The Journal of Young Investigators
The premier undergraduate research journal

Skewing the Perceptron: Modeling the Importance of System-Level Manipulation When Investigating Long-Term Potentiation

Andrew Murray, Jacob Matz
Dalhousie University
Department of Physiology and Biophysics

Advisor: Dr. Thomas Trappenberg

Volume 19, Issue 22
April 2010

Skewing the Perceptron:
Modeling The Importance of System-Level Manipulation When Investigating Long-Term Potentiation

Andrew Murray[1], Jacob Matz[1]
Correspondence: Andrew Murray, amur19@gmail.com
[1] *Dalhousie University*

Advisor: Dr. Thomas Trappenberg

Abstract

Manipulation of learning processes in the brain has proven to be experimentally challenging. Various studies have focused on specific components of Long-Term Potentiation (LTP), attempting to systematically regulate the learning progress through manipulation of its proposed elements. The manipulation of such elements is thought to result in a proportional change in the efficacy of the whole system. Based on the inadequate results of this approach, we hypothesized that changes implemented to a fraction of synapses in a network represented an incomplete investigation of LTP function. We further hypothesized that in order to successfully manipulate learning output, whole systems must be manipulated whether working with a single synapse or a neural network. We chose a relatively simple and easily applied model, the perceptron, to simulate the effects of strengthening synaptic connections between neurons, as in learning. We expected that following the learning phase, systematic changes to a fraction of synaptic weights would jeopardize the perceptron's ability to recall the learned patterns. Indeed, fractional changes of the weight distribution caused more errors in the retrieval of patterns. However, when the entire distribution was altered, error in pattern retrieval was absent.

Introduction

The connections between neurons have been shown to grow in strength (i.e. synaptic weight) in response to long trains of high frequency stimulation, or the temporal pairing of two or more stimulating signals (Kessels and Manilow, 2009; Bliss and Lømo, 1973). Such an input-dependant change in the properties of neuron communication is the definition of Long-Term Potentiation (LTP), a leading theory of memory formation and recall in the brain. LTP involves many complex cellular steps, coordinated between the pre- and postsynaptic neurons in ways that are yet not fully understood. Following induction of LTP, synapses exhibiting the increased weight values characteristic of this type of learning tend to transmit signals more readily when stimulated. A higher firing probability is subsequently observed throughout the affected pathways (Blundon and Zakharenko, 2008; Smolen, 2007). The changes in pathway activity resulting from LTP represent the specific tuning of neural response to a familiar (or learned) stimulus.

Studies have shown that neurons within the Medial Temporal Cortex of the mammalian brain are more responsive to familiar stimuli, showing either an increase or a decrease in their firing efficacy (Squire et al. 2009). Although this is good evidence for learning, it is not yet clear how this is manifested in memory retrieval. We can however speculate from this evidence that cells which have undergone increases in synaptic weight (as in learning) do develop recall ability, as a function of the increased synaptic strength.

Various studies have attempted to access the effects of altering synaptic strength via genetic and pharmacological manipulations in an effort to change either the outcome of the learning process, or retrieval ability following learning (Shunsuke et al, 2009; Tao-Cheng et al, 2001; Eichenbaum, 1996). Although LTP is a multicomponent process, these studies have suggested the presence of key constituents in synaptic weight increase. Experimenters have labeled various cellular steps as rate limiting steps, and attempted to either up- or down-regulate their activity to cause associated changes in memory systems both *in vivo* and *in vitro*. These studies have lead to very few conclusive results, suggesting that the underlying assumption, that singular key components exist within LTP systems, is inaccurate.

The single-layer perceptron (Rosenblatt, 1958) allows the basic modeling of a neural network, similar to those found in the Medial Temporal Cortex. It is a theoretical model capable of simulated learning of associations between independent patterns by adjusting the connections of its network in a way representative of LTP. The network connections are collectively represented by a matrix of weight values, each value representing connection strength. The relative simplicity of this model in theory and application, together with its amenability to experimental manipulation, makes it a good platform from which to examine general trends in neuronal connectivity. Although it has been shown that the capabilities of feed-forward learning networks are limited (Minsky&Papert, 1969; Malinowski et al. 1995), trends which can be systematically isolated in the perceptron are valuable measures of influence over basic learning (Elizondo et al. 2007).

Targeted manipulation of values in the weight matrix largely parallels previously mentioned biological experiments. Approaches in living systems have targeted synaptic weights through specific components of LTP, and in select cells within a network. Such manual increases in localized synaptic strength are presumably beneficial to the network, seeing as the overall efficiency of the system would be expected to increase. However, by altering a fraction of a network's weight matrix by a constant factor, we expected to simulate the incomplete strengthening of a network and produce effects which mirror those observed in living systems.

Using the perceptron program as a basis for supervised learning and subsequent retrieval of learned information, we hypothesized that experimental changes to the distribution of weight values after learning will cause associated changes in the state of the perceptron. We expect that these changes will reflect decreases in the efficiency of pattern retrieval.

Materials and Methods

In our configuration of the perceptron model (Figure 1.), input patterns are represented in groups of binomial values (0 and 1). Each input pattern is processed through the weight matrix during its transmission from an input node to an output node (nodes equivalent to pre and post-synaptic neurons). By this process it receives a strength value, assumed to be defined by the pattern's familiarity (and later, its ability to be recalled). A total of 20 input patterns were multiplied across the weight matrix connecting the input and output nodes. The initial weights assigned to the weight matrix were randomly selected from a uniform distribution between -0.5 and +0.5. This random selection of values results in a Gaussian distribution of weights as suggested by the Central Limit Theorem. A fixed-threshold function was used to produce a binomial output from the range of values emerging from the weight matrix. The desired output of the perceptron is chosen by the experimenter and is independent of the input. The learning process takes place by comparing the input pattern to the desired pattern and developing the weight matrix according to the Delta rule formula so that the matrix may be tailored to link the input with the output. In this way the model is able to "learn" associations between independent stimuli solely by changing the strength of connections between its nodes (comparable to LTP-based learning).

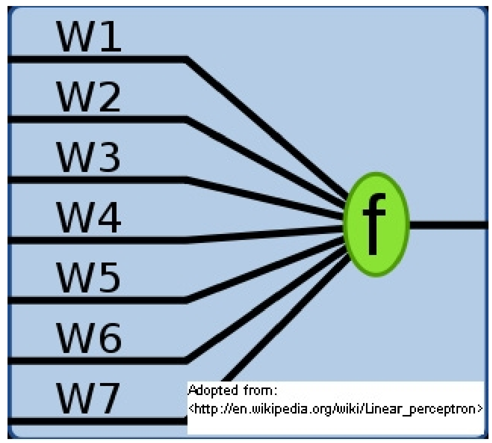

Figure 1. Schematic model representing the processing of 5 input patterns (*in*) arising from a series of 5 input nodes. The input patterns (composed of binary vectors) are compared to a predetermined "desired" output. Weight values (w*n*) are then applied to the inputs through algorithm in an effort to yield an output pattern similar to the desired output. The perceptron is able to learn patterns by adjusting these weights in order to render the desired and actual outputs identical. The memory of the input or stimulus is thereby stored in the array of updated weight values.

The error rate observed in perceptron performance is defined as the percentage of incorrectly retrieved patterns over many successive trials. The update (ΔWj) required to resolve the error between the present and the desired states of the perceptron was computed, and gradually resolved using the Delta rule:

$$\Delta Wj = \alpha(Uj - Yj)Xi$$

The learning rate (α) was chosen as 0.1 in accordance with previous protocol (Rojas, 1996). The rule requires the distance between the desired output (*Uj*) and the present state (*Yj*) be multiplied by both α and the i^{th} input (*Xi*). Delta rule updating is characterized by a gradient descent in the error rate of the perceptron as it learns a new pattern.

Following primary training of the perceptron we made systematic increases (factor of 2) to the “learned” weight matrix by initially altering 100 weight values (out of total 4000=2.5%) and increasing the number of changed weights by increments of 100, up to the point of altering the entire matrix. We operationally defined each stepwise increase in the number of altered weights as one incremental condition, for a total of 40 conditions. Following each change to the learned matrix, we tested the recall ability of the perceptron.

To measure the ability of the perceptron to recall the information it had learned, we presented the initial patterns and processed them through the transformed weight matrix, generating an output. The deviation of the generated output was compared to the desired memory and recorded as error rate. Mean error rate was plotted for the 100 iterations in the 40 conditions.

All experiments were performed within a network consisting of 20 input nodes and 20 corresponding output nodes. These were fully interconnected by 4000 modeled synaptic connections.

Results

Initial perceptron learning took place successfully and the program was able to store all associations between the patterns introduced and desired outputs within the weight matrix. The 40 incremental changes to this established matrix then produced a function of mean error rates reminiscent of a parabolic shape. As illustrated by the curve in figure 2, the program begins at a state of perfect recall, meaning that the perceptron was able to associate all previously learned patterns. Therefore no errors in recall were evident when the original weight matrix was used (number of nodes altered=0).

As changes were introduced to an increasing number of weights, a rapid decrease in the perceptron’s retrieval ability became evident. The decrease in recall ability plateaued when approximately a third of the nodes were changed (~1300/4000=32.5%). As the number of changed weights approached approximately 50% of the total matrix, an improvement in the recall ability began to take shape. By the time all 4000 weights had been adjusted, there was an error rate of zero.

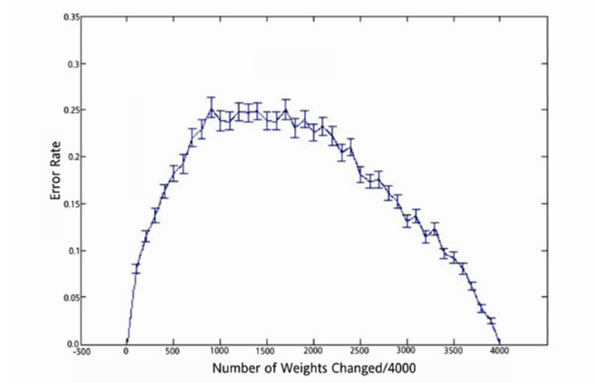

Figure 2. Modeling the error in retrieval of stored patterns as an increasing number of weights (0%-100%) is modified. *X*=0 represents the point at which no weight values have been manipulated. As the x-axis advances, an increasingly large proportion of the weight matrix is doubled. Although 0% and 100% manipulation does not result in any error in the system, intermediate proportions do cause the system to falter (maximum error at approximately 32.5% manipulation). As a larger portion of the weight values are increased, system error declines until the point of 100% manipulation, where no error is observed.

Discussion and Conclusion

The principal findings of this study were (1) Manipulating a fraction of the weight matrix resulted in errors in perceptron retrieval of learned patterns (2) Errors in perceptron pattern retrieval were maximized when 35-50% of all weights were increased (3) Beyond this point, recall performance improved as a larger fraction of synaptic weights were changed (4) When all weight values were adjusted, the error rate returned to zero.

The occurrence of initial weights is Gaussian distributed (or normally distributed) as a result of Central Limit Theorem implication. As these values undergo Delta rule-mediated updating throughout the learning process, a new matrix is conceived which exhibits a specific distribution in accordance with the stimulus learned.

In our experiment, we illustrate that the shape of this final weight distribution plays an important role in observed error during pattern retrieval. When the weight matrix is unaltered after learning (0%), the occurrence of each weight value can be expressed with a normal distribution, resulting in Gaussian distributed weight values. Interestingly, we discovered that multiplying every value of the weight matrix (4000/4000=100%) by a constant did not result in increased error. As such, 0% and 100% manipulation both result in similarly distributed weight values and zero error, consistent with perfect association between input patterns and target output. Conversely, when a proportion of the matrix (0% < χ < 100%) was multiplied by a constant, the shape of the weight distribution began to change and was associated with increased error. The error rate then declined as a larger proportion of the weights were increased (allowing the initial distribution shape to reemerge). Together, these points illustrate the importance of maintaining shape of the weight distribution established in the learning phase, throughout memory retrieval. As well as the tendency of the system to falter when the distribution of weights is experimentally skewed. We have therefore shown that even when the values of the entire weight matrix are significantly altered, the pattern retrieval performance is perfect, provided that the weight distribution is reflective of previously determined associations.

Such a result may be explained through examination of homeostatic-like mechanisms within the program. When the perceptron is trained, a level of balance is established in the weight matrix, and can be conceptualized to lie in the shape of its distribution, rather than the values themselves. It also presents with a mean value of synaptic strength, perhaps traceable to pathway function. This mean weight value is flanked by more negative and positive values drawn out by the learning phase, specific to the network, and the task learned. Disruption of this shape may lead to the error rate observed. This can be illustrated in a simplified model as shown in figure 3.

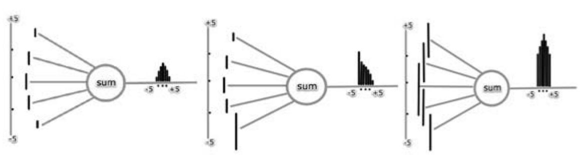

Figure 3. Schematic model of perceptron function using inputs weighted randomly according to a uniformly distributed weight matrix. The occurrence of various weight values is represented by bold vertical bars associated with inputs and the Y-axis shows the distribution of these values around the mean. The output node produces a pattern guided by the weight values.

Given the model's apparent dependence on the shape of its weight distribution in order for proper recall to take place, perhaps we can derive inferences concerning these processes in living systems. The importance of maintaining the array of weights established in learning throughout the testing of recall can be examined at a number of levels of exploration. Such alterations of specific network connections at a biomolecular level, for example, could possibly unbalance the cascade of cellular events associated with LTP. Such a disruption could potentially cause significant alterations in learning and/or recall ability; a result which may even be observable at a psychophysical level. Experimental manipulations to specific components of Long-Term Potentiation, or specific cells within a network, may disrupt the physiology of the system. According to data presented here, such a localized imbalance has potential to cause system failure in the perceptron and these results are indeed observed in many living systems following the manipulation of select steps of the memory formation process (as previously described).

The solution for perceptron learning, and possibly neuronal learning as well, is to implement changes to the system as a whole rather than to specific parts. It is evident that synaptic weights are required to change to enable successful learning. We also observe, however, that isolating select weight values for increase is detrimental to the network's ability to learn. From a biological standpoint, we can speculate that system level investigation must be carried out through means which incorporate the entire brain area or perhaps the entire brain or organism. Memory-related investigation of attention may fulfill these requirements.

Conversely, perhaps such living systems are much more versatile than the perceptron concerning synaptic strengths within a network. In this case the heightened sensitivity to balanced input may hold the perceptron back from its application to real world learning.

Many other manipulations of the weight matrix are conceivable. With them exists a chance to clarify the degree of influence each element of the weight matrix might carry. For example, an interesting approach may be to selectively change or eliminate the values of either positive or negative weights, and observe the effect on retrieval error. Based on the approximate symmetry of our data, one would expect that the positive and negative weights would carry equal influence (although this might not hold true in some more complex systems). Also, it would likely be beneficial to solely change the small weights in the matrix (either positive or negative). Under the presumption that these weights carry little effect on the overall connectivity, changes in these weights would not be expected to produce large differences in pattern retrieval.

In this study of simulated learning and recall, we have demonstrated the properties and importance of the weight distribution when working within a perceptron model. The experiments proposed above are intended to further qualify these findings, and shed more light on this feature of perceptron recall exposing what may either be an applicable trait, or strict limitation of said model.

Acknowledgements
We are extremely grateful to Dr. Thomas Trappenberg for his patience, support, advice, and instructive guidance. Thanks to the Dalhousie University Dept. of Computer Science for use of the facilities.

References

Baxter, M.G. (2009) Involvement of medial temporal lobe structures in memory and perception. *Neuron.* **61**(5): 667-77.

Bliss, T. and T. Lømo (1973) Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *Journal of Physiology*. **232:** 331-56.

Block, H.D. (1962), The perceptron: A model for brain functioning. *I. Reviews of Modern Physics* **34**:123-135.

Blundon, J.A. and S.S. Zakharenko (2008) Dissecting the components of long-term potentiation. *Neuroscientist.* **14** (6): 598-608.

Eichenbaum, H. (1996) Learning from LTP: A comment on recent attempts to identify cellular and molecular mechanisms of memory. *Learning and Memory.* 3: 61-73.

Elizondo et al. (2007) Analysis and test of efficient methods for building recursive deterministic perceptron neural networks. *Neural Networks* **20**: 1095-108.

Kessels, H.W. and R. Manilow (2009) Synaptic AMPA receptor plasticity and behavior. *Neuron.* **61** (3): 340-350.

Malinowski, L. et al. (1995) Capabilities and limitations of feed-forward neural networks with multilevel neurons. *IEEE.* **(1)**: 131-134.

Minsky, M.L. and S. Papert (1969) *Perceptrons*. MIT Press.

Qiuzhen, X.H.H. and W.J. Tompkins (1991) Structural simplification of a feed-forward, multilayer perceptron artificial neural network. *IEEE.* **2.5**: 1061-4.

Rojas, R. (1996) *Neural Networks*, Spinger-Verlag. p. 77-99.

Rosenblatt, F. (1958) The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*. **65** (6): 386-408.

Scoville, W.B. and B. Milner (1957) Loss of recent memory after bilateral hippocampal lesions. *Journal of Neurology, Neurosurgery and Psychiatry.* **20:** 11-21.
Shunsuke et al. (2009) Transgenic up-regulation of alpha-CamKII in forebrain leads to anxiety-like behaviors and aggression. *Molecular Brain.* **2**:6.

Smolen, P. (2007) A model for late long-term potentiation simulates aspects of memory maintenance. *PloS ONE*. **2** (5): e445

Squire, L.R., J.T. Wixted and R.E. Clark (2007) Recognition memory and the medial temporal lobe: a new perspective. *Nature.* **8**: 872-883.

Tao-Cheng, J.H. et al. (2001) Sustained elevation of calcium induces calcium/calmodulin-dependant kinase II clusters in hippocampal neurons. *Neuroscience.* **106** (1): 69-78.

Supplementary Materials

The perceptron (20 input/output nodes with 200-unit patterns) converges to a solution: `(Supplement.jpg)`

Program for the Error Rate vs. number of altered nodes curve:

```
clear;
perceptron;
n=0:100:4000;

for u=1:41;
d=n(u);
if d < 1;
    rOut=wOut*rIn>0;
    distzero=sum(sum((rDes-rOut).^2))/26;
    error1=distzero;
else
    for c=1:100;
        x=randperm(4000);
        select(:,c)=x(1:d);
    end
    for y=1:100;
        wOut(select(:,y))=
wOut(select(:,y))*2;
        rOut=wOut*rIn>0;
        distH1=sum(sum((rDes-rOut).^2))/26;
        error1(y)=distH1;
        wOut(select(:,y))=wOut(select(:,y))/2;
    end
    int(u)=mean(error1);
    dev(u)=se(error1);
clear select
```

```
end
end

errorbar(n, int, dev)
xlabel('Number of Increased Weights (/4000)');
ylabel('Average Error Rate');
```

Programmed error measured as “Hamming distance” ;
n correcting substitutions required (ie. distH1)/n total values =
proportion of error reported