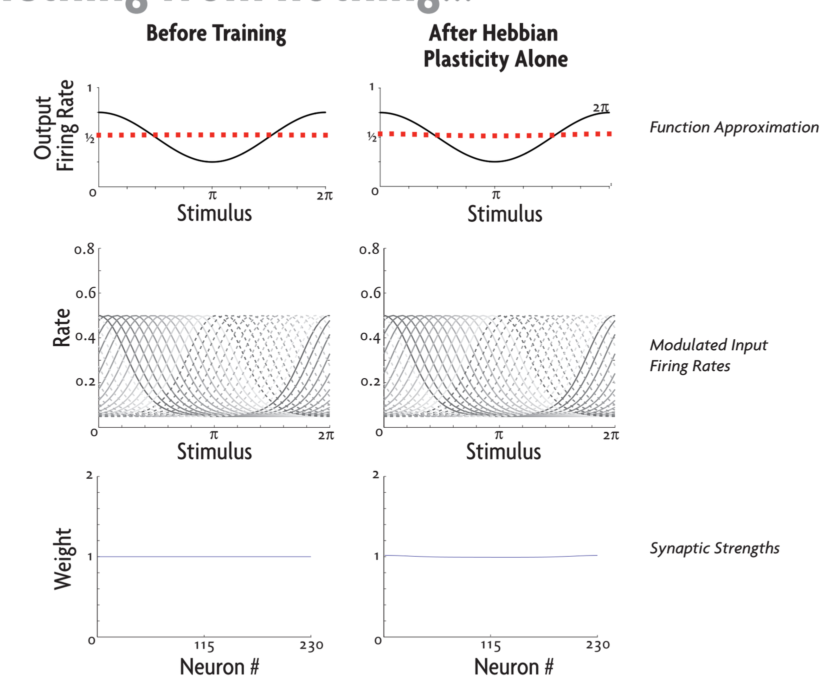

As a correlation-based mechanism, Hebbian plasticity modifies connections as a function of preexisting patterns of activity within the network.

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

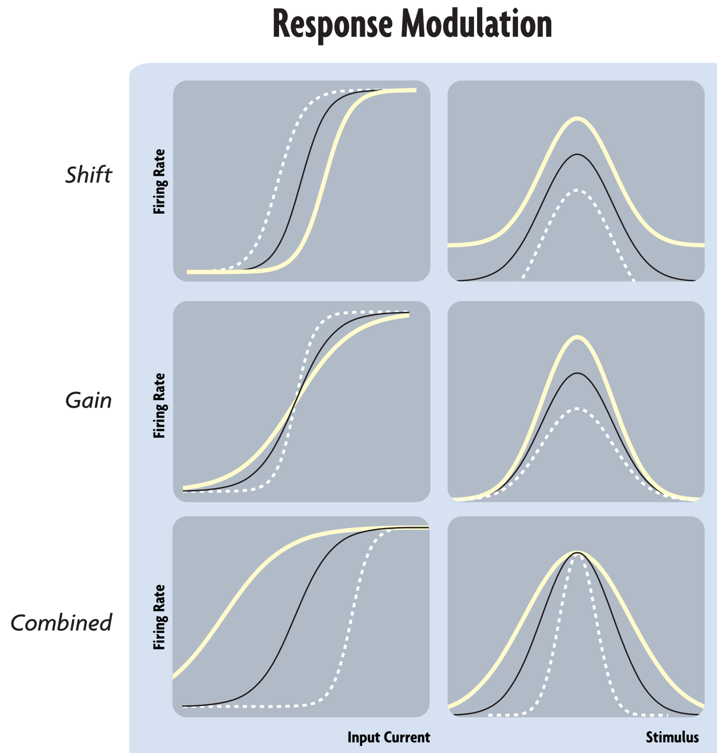

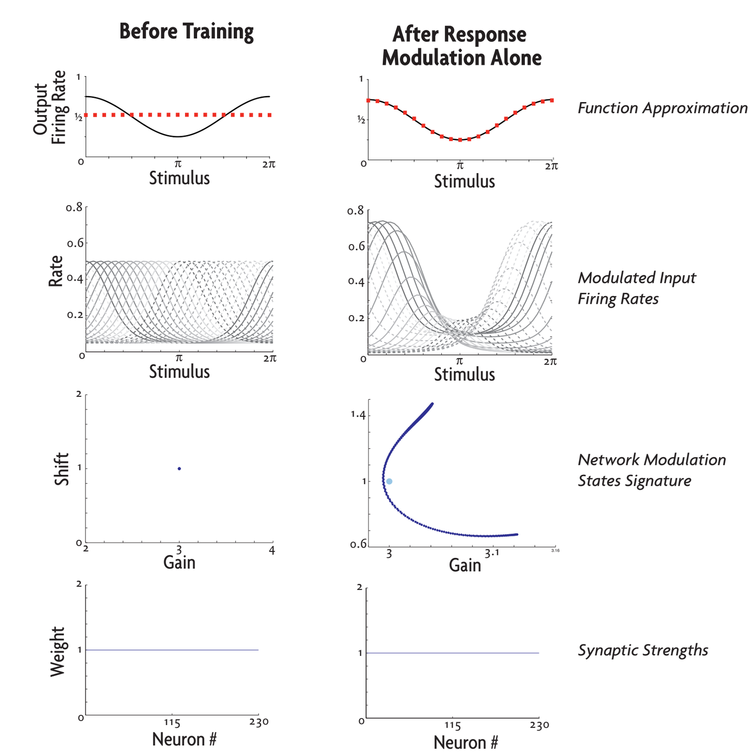

Response Modulation-biased Hebbian Learning

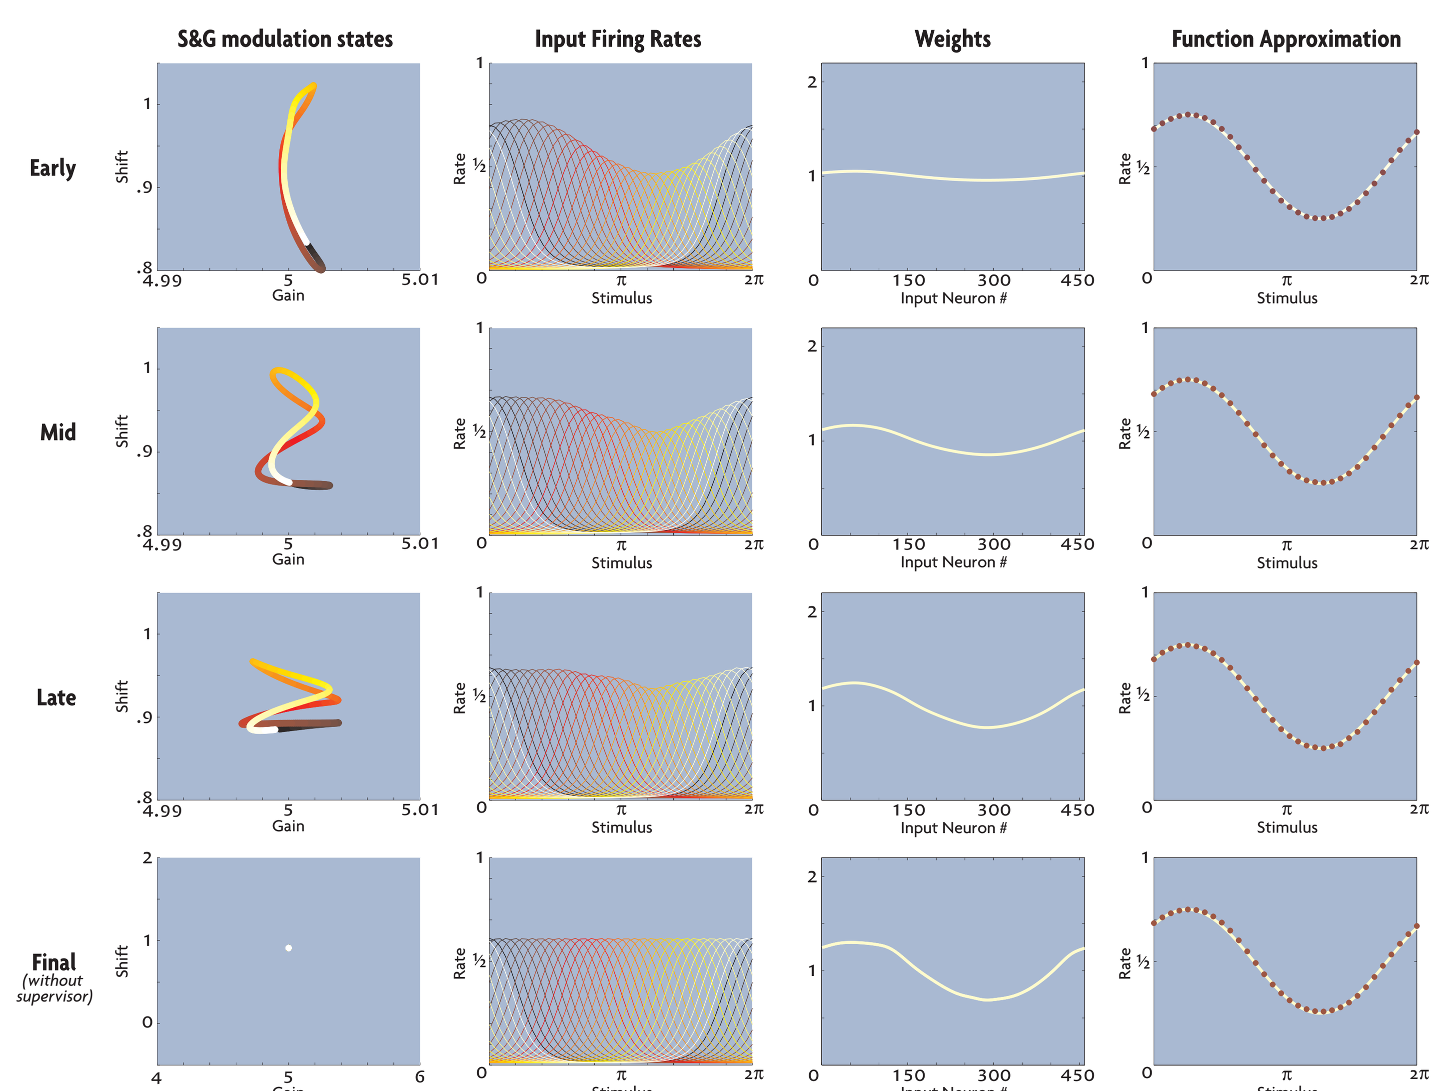

Learning in Single Output Unit Networks

Range & Limitations

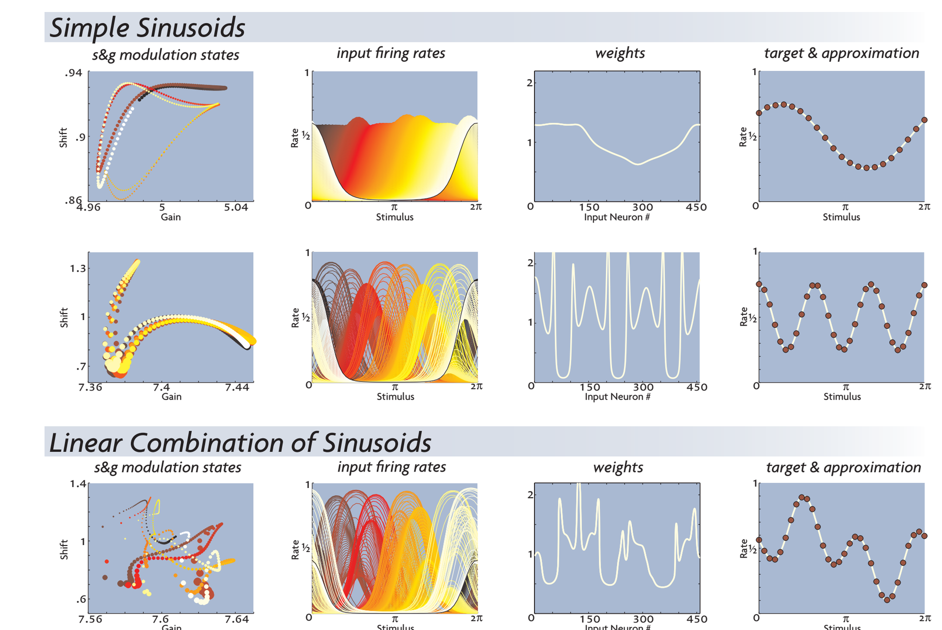

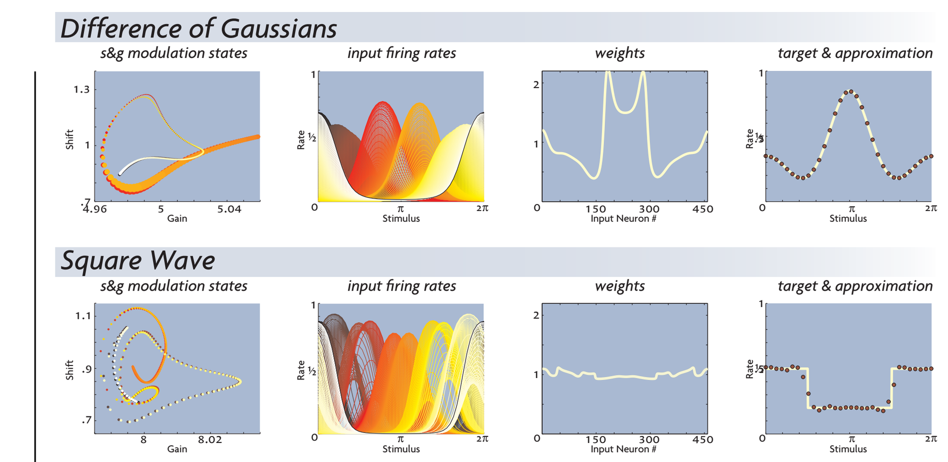

Response Modulation coupled with Hebbian plasticity can approximate a wide variety of smooth, continuous functions. However it is limited by the Gaussian tunings of the input cells. This is most apparent in cases where the target function's rate of change exceeds the maximal slope of the modulated gaussians.

Multiple Outputs, Interference, and Bias

Function Approximations with Two Output Units

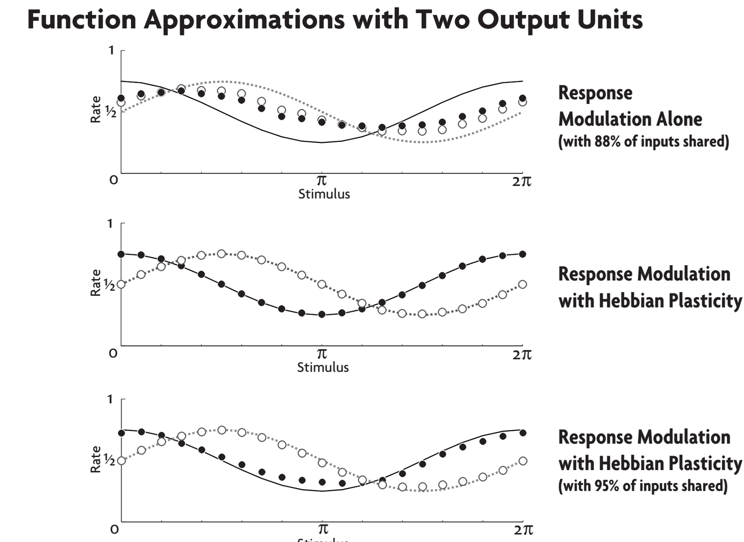

In networks with more than one output, each input cell may project to multiple target neurons—complicating the supervisor's task. In the two-unit case, there must be a sufficient number of uniquely connected input neurons to allow for different target functions to be approximated.

However, even in cases with too much overlap for the supervisor alone to impose a proper activity pattern, the relatively weak bias it provides can direct Hebbian plasticity to approximate the two target functions.

Function Approximations with >2 Output Units

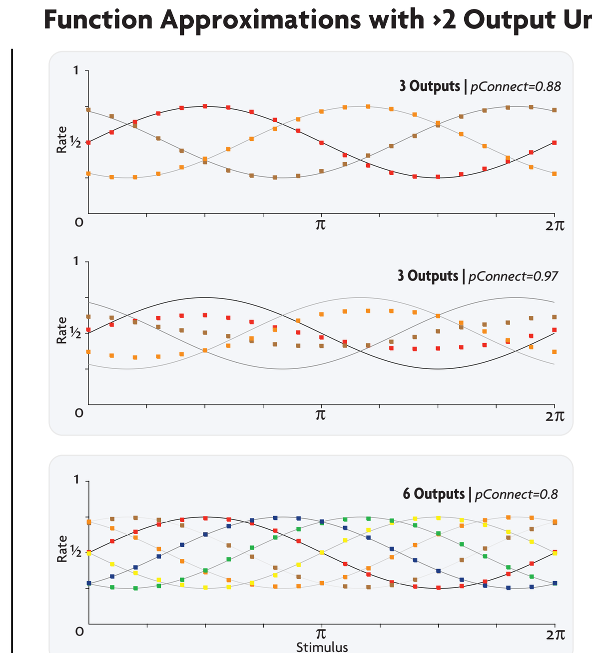

Extending the model beyond two outputs shows that it is in fact *sparseness* of connections that is the controlling factor—totally *unique*, non-divergent inputs are not required.

For these simulations, networks were constructed stochastically, with each possible feedforward connection being created with a fixed probability.

Interestingly, network performance was primarily predicted by this probability of connection factor, and was essentially constant across different input- and output-unit population sizes, as well as number of target functions to be approximated.

Results suggests that:
(a) this approach could scale to large networks, and
(b) connections can be made randomly as long as the probability of connection remains below a critical value of approximately 0.95.

Error vs. Number & Proportion of Unique Inputs

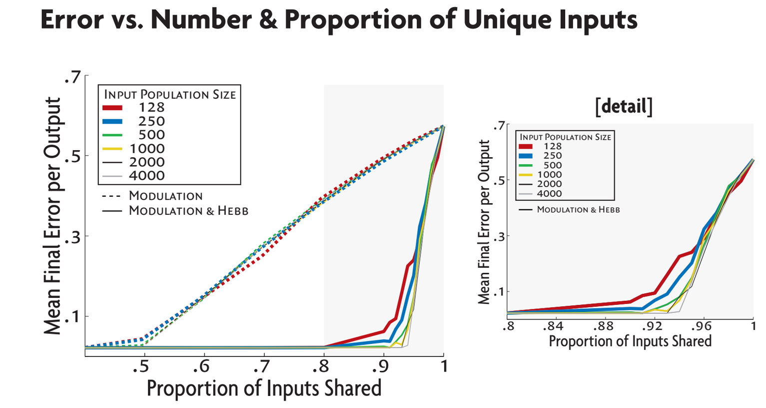

Error vs. Number of Output Units

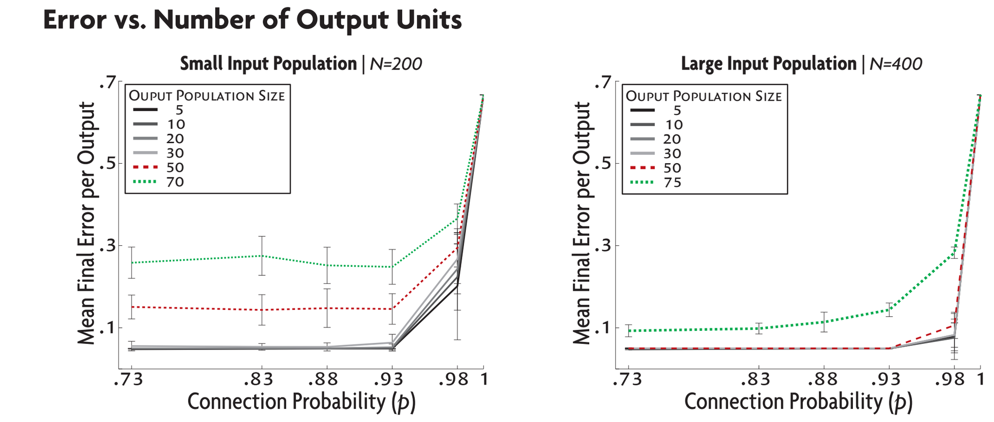

Nature of the 'Supervisor'

Random Walk/Reinforcement Learning Supervisor

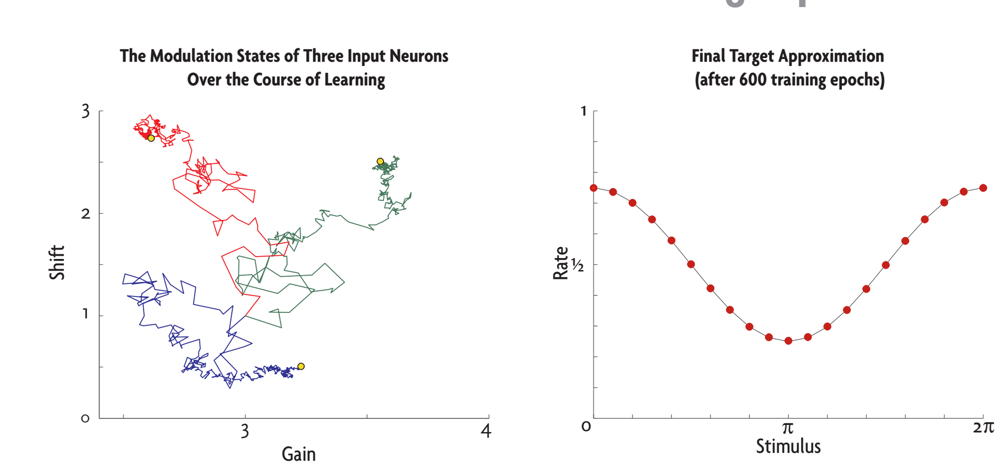

References

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)*.