

UNIVERSITY OF CALIFORNIA SAN DIEGO

Evaluating Modern STDP Rules and Their Role in Embedding Information in the Brain

A Dissertation submitted in partial satisfaction of the requirements
for the degree Doctor of Philosophy

in

Bioengineering

by

Vikash Nagin Morar

Committee in charge:

Professor Gabriel A. Silva, Chair
Professor Maksim Bazhenov
Professor Marcus Benna
Professor Bruce Wheeler

2023

Copyright

Vikash Nagin Morar, 2023

All rights reserved.

The Dissertation of Vikash Nagin Morar is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2023

TABLE OF CONTENTS

DISSERTATION APPROVAL PAGE.....iii

TABLE OF CONTENTSiv

LIST OF FIGURES v

LIST OF TABLES vi

ACKNOWLEDGEMENTSvii

VITA..... xi

ABSTRACT OF THE DISSERTATION.....xii

Chapter 1 Synaptic Plasticity: From Neuron to Neural Network to Robot..... 1

**Chapter 2 A Computational Model for Storing and Embedding Data in Dynamically
Evolving Network Synaptic Structures 59**

Chapter 3 Simplifying Biology Vocabulary via Morphology 100

CONCLUSION 110

LIST OF FIGURES

Figure 2.1 Overview of the network model and analysis methods. (A) We used the leaky-integrate-and-fire model of a neuron to construct networks with a time delay (temporal latency) and weights for each edge. MNIST images consisting of 784 pixels were fed into the input layer with each input neuron corresponding to one pixel. Each input layer neuron 63

Figure 2.2 Biologically-inspired STDP rules we used in our networks. If an excitatory signal from a presynaptic neuron arrives at a postsynaptic neuron before it fires, the weight of the edge between them increases. If the signal arrives at the postsynaptic neuron while it is refractory, the weight of that edge decreases. Inhibitory edges are modified with a different 65

Figure 2.3 Qualitative characteristics of the model. (A) Spike raster for the first 30 activations for each neuron in the network. The activity of the network starts slowly due to the delays from the input nodes, but quickly increases. In addition, even with just 200 recurrent nodes, there are approximately as many as 20,000 activations per second in the simulation. As..... 72

Figure 2.4 Summary of KNN classification accuracies over time. We explored how different parameter values impacted performance. (A) We tested various recurrent layer sizes to balance the classification accuracy with simulation efficiency. The peak classification accuracy improved as the size of the network grew, reaching a constant at 300 ms of 74

Figure 2.5 Testing the individual contributions of potentiation and depression on edge weight changes and classification accuracies. Depression alone was sufficient to maintain high classification accuracy, but potentiation by itself reduced the peak accuracy substantially. We then switched the parameter values for each learning rule to see if they were the cause, but . 77

Figure 2.6 Increased and decreased refractory periods of neurons within the recurrent layer suggests a quantifiable reason for the scale of biological refractory periods. Regardless of how large or small the refractory period is, with both potentiation and depression changing weights, the classification accuracy remains consistent. However, with just depression 78

Figure 2.7 Distances of nearest neighbors to a single image based on their raw pixel values sorted to be ascending (red lines). Over the course of the simulation, we used this ordering to see how it compared to distances of the same image based on the network edge weights while preserving the order from the raw pixel values. The running mean of the following nearest.. 80

Figure 2.8 (A) We compared the classification accuracy performance on an MNIST data set between our biologically-inspired neural network and an artificial neural network using a similar edge weight KNN method. We see that our network achieves better classification accuracy regardless of the number of neurons within the recurrent/hidden layer. (B) When .. 81

LIST OF TABLES

Table 2.1 Each of the parameters chosen for the latencies in the network and the STDP rules was taken from murine hippocampus experimental data. The only exception is the edge weight range and threshold potential, which were chosen to make the computations manageable.	68
--	----

ACKNOWLEDGEMENTS

I would like to acknowledge Professor Gabriel Silva for his support as my advisor for these past few years and as the chair of my committee. I will forever appreciate him taking a chance on me with my limited programming experience and through that, I've been able to learn an immense amount in graduate school. Professor Silva letting me take the time to teach over the years has led to me growing immensely as an academic and as a person. Without those opportunities, I may not have discovered my passion for teaching and chosen to pursue the career path that I have.

I would like to also acknowledge Dr. Vivek George for his unwavering support and guidance throughout the program as well. Whenever I needed help troubleshooting code or was looking for advice on how to navigate life, especially during the pandemic, he was always warm and welcoming. Without his support, I would not have been able to pick up nearly as much programming knowledge as I did.

Next, I would like to acknowledge Dr. Alyssa Taylor for her generosity spending so much of her time meeting with me to discuss pedagogy and career advice. Without her belief in me, I would not have been nearly as confident as an instructor and I would not have begun looking at career opportunities as early as I did. Having her as a role model gave me invaluable perspective into the responsibilities of a teaching professor. Beyond that, her commitment to service and extracurricular committees concretized my aspiration for that career path for myself.

I would like to acknowledge Dr. Adam Engler for fostering my interests in teaching as department chair even when he did not have to. He took the time to meet with me and make exceptions when appropriate to let me practice teaching beyond the normal requirements for

PhD students. With his blessing, I was even able to teach my own class on behalf of the department, which was an amazing experience and also critical for my CV. On a personal and professional level, I'll always be grateful for him trusting me with that opportunity.

While ambiguous, I would also like to acknowledge all of the students that I had the privilege of teaching over the years. I learned an immense amount about pedagogy, empathy, and privilege from those experiences and they mean the world to me. Many of my students were incredibly sweet and I sincerely appreciate their trust in me to teach them well and with compassion.

On the personal side, I would like to acknowledge my parents and older brother, as they constantly tried to shield me for the duration of the program. Their priorities shifted to doing whatever was necessary so I could focus on my degree without external distractions. Their effort did not go unnoticed. With their help, I was able to spend as much time as I could on research and teaching without having to worry about too many other externalities. The privilege and stability they allowed me to have cannot be understated and I will never take it for granted.

Lastly, I want to acknowledge all of the friends who have supported and inspired me along this journey. To Eti and Katherine, having a virtual cohort of us three PhD students at three different institutions has been refreshing and grounding in many ways. Even when we navigate “qualifying exams” meaning completely different things and some of us having “data meetings” and others not, the solidarity is powerful and comforting. I never expected to be the first to defend, but I know your times will come sooner than you think and your dissertations will turn out incredible. Even now, we're still in this together. To Phoebe, your support as my best friend has been more than I could ever deserve. From the initial moment

when I was accepted to my transition from MS to PhD even to me now graduating, your excitement on my behalf is humbling. Truly, without you, I would never remember to celebrate even the small victories and I would get caught up in chasing each new challenge forever. Your perspective well-balances my flair for the dramatic and I hope to give you many more things to be excited about in the future. To Jonathan, Connie, and Dexter, I have always appreciated our chats in particular since they have truly been an escape from the labors of graduate school. It can sometimes feel cleansing to spend so much time with three people who have never pursued research to begin with. Hearing about normal office drama is something of a panacea after your fourth consecutive experiment has gone awry. To Jessica, you've truly been an inspiration throughout these years and I am so proud of all you accomplished. Things have never been easy for you and yet you managed to overcome constantly. I've aspired to have your resiliency, work ethic, and commitment and hopefully, as time goes on, I'll keep getting closer. You can also forever hold it over me that you received your doctorate before me. To Rachael, my longest tenured friend, I certainly don't have a choice but to include you. I find your struggle and tenacity as a teacher quite remarkable. While we teach extremely different subjects, I do sincerely hope to be like you when I have more classrooms of my own. I appreciate you always praising my capabilities even when I am critical of myself. I, of course, love few things more than our talks about pedagogy. To Angelica, I am so grateful for the unwavering enthusiasm you've brought me throughout even the hardest parts of this program. No matter what I'm struggling with, I can't help but feel better after talking to you and feeling the strength of your positivity. There are few people who really make me feel as appreciated as you have and that is the energy and sentiment that I try to channel right back into my students and peers. To conclude, there are even more people that I'd love to

acknowledge (JW, VX, RZ, IZ, SG, CU, RT, JH, KT), but I don't have enough space or time to keep going. All of your impact on me has been no less meaningful and I think you all deserve acknowledgement as well. I could not have done this without all of you.

Chapter 1, in full, is currently being prepared for submission for publication of the material. Morar, Vikash; Silva, Gabriel A. The dissertation author was the primary author of this material.

Chapter 2, in full, has been submitted for publication of the material as it may appear in Nature Scientific Reports, 2023, George, Vivek K.; Morar, Vikash; Silva, Gabriel A., Nature Portfolio, 2023. The dissertation author was one of two primary researchers and authors of this paper.

Chapter 3, in full, is a reprint of the material as it appears in The American Biology Teacher, Morar, Vikash, University of California Press, 2023. The dissertation author was the primary author of this paper.

VITA

2016 Bachelor of Science in Biochemistry and Cell Biology with Minor in Mathematics, University of California San Diego

2016 – 2017 Lab Assistant, University of California San Diego

2017 – 2019 Bio-technician, Diazyme Laboratories, Inc.

2023 Associate (Teaching A Course), University of California San Diego

2023 Doctor of Philosophy in Bioengineering, University of California San Diego

PUBLICATIONS

George VK*, **Morar V***, Silva GA. (In Review). A Computational Model for Storing and Embedding Data in Dynamically Evolving Network Synaptic Structures. *Nature Scientific Reports*.

Guerra SP, George VK, **Morar V**, Roldan J, Silva GA. (In Review). Extending undirected graphs techniques to directed graphs via Category Theory. *Theory and Applications of Graphs*.

Morar V. (2023). Simplifying Biology Vocabulary via Morphology. *The American Biology Teacher*.

George VK*, **Morar V***, Silva GA. (2022). A Computational Model for Storing Memories in the Synaptic Structures of the Brain. *bioRxiv*.

Gomez-Godinez V, **Morar V**, Carmona C, Gu Y, Sung K, Shi Z, Wu C, Preece D, Berns M. (2021) Laser induced shockwave paired with FRET to study neuronal responses to shear stress. *SPIE 2021 Nanoscience + Engineering Conference*.

George VK, **Morar V**, Yang W, Larson J, Tower B, Mahajan S, Gupta A, White C, Silva GA. (2021). Learning without gradient descent encoded by the dynamics of a neurobiological model. *arXiv*.

Gomez-Godinez V, **Morar V**, Carmona C, Gu Y, Sung K, Shi Z, Wu C, Preece D, Berns M. (2021). Laser-Induced Shockwave (LIS) to Study Neuronal Ca²⁺ Responses. *Frontiers in Bioengineering and Biotechnology*.

ABSTRACT OF THE DISSERTATION

Evaluating Modern STDP Rules and Their Role in Embedding Information in the Brain

by

Vikash Nagin Morar

Doctor of Philosophy in Bioengineering

University of California San Diego, 2023

Professor Gabriel A. Silva, Chair

Pattern recognition is a core facet of learning any concept, whether it be a literal pattern in an image or a predictable stimulus-response pattern as in operant conditioning. From the molecular mechanisms up to the behavioral consequences, patterns can be embedded to allow humans to learn from past experiences to better predict future outcomes. A key player in this embedding process is the spike-timing-dependent plasticity (STDP) of the connections between neurons within the brain. STDP is theorized to leverage the spiking

activity of neurons resulting from various stimuli according to the initial rule proposed by Donald Hebb: neurons that fire together wire together. However, the actual process by which STDP learning takes place remains disputed, especially in cases such as inhibition, where neurons do not actually support subsequent firing. This dissertation tackles a few of these issues from multiple angles: (1) An evaluation of many different proposed STDP rules and how they have been justified experimentally and employed computationally. (2) The creation of a computational model that highlights the value of the STDP-mediated embedding of a stimulus into the architecture of a spiking neural network in a manner that could resemble brain embeddings. (3) A proposed approach to teaching students by taking advantage of the natural tendency of the brain to identify patterns and use them to create connections in new contexts. Here, we employ a biologically-inspired spiking neural network model with edge latencies, refractory periods, and modern STDP rules to demonstrate the potential of the brain to embed information within its architecture via synaptic weights. These embeddings may be quite relevant in the brain where the astrocyte syncytium may be able to access them to communicate local learning information on a more global scale.

Chapter 1 Synaptic Plasticity: From Neuron to Neural Network to Robot

1.1 Introduction

Learning and memory are phenomena of great interest to both neuroscience and machine learning. On the biological side, the ongoing electrical activity within the brain culminating in perception and recognition leads to numerous theoretical questions and experiments. With computer science, the quest to design smarter, more capable algorithms leads many to assess which features and structures lead to efficient computation on designated hardware. These two fields navigate similar inquiries with different goals, but there is much overlap. One area where both fields can mutually benefit from each other is neural computation. By investigating how the brain processes information, machine learning models can attempt to match its computational efficiency. Similarly, by using computational models, complementary functions within the brain can be assessed in a way that current experimental methods are unable to isolate^{1,2}.

A key aspect of neural computation is the structural changes that take place to produce necessary computational modifications³. These changes, which occur at the synapses between neurons in the brain, have a storied history of being critical to learning⁴⁻⁶. Decades ago, synaptic plasticity was challenged as a method of learning based on four fundamental criteria: whether synaptic changes were detectable, if they could be mimicked to induce false memories, if preventing synaptic changes prevented memory formation, and if synaptic changes could be retroactively undone to remove memories⁷. The first criterion, detectability, has been easily achieved with modern imaging technologies *in vivo* as mentioned previously as well as *in vitro*⁸⁻¹⁰. Numerous experimental studies have demonstrated the importance of structural changes within the brain for auditory, visual, and motor function and recognition¹¹⁻¹⁴. To meet the second

criteria, false fear memories have been created through optogenetics in murine experiments as well^{15,16}. For the third, several studies have demonstrated that by blocking the molecular mechanisms of synaptic plasticity, learning is inhibited as well¹⁷⁻¹⁹. Lastly, one study has even demonstrated memory loss as a direct effect of dendritic spine loss in the murine hippocampus²⁰. Another has correlated the lifespan of episodic memory with the lifespan of hippocampal dendritic spines²¹. With mounting evidence that structural changes tie directly to memory formation and learning, there becomes greater need to investigate the mechanism at work.

Spike-timing-dependent plasticity (STDP) was originally proposed by Dr. Donald Hebb through his theory often paraphrased as “neurons that fire together wire together”²². Hebbian STDP claims that if one neuron firing leads to the subsequent firing of another, the connection (synapse) between the two neurons should be strengthened, as it seems to be purposeful and functional. Extended further, the magnitude of the change can be quantified as a function of the time between the spikes²³. There have been many different proposed rules for STDP, either justified through experimental data or computational performance^{24,25}.

This review serves to investigate the neuroscientific bases and highlight the computational utilities of these various rules. In addition, neuromorphic and robotics applications will also be assessed, as many have taken inspiration from biological and artificial neural network implementations of STDP²⁶.

1.2 Experimental Evidence for Hebbian STDP

For a discussion of STDP rules, the curve to begin with is the classic hyperbolic cotangent curve experimentally observed in murine hippocampal neurons^{27,28}. This curve, sometimes referred to as “Hebbian STDP” is the simplest interpretation of Hebb’s theory.

According to Hebbian STDP, if a presynaptic neuron fires prior to a postsynaptic neuron, the synapse between them strengthens as a function of the delay between the firings (known as long-term potentiation or LTP). In contrast to this, if the postsynaptic neuron fires first, the synapse is weakened as a function of the time between those firings (known as long-term depression or LTD). As such, each synapse is evaluated via the pairs of firings that take place between any two adjacent neurons. Shorter delays result in larger changes in either direction depending on which neuron fired first. Further experimental data for this rule has been observed in murine pyramidal cells by observing the increases in excitatory postsynaptic potentials (EPSPs) resulting from corresponding activity²⁹.

To dive deeper into the mechanism behind this STDP rule, studies have investigated the different biochemical pathways by which these structural changes can take place. One hybrid experimental computational study assessed the impact of noise on N-methyl-D-aspartate(NMDA)-mediated STDP when compared to endocannabinoid-mediated STDP³⁰. The study concluded that endocannabinoid-STDP appeared to be more robust to noise until the activity of the neurons increased past a certain threshold. This observation demonstrates that redundancy can be quite beneficial in the brain. With two different pathways to implement STDP, neurons can respond appropriately regardless of the situation.

When trying to model STDP rules in the brain, it is quite uncommon to observe isolated spike pairs. To address the realism of the Hebbian STDP rule, one group monitored murine visual cortex slices to track the activity-induced plasticity that resulted from the spike trains³¹. While they were still assessed as pairs, as the rule requires, the spikes measured were all part of larger contingents. This allowed for a more realistic scenario for plasticity to occur within. The

observed synaptic changes followed the proposed Hebbian STDP rule, strengthening the evidence that this rule could be accurate, at least in the visual cortex.

As technology has advanced, the Hebbian STDP rule has been investigated down to its molecular mechanisms. One group employed two-photon microscopy to investigate the relationship between STDP-mediating NMDA receptors and actin polymerization in murine pyramidal neurons³². As it has been proposed that structural changes to dendritic spines correspond to synaptic learning, this study aimed to identify the biological mechanism behind the connection. Interestingly there was variable behavior depending on how close different dendritic spines were to each other. Along with the different biochemical mediators of STDP, it seems that spine proximity and size may also play a role in how predictable neuronal plasticity may be.

1.3 Computational Models Incorporating Hebbian STDP

1.3.1 Data Modeling

As Hebbian STDP is based off of experimental observations, a natural implementation for it would be to model brain activity data. To this end, one group implemented Hodgkin-Huxley neurons with a classic Hebbian STDP rule to model brain activity in non-rapid eye movement (NREM) sleep³³. The study showed that sleep activity leveraged STDP to encode memories via a spike-timing based representation for long-term storage.

Another group designed the NeuCube, which is a spiking neural network with an adapted STDP rule designed to analyze brain data, starting with electroencephalogram (EEG) data³⁴. The NeuCube used leaky integrate-and-fire (LIF) neurons and modified the STDP rule slightly by augmenting the weight updates with rank order on top of spike delay time. Through this, the first spike to arrive at a postsynaptic neuron will not be rewarded with a strengthened synapse when

the neuron fires, but the final spike which triggers the action potential will receive a heavily weighted increase. This adjustment could assist with event-based simulators that do not employ discrete time as they may only have finite delays to update with.

To evaluate biological networks even further, one group used Izhikevich neurons^{35,36} with connections of random lengths to model murine hippocampal activity³⁷. This study implemented Hebbian STDP with a buffer for the impact of the delays. For example, if it would take 2 ms for a spike to reach a postsynaptic neuron, then a presynaptic neuron that fired 1 ms before the postsynaptic neuron fired should not be rewarded. That presynaptic neuron's spike would not have arrived by the time the postsynaptic neuron fired, so the modified rule is crucial to nullifying its impact and reward. Overall, this model highlighted the ability of STDP to synchronize chaotic activity in a hippocampal network, which may be a desirable feature in future spiking neural network implementations.

In a different direction, one group employed spike response neurons with more biological parameters to model the hippocampus in a different way³⁸. Instead of using a simple Hebbian rule, this study added in parameters for the NMDA receptor kinetics, which would directly influence the capacity of STDP. This implementation allowed the network to demonstrate its memory storage via coincident firings resulting from a variety of input patterns.

1.3.2 Classification Tasks

To transition to a machine learning application, one group tested multiple iterations of a STDP rule on LIF neurons to classify the MNIST handwritten digit dataset^{39,40}. The first rule they tested employed a frequency-based synaptic trace based on incoming signals. For synapses with a high frequency of activity, this trace term would result in the weight increasing quickly. A

second rule involved a synaptic trace for outgoing signals as well, which results in the network focusing more on the spiking rates rather than the spike timings. The last rule this study applied was a triplet STDP rule, but the details of this rule will be discussed in Section 1.6. After experimenting with the different configurations, this study ultimately reached a classification accuracy of MNIST of 95.0% after training 6,400 neurons on 40,000 images.

In a similar vein, a different group used spike response neurons and the adjusted edge delay STDP rule to embed images from the Semeion dataset⁴¹. Like the data modeling neural network above, this study incorporated edge lengths and as such, STDP was only applied when it made sense. For spikes that were not causal, even if the presynaptic spike occurred first, there would be no potentiation. As a result of the learning due to STDP, the authors were able to show an embedding of the Semeion handwritten digit dataset into the activity of the network. While explicit classification accuracy wasn't provided, an emphasis was placed on the efficiency of the embedding, especially relative to network and input size.

Extending NeuCube from the previous section for a machine learning task, one study involved treating the neural network similarly to the human visual system⁴². With 732 LIF neurons, the NeuCube uses both unsupervised STDP and supervised rank-order learning to learn dynamic vision input. Rather than using these methods to match empirical data, the authors extended NeuCube to learn and attempt machine learning tasks as well. Since NeuCube is designed for dynamic visual data, it was used to classify MNIST-DVS, which are videos of the MNIST handwritten digits. With 9,000 scale-16 training videos, NeuCube classified 86.09% of videos correctly. This demonstrates both that biological rules can improve the performance of machine learning models on harder tasks, but also that these rules may be more useful for tasks

that are inherently more biological. Vision in motion, like MNIST-DVS, may suit STDP better than static images, which less represent perception in the brain.

To connect Hebbian STDP to modern deep learning techniques, one group devised a network of non-LIF neurons to resolve multiple machine learning tasks^{43,44}. The conventional Hebbian STDP rule was used to train the weights between two convolutional layers and a global pooling layer was used afterwards for classification. Various tasks were performed with the network, including the Caltech face/motorbike (99.1%), ETH-80 natural object (82.8%), and MNIST (98.4%) classification tasks. The study briefly discusses how STDP serves the role of an unsupervised learning method more effectively than others that were tested. This network serves as a staunch example of how valuable biological paradigms can be for machine learning models.

In another deep learning approach, a group employed LIF neurons with two convolutional layers and two pooling layers⁴⁵. A standard Hebbian STDP rule was simplified for the network to deploy in a feedforward manner. This simplification involved binary weight updates instead of the updates being a function of the spike delays. To add to this simplification, a normalization scheme was added to prevent any edges from increasing indefinitely and creating unfair competition with other edges. With all of these pieces together, the network was able to classify 98.49% of MNIST images, 75.2% of ETH-80 objects, 71.2% of CIFAR-10 natural images, and 60.1% of STL-10 natural images. While this method suppresses many biological features for the machine learning approach, it results in a network that is quite successful at the intended tasks.

Considering a normalization scheme, another group implemented a normalized STDP rule with LIF neurons and inhibitory interactions to create a simpler network that could classify datasets efficiently⁴⁶. This model is unique in that it includes a refractory period for the neurons

to increase their resemblance to biological neurons. On top of that, the STDP rule was implemented such that any presynaptic spike causes LTD and subsequent postsynaptic spikes cause LTP. If the spikes are close enough in time, the LTP will override the LTD, but if not, the synapse will end up weaker than it was prior. This results in a clever way of ensuring that only the most synchronized spike pairs are able to strengthen their synapses. For normalization, the entire network is assigned a target average synaptic weight after each weight update. While the average weight is manually chosen, so it is not very biological, the method does provide more control over the network activity. With 9,000 neurons trained on 60,000 MNIST images, this network is able to achieve a classification accuracy of 95.07%.

To create a network resistant to forgetting through spontaneous activity, one group trained LIF neurons with Hebbian STDP in a semi-supervised fashion to classify MNIST⁴⁷. With only two layers in the network, STDP would change synapses while a teaching signal was active and firing rates would probabilistically control LTP and LTD based on a threshold. While the learning rule was based on spike timings, the trigger for learning was based on spike rates, which made this versatile in design. With 20,000 MNIST training images and 15 potential output classes, this network was able to achieve a classification accuracy of 96.8%. Overall, this is one of many models that demonstrate how spike-timing methods and rate-based methods do not need to be distinct and in fact could complement each other well⁴⁸.

To demonstrate the importance of the parameters chosen for these spiking neural networks, one group used a pair of conductance-based neurons to demonstrate the variability⁴⁹. While not solving a particular machine learning task, this study employed a typical STDP model to highlight how the mean synaptic strength was largely influenced by the initial standard deviation of the weights. This was more influential than the initial mean of the weights,

demonstrating how much thought needs to go into the weight parameters for STDP models. Without the proper balancing, the activity and weight distribution can quickly become undesirable.

As STDP rules are assessed as machine learning tools, one group concluded that they were fit to create Bayesian computation in certain networks if designed appropriately⁵⁰. Mathematically, it has been shown that Hebbian STDP can be extended to learn Bayesian decisions⁵¹. Building off of that, a spiking neural network was devised that could implement the biological rule in a way that could harness expectation maximization. With the conventional Hebbian STDP rewritten in a new form, the authors created a stochastic model to tackle machine learning problems with. With each image of MNIST converted to 708 input spike trains, the study was able to achieve a classification accuracy of 80.14%. This is a machine learning finding that may lend itself to the neuroscience as biological neurons may have probabilistic characteristics that are worth investigating.

These select computational models are not the only ones to incorporate Hebbian STDP and there are many other reviews that have focused in on examples that specialize in other validation tasks⁵²⁻⁵⁵. As it is the original proposed rule for STDP, there are myriad models that have incorporated it to fit data or to solve specific machine learning challenges.

1.3.3 Neuromorphic and Robotics Implementations

In more recent developments, the Hebbian STDP rule has been adapted to hardware implementations, such as neuromorphic chips and robotics. In a similar vein to the previously described models, one group adapted the rule to a neuromorphic chip to use voltage to attempt the MNIST classification task⁵⁶. In the proposed voltage-dependent synaptic plasticity rule,

Hebbian STDP is faithfully preserved, but a new term is added to represent the voltage of the presynaptic neuron for the weight updates. With 500 output neurons, this model was able to reach a classification accuracy of 90.56% on MNIST. As neuromorphic implementations are iterated upon, it may not be long before this accuracy is exceeded and more difficult datasets are attempted using neuromorphic hardware.

For an implementation of Hebbian STDP in robotics, a group took advantage of the rule through classical conditioning of robots to navigate through obstacles and find food⁵⁷. Using non-LIF neurons, a simulated insect was guided through a virtual world and taught to avoid collisions. After sufficient tuning, the same neural network was applied to a real Lego Mindstorms EV3 robot with similar success for locomotion. Beyond STDP and the neuron model, the network architecture itself was complex with olfaction and nociception included. To extend the idea of robotic locomotion, a separate group put together a robot with a brain, sensory system, nervous system, and body and taught it how to navigate in a maze using STDP as well⁵⁸. A third group used a Khepera IV robot with wheels, infrared sensors, LEDs, and a camera controlled by STDP to navigate using binary decisions⁵⁹. In a more straightforward model, a group designed a simulated two-wheeled robot to use phototaxis for locomotion⁶⁰. As STDP results in positive reinforcement, successful phototaxis cascades to the intended autonomy relatively quickly.

Aside from locomotion, Hebbian STDP has been used in robotics for other tasks as well. For example, one group designed a robot to identify the color of a Lego and sort it based on whether it is blue or red⁶¹. In a completely different context, another group navigated a simulated drone using LIF neurons with Hebbian STDP⁶². Beyond just navigating through the virtual world, the drone was tasked with calculating the appropriate trajectory to counter wind resistance

and to determine the shortest flight path to the destination. With many neurons serving as place cells, STDP was sufficient to navigate the drone effectively.

Towards anthropomorphic endeavors, one group designed a simulated robot that could identify and travel towards a target using the STDP in its neural networks⁶³. To extend the idea of multiple muscles working in concert, that group continued with a simulated arm that had to reach out to a target by alternating relaxing and extending flexors and extensors. Also, another group designed a robot that could simulate nociception and respond accordingly to avoid the insult in the future⁶⁴. To accomplish this, the study involved a robot with arms that had proprioception to determine when joints were at inappropriate angles. As the experiment advanced, the robot began to identify and exhibit avoidance behavior towards stimuli that would result in nociception.

1.4 Experimental Evidence for Reward-Modulated STDP

While the classical Hebbian rule for STDP holds true for many experimental models, there are certain regions of the brain or observed patterns that aren't well-modeled by it alone. Beyond that, many experimental observations have identified numerous chemicals that augment the effect of STDP even under identical conditions^{65,66}. While this opens up numerous questions and experiments for neuroscientists, this also provides a unique opportunity for machine learning models to take advantage of the various avenues of modulation of STDP. Computational models can integrate various types of parameters while maintaining biological feasibility if they can map to at least one neuromodulator.

While there are many neuromodulators in the brain, dopamine is the biggest one and it has an extensive catalog of research around it⁶⁷. One group used whole-cell recordings of murine

hippocampal cells to identify that dopamine increased the window for LTP along with the resultant increase in synaptic strength⁶⁸. To ensure the results weren't spurious, dopamine receptor antagonists immediately returned STDP to normal, implicating only dopamine in the observed results. Similarly, another group found that dopamine turned a previously depressed synapse into a potentiated one even retroactively⁶⁹. Even spike pairs where the postsynaptic spike fired slightly before the presynaptic spike, which would normally trigger LTD, resulted in significant LTP. However, the time window within which synapses are vulnerable to dopamine modulation may not be longer than a couple of seconds⁷⁰. Using two-photon uncaging with optogenetic tools, a group imaged the dendrites of murine striatal medium spiny neurons. The results showed that dopamine provided over two seconds after STDP glutamate release had minimal impact on dendritic spine growth. However, on a longer scale, it seems that dopamine has a large impact as well⁷¹. One experiment demonstrated that presenting a dopaminergic antagonist prior to encoding in a murine brain prevents recall of the encoded information after 24 hours, but not after 30 minutes⁷². It seems for the short- and long-term, dopamine makes a key impact on plasticity, but not in the medium-term from seconds to minutes. For both neuroscience and machine learning, dopamine's suppression of LTD and promotion of LTP could prove quite advantageous for certain purposes.

In addition to dopamine, acetylcholine is commonly implicated in affecting the dynamics of STDP within the brain⁷³. Unlike dopamine, acetylcholine doesn't universally increase LTP and reduce LTD. Acetylcholine seems to have a particular potentiating impact on behaviors that require immediate attention and a depressing impact on behaviors in the background. In one study, murine basal forebrains were stimulated using optogenetics, and imaging demonstrated that acetylcholine neurons actually responded faster to stimuli than dopamine neurons did,

perhaps to direct the rats' attention⁷⁴. As the visual cortex is heavily involved in attention depending on the visual cues that arise, it makes sense that murine visual cortex brain slices demonstrated a reliance on cholinergic reinforcement signals⁷⁵. Similarly, a murine acetylcholine antagonist experiment in the perirhinal region of the brain demonstrated how important cholinergic neurotransmission is for discerning between novel and familiar stimuli⁷⁶. Altogether, it seems like acetylcholine is heavily involved in discerning whether a stimulus is unexpected and, if it is, drawing attention to it quickly. In addition to affecting neurons directly, acetylcholine may impact astrocytes, which could increase gliotransmitter release⁷³. This could result in different downstream effects and may explain the variability in acetylcholine's impact on STDP. To add another wrinkle, whole-cell recordings of murine pyramidal cells demonstrated that acetylcholine biased towards LTD until dopamine was applied as well⁷⁷. As dopamine has been shown to expand the window for LTP on its own, it makes sense that this may overpower acetylcholine's slant towards LTD to promote potentiation in the end.

Within the midbrain, serotonin serves as a strong neuromodulator for spatial learning and sensory discrimination⁷⁸. Specifically, whole-cell recordings of murine dorsal raphe nuclei were observed to identify the role that serotonin played on learning. Interestingly, as the dorsal raphe nuclei send reward signals out to other brain areas, reward sensations at large rely heavily on serotonin signaling within the midbrain. In serotonin-depleted mice, glutamate released by dorsal raphe nuclei neurons has limited downstream effect and the neurons themselves are less active. With the focus of classical conditioning, one study compared the response in monkeys of the dorsal raphe nucleus to the response of the substantia nigra pars compacta, source of dopamine, upon presentation of a reward⁷⁹. The serotonin neurons scaled with the size of the reward, but the dopamine neurons only responded strongly if the size of the reward was unexpected. As such,

each neuromodulator seems to play a key role with serotonin neurons assessing the gross reward while dopamine neurons assess if predictions were accurate.

To step back and generalize, with all of the different neuromodulators acting in concert with one another, it is often easier to observe the resultant dynamics than to try to pick apart each individual actor. As such, one study took whole-cell recordings of murine cortical mice to observe the “eligibility traces” that designate spike pairs for LTP or LTD⁸⁰. In these brain slices, beta 2 adrenergic receptors triggered LTP while serotonin receptors triggered LTD. Other neuromodulators were observed, but these were the two that made the most impact on eligibility trace conversion. Specifically, the synaptic tagging and capture hypothesis fits nicely into the theory of these eligibility traces steering the direction of plasticity. Rather than synaptic tagging determining the fate of a synapse a priori, modern discussions of the hypothesis propose synaptic tagging as a floating state that can be influenced by neuromodulation, which is the same idea laid out by eligibility traces^{81,82}. With this theory added to STDP, plasticity becomes more nuanced than simply firing order, as the spike timings can transition the synapse into a vulnerable state whereupon neuromodulators in the area may determine its fate.

1.5 Computational Models Incorporating Reward-Modulated STDP

1.5.1 Data Modeling

When modeling reward-modulated STDP (including synaptic tagging, eligibility traces, etc.), the computational cost can be challenging due to the number of neuromodulators that may be at play in the particular brain region. However, one group attempted to simplify the model by incorporating Hebbian STDP with a global reward and noise to model data from the motor cortex of monkeys performing a virtual reality cursor trajectory experiment⁸³. In their model, the

authors combined STDP with a global neuromodulation term, which created an equation that resembles gradient descent. One major conclusion from the model was that the introduced noise was actually necessary to model the experiment data accurately, as the model wasn't able to correct error optimally without it.

1.5.2 Classification Tasks

To apply reward-modulated STDP to machine learning problems, the network needs to be very intentionally designed to simulate the presence of dopamine or other neuromodulators. In one model, a group used LIF neurons in a three-layer feedforward neural network to solve machine learning challenges⁸⁴. Rather than using a generic Hebbian STDP rule, this model incorporates an updated almost symmetric dopamine-STDP rule. With this rule, there is no LTD, but the STDP curve is adjusted so it is not perfectly symmetric. The method begins with the network going through a training process within which the label information of a stimulus is used as a teacher signal to become embedded via one-shot learning. When tested to classify MNIST images, with a hidden layer of 10,000 neurons and an output layer of 10 neurons, this model achieved a classification accuracy of 96.73%. For the Fashion-MNIST dataset, with 6,400 hidden neurons, the network achieved a classification accuracy of 85.31%.

In an effort to tackle an exceptionally challenging problem, one group proposed a Brain-inspired Sequence Production Spiking Neural Network (SP-SNN) to use reward-modulated STDP to learn grammar in languages via a chunking mechanism⁸⁵. This model, which employs LIF neurons, expands on the simple Hebbian STDP rule with a synaptic eligibility trace function. This function manifests as a reinforcement learning paradigm that is weighted by the number of recent subsequent firings of the spike pair. Through this, synapses are updated via STDP as well

as a reinforcement learning rule based off of a longer history of the activity along the synapse. With this model, the authors assessed the sequence memory and production of the model using a chunking method similar to how people chunk information to remember and learn more efficiently.

Since loss functions are incompatible with STDP, a group used LIF neurons to create a STDP-inspired temporal local learning rule (S-TLLR)⁸⁶. This rule incorporates eligibility traces and the principles of backpropagation to allow for precise credit assignment without the computational burden backpropagation alone would incur. S-TLLR proposes that a learning signal, perhaps from glia, represents the error returned down from the top layer. In addition, with the eligibility trace arising from the STDP equation, the authors composed a weight update equation that resembles backpropagation with spatial locality. To validate this combination, the study included the model tested on CIFAR-10 DVS, which resulted in a classification accuracy of 73.93%.

In a similar vein of trying to remove any teacher signal, another group took advantage of the dopamine state of the network as a representation of its change in performance⁸⁷. In essence, if the current performance of the network (based on mean squared error) is worse than its recent performance, the dopamine signal will be activated and synaptic weight changes will be more drastic. The reward-modulated STDP employed here is the same as in the monkey task model⁸³ with the only change being the regulation of “dopamine”. With 1,000 neurons in a single recurrent layer, this network model was used to demonstrate periodic pattern generation, memory-dependent computations, and network state-dependent routing of information. While there were no machine learning tasks directly addressed, this network serves as a promising

proof of concept for using reward-modulated STDP to embed and transfer information within a neural network.

To evaluate a simpler model, a group used non-LIF neurons and reward-modulated STDP to demonstrate reinforcement learning as a useful method to solve an XOR problem⁸⁸. This model proposes including a multiplicative reward function with Hebbian STDP to convert it to a reward-modulated form. With a three-layer feedforward neural network composed of 60 input neurons, 60 hidden neurons, and 1 output neuron, this model was tested on a benchmark XOR problem. Binary inputs were presented and the firing rate of the output neuron was used to determine which class was selected. The network was capable of discerning one class from the other in 99.1% of the simulated experiments that were performed. Beyond the optimal network configuration, the author goes into great detail in the mathematical basis for how reward-modulated STDP maps onto reinforcement learning and how that is a major boon for solving these machine learning challenges.

1.5.3 Neuromorphic and Robotics Implementations

Just as many of the machine learning models extended their utility with reward-modulated STDP over time, robotics models followed a similar trajectory. As an extension of the complex robot with multiple communicating systems⁵⁸, a new reward-driven STDP was implemented into the robot's brain using the synapto-dendritic kernel adapting neuron (SKAN) model. To challenge this more intelligent robot, the authors placed it in a maze with changing features to assess if it could learn information and then overwrite that when presented with a new context. The robot was successful, demonstrating the improvements gained by the new STDP rule. Similarly, another group trained a simulated robot with Izhikevich neurons to find food and

avoid poison in a collision-free environment⁸⁹. To avoid learning windows not being optimal for dopamine-modulated plasticity, this model triggered dopamine release upon reward expectation, similar to what was observed in the murine midbrain⁷⁹. In an effort to solve the challenge of collision avoidance, one group designed a simulated robot model with LIF neurons to navigate past obstacles towards a target⁹⁰. Using a three-layer feedforward network with 75 input neurons, 210 hidden neurons, and 2 output neurons, the authors compared their network model to a more traditional Q-learning approach. With reward-modulated STDP guiding the robot, it was able to perform better than when trained with Q-learning with the same number of training iterations.

For a more time-sensitive task, a different group designed a wheeled robot with reward-modulated STDP to traverse along a course without leaving its lane⁹¹. With a neuromorphic vision sensor, a Pioneer robot was able to quickly acquire lane data and process it using the STDP rule to follow the road patterns. Internally, there are 32 input neurons that connect to 2 LIF output neurons in an all to all fashion with updating synapses. As the robot proceeds through the training process, it will more quickly correlate specific signals with needing to turn left or right to improve the lane adjustment process even for unfamiliar paths. Because the STDP method is based on a local rule with minimal communication, training a robot to make lane changes promptly is easier than when compared to classical reinforcement learning algorithms that need to process a state evaluation step.

Extending navigation tasks beyond land, reward-modulated STDP was used to navigate a simulated sailboat to a target with realistic boating mobility, such as tacking and gybing strategies⁹². Using LIF neurons in two-layer fully connected feedforward network, around 42% of the 1024 experiments reached all target points within the allotted time. Ultimately, the model

did not travel faster than ANN-based methods, but it resulted in less deviation across trials, which could be a boon in certain contexts.

To complicate locomotion further, one group designed a flapping robot to fly between specific targets⁹³. This robot incorporated LIF neurons and a reward-modulated STDP rule to hover, follow trajectories, and land safely without collisions or damage. With a three-layer feedforward neural network and various onboard sensors, the RoboBee is able to traverse along specific paths safely.

1.6 Experimental Evidence for a Triplet STDP Rule

So far, we have only considered STDP through a strict interpretation of Hebb's theory involving spike pairs. However, in networks with a lot of activity, such as with spike trains or high-frequency oscillations, it is possible that plasticity may be influenced in a more nuanced way. With access to information of multiple spike-timings at once, plasticity rules can leverage even more local information for embedding.

In murine visual cortical brain slices, one group was able to experimentally observe variations in synaptic modifications depending on the timings of preceding spikes²⁹. Depending on whether the third spike was presynaptic or postsynaptic, new outcomes were observed. With a preceding postsynaptic spike, it was observed that the synapse weakened more than expected by the two spike pairs themselves. With a succeeding presynaptic spike added, the synapse strengthened beyond what the two spike pairs would predict. For both scenarios, a maximum spike window of 30 ms was observed before the triplet STDP phenomena disappeared and spike pair STDP took over.

Similarly, in murine hippocampal neurons, another group observed potentiation via calmodulin-dependent protein kinase II (CaMKII) and depression via calcineurin resultant from spike triplets⁹⁴. Matching the prior data, succeeding presynaptic spikes resulted in stronger potentiation of the synapse. However, in this experiment, a preceding postsynaptic spike either resulted in potentiation or no net change to the synapse. In addition, the maximum spike window for triplet STDP here was 70 ms, so it is possible that the hippocampus responds differently to spike trains than the visual cortex does. Finally, the influence of calcium-mediated enzymes may point towards calcium's control over the plasticity of various synapses. This idea will be revisited in Section 1.12.

1.7 Computational Models incorporating a Triplet STDP Rule

1.7.1 Data Modeling

Because there is no consensus on the most accurate triplet STDP rule, groups looking to incorporate it into computational models must discern for themselves which makes most sense⁹⁵. One such group implemented a triplet STDP rule to model visual cortex and hippocampus data⁹⁶. For their model, they decided to implement the triplet rule with both preceding postsynaptic spikes and succeeding presynaptic spikes resulting in potentiation⁹⁴. This model was capable of reproducing all but one of the experiments from the source papers with the equations that were proposed. In addition, it was even capable of modeling data for a proposed quadruplet rule, demonstrating that triplets may contain the maximum complexity required. Lastly, the authors noted that if the spike trains are thought of as Poisson processes, this triplet rule resembles the Bienenstock-Cooper-Munro (BCM) learning rule⁹⁷.

In an effort to demonstrate the impact of dendritic gating, one group put together a network model of pyramidal cells incorporating triplet STDP⁹⁸. While the focus of the paper was on the compartmentalization of the dendrite and soma in a neuron, the triplet STDP rule was chosen to effectively replicate pyramidal cell activity. Building off of the previous model, this one also used the triplet STDP rule with preceding postsynaptic spikes causing potentiation⁹⁴.

1.7.2 Classification Tasks

In order to assess the value that the triplet rule could provide, one group compared it against two different STDP schemes to assess a network's ability to learn a spatial pattern⁹⁹. Using spike response model neurons, this model compared pairwise STDP, triplet STDP, and a global STDP in which every spike informs every neuron. Based on the heuristic search algorithm that was used to test the networks, the triplet STDP rule performed the best, but it had issues with falling into local extrema that required extra training in some cases.

1.7.3 Neuromorphic and Robotics Implementations

Extending the BCM idea further, one group implemented a triplet STDP rule in a memristor to intentionally resemble BCM¹⁰⁰. This model built off of the potentiating preceding postsynaptic spike triplet STDP rule⁹⁴ and recapitulated the enhanced depression effect that prior memristor models were unable to by that point. As such, the memristor model took a major step towards memristors resembling biological data even closer, particularly via the BCM approach.

1.8 Experimental Evidence for Anti-Hebbian STDP

Up to this point, all of our STDP rules have built off of the classical Hebbian curve for plasticity. However, there is strong experimental evidence and computational value in the flipped version of the curve. The negative hyperbolic cotangent curve is referred to as the anti-Hebbian STDP rule¹⁰¹ and it can be difficult to reason with, as it implies neurons that fire together suppress the synapse between them. Rather than a learning rule intended to support the proliferation of signaling within a network, anti-Hebbian STDP focuses on suppressing activity in regions where activity is undesirable.

In a hybrid experimental-computational study on murine pyramidal neurons, one group discovered that the switch between Hebbian and anti-Hebbian plasticity is actually controlled by the location of the dendrite of the postsynaptic neuron¹⁰². Distal dendrites result in the recruitment of anti-Hebbian plasticity machinery while proximal dendrites lead to the initiation of Hebbian plasticity. Beyond the importance of network geometry on signaling latencies and connectivity, dendritic heterogeneity adds a new consideration to where synapses form between neurons. It opens the possibility that this heterogeneity extends beyond a binary and could highlight a spectrum of dendritic domains that could respond uniquely to incoming stimuli¹⁰³.

To investigate the mechanism of the heterogeneity, another group employed patch-clamp recordings and two-photon imaging on murine neocortical brain slices¹⁰⁴. With these tools, the authors discovered that ionic backpropagation from the action potential contributes to the function of anti-Hebbian STDP. However, they also observed that the ionic backpropagation was essential for Hebbian STDP at proximal dendrites as well, so it seems to be critical for plasticity at large. The role that calcium and voltage-gated sodium channels play in action potential

backpropagation may be significant and may determine whether neurons in other regions experience anti-Hebbian STDP or not.

To interrogate the role of calcium in anti-Hebbian STDP, a group investigated the response of murine somatosensory cortical brain slices to high- and low-frequency action potential bursts¹⁰⁵. During these action potential bursts, it is possible to control the induction of LTP and LTD in both proximal and distal dendrites, but the authors show that the spiking timing requirements are completely different depending on location. They observed that due to the geometry of action potential propagation, distal dendrites are impacted more heavily by calcium electrogenesis resulting from downstream depolarization. To convert anti-Hebbian STDP to Hebbian STDP in the distal dendrites, the authors propose the addition of T- or R-type voltage-gated calcium channels to create a voltage that unblocks NMDA receptors.

Moving beyond animal models, anti-Hebbian STDP has also been observed in the human brain in long-range synapses¹⁰⁶. Transcranial magnetic stimulation was applied to a region between the posterior parietal cortex and the primary motor cortex in human subjects while electromyographic traces were recorded from the right first dorsal interosseous. As the brain was stimulated with 100 spike pairs at 0.2 Hz, the motor-evoked potentials were measured and their amplitudes were compared. In response to the stimulation, the MEPs reduced in amplitude when a presynaptic spike preceded a postsynaptic spike and increased in amplitude when a postsynaptic spike preceded a presynaptic spike. As such, this long-range synapse seems to follow anti-Hebbian STDP. This aligns with the other evidence for dendritic heterogeneity, as long connections likely involve dendrites that are far removed from the soma, which would explain why anti-Hebbian STDP would be active.

1.9 Computational Models Incorporating Anti-Hebbian STDP

1.9.1 Data Modeling

To observe anti-Hebbian STDP's impact on computational models, one group create a LIF neuron-based model of a two-tiered neocortical area¹⁰⁷. With a two-layer network with 100 neurons in each layer, the authors compared the performance of Hebbian STDP to anti-Hebbian STDP in terms of stable weights. Anti-Hebbian STDP demonstrated superior edge weight stability, as the plasticity rule preserves a fixed point for the edge weights to orbit around. Since anti-Hebbian STDP can't create a positive feedback loop, it may be a useful tool for computational models that are sensitive to hyperactivity.

1.9.2 Classification Tasks

Combining an interest in modeling with machine learning, one group designed a network with Hebbian and anti-Hebbian STDP to mimic an autoencoder¹⁰⁸. Rather than implementing both STDP rules, the author combines them into mirrored STDP, which is a function of hidden layer firing time minus visible layer firing time. As such, connections toward the hidden layer treat the hidden layer as the postsynaptic neuron and connections away from it treat it as the presynaptic neuron. With a two-layer network full of LIF neurons, reconstructions of MNIST images were able to achieve correlations of up to 91%. These reconstructions were deciphered based on the spiking activity of the network when it was stimulated by a particular image with plasticity turned off. The embedding within the hidden layer contains receptive fields similar to the ones found in the visual cortex. This serves as an effective proof of concept for other models that may be looking to incorporate a receptive field analogue into their approach.

Rather than isolate anti-Hebbian STDP on its own, computational models can leverage employing it along with the previous STDP rules that have been discussed. With that thought, one group devised a reward-modulated STDP rule that triggered Hebbian STDP with a reward signal and anti-Hebbian STDP with a punishment signal¹⁰⁹. With their four-layer feedforward network, they were able to achieve classification accuracies of 98.9% on the Caltech face motorbike dataset, 89.5% on the ETH-80 natural objects dataset, and 88.4% on the NORB 3D toy dataset.

In a similar separation model, a group aimed to recreate backpropagation through STDP in non-LIF neurons by adapting both Hebbian and anti-Hebbian rules with a teacher signal¹¹⁰. To explain, while the teacher signal is active, target neurons are rewarded with Hebbian potentiation while non-target neurons are punished with anti-Hebbian depression. As such, once the teacher signal is off, the rewarded neurons will be more likely to respond to similar stimuli. Since there is a teacher signal, this approach is supervised, but it provides an example of biological rules proving useful in a supervised model. Even though STDP is usually touted as a solution for an unsupervised local rule, there are other contexts where it may be useful. With one hidden layer with 500 neurons and a second with 150 neurons, this network was able to classify MNIST with 97.2% accuracy.

Furthering the connection between spike-timing and rate-based methods, one group implemented anti-Hebbian STDP in a way that sorted neurons into assemblies¹¹¹. Upon presentation of the input, neurons in the hidden layer with the highest firing rate were designated a degree of belonging in a particular assembly. When STDP is triggered to change the weights of a particular neuron, that neuron's weights may be decreased via anti-Hebbian STDP if it is active with the wrong assembly. This creates quick separation between neurons in assemblies

representing particular stimuli. As before, the spike rate informs whether a synapse should be augmented, but spike-timing rules determine the magnitude of change. In certain models, such as Assembly STDP, this synergy results in efficient separation of neurons that respond to one stimulus from another. This also highlights a key case where both Hebbian and anti-Hebbian STDP can be employed to maximize divergence between subpopulations.

1.9.3 Neuromorphic and Robotics Implementations

To create a neuromorphic implementation of the contrastive divergence algorithm, one group incorporated Hebbian and anti-Hebbian STDP rules with LIF neurons¹¹². Instead of incorporating separate phases for contrastive divergence, the proposed rule updates the weights after every spiking event. The rule is able to mimic contrastive divergence since there is a global modulatory rule that switches between Hebbian and anti-Hebbian plasticity rules at each synapse. With a restricted Boltzmann machine with 500 hidden neurons and 40 output classes, the authors were able to reach 91.9% classification accuracy of MNIST images.

1.10 Experimental Evidence for Inhibitory STDP

Within the brain, there exist more than just synapses that contribute to the membrane potential of the postsynaptic neuron. Some connections between neurons are inhibitory, which means a signal will actually hinder the future activation of the downstream neuron. However, just like synapses between excitatory neurons, these connections are plastic and they also may adapt to different stimuli. Learning rules of these synapses are harder to decode, as if an inhibitory neuron is successful, the postsynaptic neuron won't fire and there won't be any postsynaptic potential to measure.

To directly tackle the challenge of inhibitory plasticity behavior, one group used two-photon microscopy and whole-cell recordings on murine hippocampal slices¹¹³. Immediately, the biggest observation was how dynamic inhibitory synaptic boutons were. Throughout the imaging data, the boutons were constantly probing for potential new synapses to try to expand their network. Interestingly, neuronal activity made a major impact on the behavior of synaptic boutons. The number of transient, probing boutons decreased when there was a reduction in spiking activity. Inversely, when there was hyperactivity, the number of persistent, stable boutons declined. Through both cases, it is clear that inhibition is quite plastic as a result of network activity.

In an attempt to model a STDP rule, a group studied murine midbrain slices and attempted to induce observable plasticity changes¹¹⁴. GABAergic synapses, which are the inhibitory synapses in the brain, demonstrated conventional Hebbian LTP and LTD towards ventral tegmental area dopamine neurons. In somewhat counterintuitive fashion, when the GABAergic synapse would fire to reduce the membrane potential of the postsynaptic neuron and then that neuron would fire anyway, the synapse would strengthen. This may be to prevent firing in the future, as the strengthened synapse would release more GABA, which would cause a larger reduction of the membrane potential of the dopamine neuron.

Zooming in further, one group investigated the plasticity rules of two particular subtypes of inhibitory interneurons based on patch-clamp recordings of murine hippocampal brain slices¹¹⁵. Specifically, the two interneuron plasticity rules studied were the LTD in parvalbumin interneurons and the LTP in somatostatin interneurons. Despite the fact that both subtypes only take part in one particular plasticity function, they do follow Hebbian STDP for it. The fact that parvalbumin LTP and somatostatin LTD isn't naturally observed opens up an interesting

question about interneuron subtypes and inhibitory STDP. It is possible that some of the heterogeneity that is observed in inhibitory STDP plasticity data is a result of the specific subtypes in the experiment and their unique plasticity rules and limitations^{116,117}. Extending this idea, the interneuron subtypes may have been useful for special tuning of different functional networks in the brain. To support this idea, a study on murine hippocampal neurons demonstrated that certain spatial exploration tasks resulted in potentiation of parvalbumin-expressing inhibitory interneurons while depressing cholecystokinin-expressing inhibitory neurons¹¹⁸. Even though both subtypes were inhibitory, learning in the brain led one type to potentiate and the other to depress. It is possible that the interneurons' downstream targets were the main reason behind this, but there could also be other characteristics of each subtype that would lead to this prioritization.

Inverse to these past two rules, a different group observed anti-Hebbian STDP in recordings of the murine auditory cortex¹¹⁹. When recoding parvalbumin-interneurons and their downstream principal neurons, the authors observed that interneurons spiking first resulted in the inhibitory synapse strengthening while the opposite occurred if the interneurons spiked second. Even more interesting is the fact that this anti-Hebbian LTD switched to LTP when the mice aged. It is possible that during brain development, LTD serves an important role in the auditory cortex, but upon maturity and exposure to more sounds, LTP is more beneficial.

Continuing the effort, one group used whole-cell recordings of murine entorhinal cortical cells to put together a different inhibitory STDP rule to consider¹²⁰. Based on the data that was compiled, the STDP curve looked similar to a cube function placed between the conventional hyperbolic cotangent function. Functionally, this is similar to the previously described rule in that a successful spike pair would increase the inhibitory weight of the synapse even though the

presynaptic neuron was originally incapable of preventing the action potential. A secondary important observation from this study was the importance of calcium in inhibitory STDP, which is consistent with calcium's role in excitatory STDP as well.

Using whole-cell recordings of murine auditory cortex brain slices, inhibitory synapses were observed to potentiate based on spike timing, but independent of the spike order¹²¹. This symmetric STDP rule was observed to cause the largest increase in synaptic weight when the two spikes were almost simultaneous with smaller increases as the gap between them grew. Following this, depression of the synaptic weight only occurred when the window between the spike pair was sufficiently large in either direction. This rule, therefore, resembles a coincidence detector more than a mechanism to guide the activity in a particular direction. To conclude, the authors discussed the importance of the excitatory/inhibitory (E/I) balance and how synaptic plasticity regulates it. Building off that, the same group investigated how plasticity calibrates this balance in murine auditory cortical pyramidal neurons¹²². Using whole-cell recordings and two-photon microscopy, the authors observed how resilient synapses are to achieve a particular E/I balance. A separate group even observed that inhibitory synaptic spines will augment their size to preserve a certain volume as their total number changes or in response to the overall volume of excitatory synaptic spines changing¹²³. The myriad functions and biological machinery at play for that demonstrate that inhibitory plasticity may serve a role beyond just embedding information. Inhibitory neurons are critical to maintain a healthy amount of activity within the network and to allow for new connections to have a chance at development¹²⁴.

To contrast with the STDP rules stated above, one group found a form of inhibitory plasticity in murine medial superior olive neurons that was based on persistent activity rather than specific spike timings¹²⁵. Using patch-clamp methods, the authors were able to identify that

potentiation was insensitive to the relative timings of the inhibitory postsynaptic potentials. However, while the timings had minimal effect, the number of pairings in general did influence the changes that were observed. It is possible that this region had a unique STDP rule because the synapses were glycinergic instead of GABAergic or because it was specially tuned for auditory sensing. Ultimately, the heterogeneity in the pathways that lead to inhibitory plasticity result in a large variability in the data that is collected in different brain regions.

To look at networks as a system, one group investigated the impact that inhibitory STDP would have on excitatory STDP in the murine visual cortex¹²⁶. With whole-cell recordings, the authors observed how inhibitory LTP attempts to interfere with excitatory LTP via myriad signaling pathways. It makes sense that after increasing the ability to silence a postsynaptic neuron, a network may not be keen to increase the strength of outgoing signals towards it. However, with the prevalence of inhibitory interneurons in important brain regions, this complicates plasticity modeling dramatically. If inhibitory plasticity attempts to suppress excitatory plasticity and vice versa, then the most realistic computational models must account for these interactions¹²⁷. Further, it solidifies the theory that these two rules share biological machinery or, at the very least, take advantage of the fact that their pathways can interact at numerous points.

1.11 Computational Models Incorporating Inhibitory STDP

1.11.1 Data Modeling

In a similar fashion to the triplet STDP rules, there are numerous inhibitory STDP rules to choose from for computational models and none of them serve as a frontrunner for realism. In addition, since inhibitory STDP is so heterogeneous depending on the brain region and

interneuron subtype, there is no one-size-fits-all solution for incorporating inhibitory STDP into a computational network.

Using LIF neurons and a custom differential inhibitory STDP rule, a hippocampal network was simulated to observe the stability of place fields as a mouse traversed a track¹²⁸. This implementation of STDP was quite unique, as it allowed for inhibition to adjust to the synaptic input, rather than updating a specific weight due to spike pair timings. This model proved to be quite successful at recapitulating the robustness found in the murine hippocampus. Place cells naturally emerged to represent specific regions of the track in the mouse's memory, but even when the initial place cells were silenced, the mouse was able to embed that information in new cells.

Rather than focusing on inhibitory STDP's impact on spiking activity, the same group designed a model to demonstrate the capability of inhibitory neurons to store stimulus information¹²⁹. With LIF neurons representing pyramidal cells, parvalbumin interneurons, somatostatin interneurons, and vasoactive intestinal polypeptide-expressing neurons, a network with stochastic connectivity was generated. With these varied cell types, the model demonstrated that in the absence of reward stimulus, disinhibition can lead excitatory plasticity to embed stimulus information. This scheme could be possible in multiple areas of the brain that contain a top-down reward architecture.

In a seminal paper regarding inhibitory STDP, one group employed a LIF neuron with the symmetric inhibitory STDP rule to recapitulate the self-organizing dynamics in the brain¹³⁰. As we discussed in the previous section, the excitatory/inhibitory (E/I) balance of a network is extremely important for avoiding hyperactivity while also allowing for enough activity to embed stimulus information. In this study, the authors observed that the response to stimuli is sparse,

which allows for enough activity to transfer information without persistent volleys of activity. This behavior has been observed in the murine auditory cortex to efficiently learn information and capturing it could prove extremely useful for computational models.

To combine many of the STDP rules discussed in this review, a group devised a network of perceptrons with Hebbian STDP, neuromodulated STDP, and inhibitory STDP¹³¹. The purpose of such a complex model was to highlight the value of the E/I balance of networks with inhibitory plasticity while also connecting these biological learning rules to the machine learning credit assignment problem. With these learning rules in place, the authors re-evaluated some of their previously published experimental data to discern reasonable E/I ratios and what they would mean for the learning rates of the networks.

1.11.2 Classification Tasks

As it was for data modeling, credit assignment is a method of great interest for designing networks to solve machine learning problems. Since biologically plausible networks are unable to use backpropagation to capture credit assignment, complex inhibitory microcircuits can be employed to reach a similar outcome. In one model, high-frequency bursting was incorporated with a burst-dependent plasticity rule and multiple types of inhibitory interneurons to effectively embed information within a network¹³². With bursting allowing for more significant weight changes due to STDP, the model was able to solve the XOR task and achieve classification accuracies of 79.9% on CIFAR-10 and 98.9% on MNIST. While all of the neuron types and activity rules result in a model with high complexity, the tunable features could result in robustness to different tasks or difficult data that could be useful for future computational models.

In another model, pyramidal cells are broken up into different compartments and connected to inhibitory interneurons in an effort to approximate backpropagation to succeed at recognition tasks¹³³. While this model also doesn't adapt any specific inhibitory STDP that was mentioned in the previous section, it does define unique rules for each type of neuron. To test an established benchmark, the network reached a classification accuracy of 98.04% of MNIST with two 500-neuron hidden layers and one output layer.

To test the boon against catastrophic forgetting that inhibitory STDP provides, a group put LIF neurons in a three-layer network with inhibitory interneurons to discern novel inputs from familiar ones¹³⁴. In this model, Hebbian STDP was used for the excitatory synapses and the symmetric STDP rule¹²¹ was used for the inhibitory synapses. Different 2D arrays of pixels were presented to the network periodically and the firing rate of the readout neurons were monitored to determine the network's classification of the stimulus. One promising result from this method was that the embedded information was stable for a long duration, which could have many practical applications. In addition, this network retained the ability to learn new inputs even after training was completed and it did so without catastrophic forgetting of prior learned inputs.

Shifting to a supervised approach, one group incorporated backpropagation with STDP rules in a network with unique edge delays to solve the XOR problem¹³⁵. Using spike response model neurons, the network employed a normalization function on the weights after every update to ensure that they didn't get too large. To test the model further, it was challenged with the Iris dataset and the number of iterations required to meet a convergence rate of 80% was recorded for comparison.

1.11.3 Neuromorphic and Robotics Implementations

To apply inhibitory STDP to a specific challenge, a different group devised a five-layer neural network of LIF neurons which incorporates Hebbian excitatory STDP, symmetric inhibitory STDP, and a reward-modulated STDP rule¹³⁶. This network was designed to monitor personal ECG devices in real-time to make quick, effective decisions. Classification performance of the different states of heartbeats was competitive to other algorithms, but there was a major boon to the energy consumption required to complete the computations. For ECG classification and other wearable-related algorithms, this combination of STDP rules in different network layers could prove to be quite advantageous.

1.12 Experimental Evidence for Gliotransmitter Plasticity

When discussing the dynamics of synapses in the brain, it is inaccurate to omit the impact that astrocytes have, either by cleaning excess neurotransmitter or secreting modulating gliotransmitters¹³⁷. The tripartite synapse has been defined to account for the role that glia has on network activity and plasticity, but experiments are still being done to get a better understanding of the full scope of that role¹³⁸.

To prove how important astrocytes are to synaptic plasticity, one group induced astrocyte dysfunction in a murine model and demonstrated the impact on the hippocampus function¹³⁹. By eliminating gap junctions between astrocytes, essentially isolating each one, the authors removed a dimension from the network in the hippocampus. Rather than there being a network of neurons and the syncytium of astrocytes that could communicate separately and with each other, only the network of neurons remained in concert. In the mutant mice, pyramidal cells demonstrated less excitability, weaker potentiation, and an increase in AMPA receptor-mediated

miniature excitatory postsynaptic currents. On a behavioral level, these mice developed subtle sensorimotor deficits and deficits in spatial learning after 90 days. Even though the astrocytes are still physically present, losing their ability to communicate with each other and transfer that information to neurons absolutely hinders the mice's learning and memory.

In various brain regions, the impact of astrocytes on synaptic plasticity has been well-documented¹⁴⁰. While the simplest method of astrocyte interference is gliotransmitter release, it is not the only process that astrocytes take part in. The catalog of gliotransmitters that can be released is impressive to begin with, but they can also coordinate between neurons in different regions. Since a single astrocyte is connected to thousands of individual synapses and astrocytes are connected to each other via gap junctions in the syncytium, each one can monitor a large array of activity. Through calcium transients in the syncytium, one astrocyte can signal to another to release glutamate at a faraway synapse to increase the network activity based on specific situations.

Because of the ubiquity and variety of astrocyte function, there are no clear, widely accepted plasticity rules that can be proposed for computational models. The utility of astrocyte function is undeniable in a computational context¹⁴¹, but like with inhibitory STDP rules, the biology needs to match the computational challenge for it to make sense.

1.13 Computational Models of Gliotransmitter Plasticity

1.13.1 Data Modeling

While there are few consensus theoretical rules for astrocyte function in the brain, the diversity in function has made astrocyte models quite appealing for a variety of purposes^{142,143}. To investigate the impact of calcium dynamics within astrocytes, one group devised a model of

100 Hodgkin-Huxley neurons connected to two astrocytes¹⁴⁴. The astrocytes are represented by their cytosolic calcium dynamics as they change over time in response to the network activity. When the calcium concentration crosses a threshold, gliotransmitters are released, which multiplies a scaling parameter to the subsequent postsynaptic currents. Lastly, this model even included a function for gap junctions so calcium ions can be distributed from one astrocyte to the other. This implementation consisted of many dynamic equations, but the complexity may prove useful for tackling intricate computational problems.

Connecting astrocytes to other plasticity rules, one group employed Tsodyks-Markram neurons to model the impact that gliotransmitters can have on potentiation and depression¹⁴⁵. Their model was able to clearly demonstrate how gliotransmitters could flip the direction of synaptic plasticity depending on how the astrocyte was trying to influence it. This is similar to how reward-modulated STDP worked both experimentally and computationally, but astrocytes can serve even more functions on top of that.

Since working memory is an area of interest in both neuroscience and machine learning, a group created a model of 1000 synapses and one astrocyte to observe the impact that gliotransmission may have on working memory¹⁴⁶. Using Bernoulli equations to control neuron and glia spiking, the authors were able to demonstrate how gliotransmission allowed for a slower decay of working memory. Not only does this match the experimental data of astrocyte dysfunction, but it also could prove very beneficial for current models that require longer-term memory. A separate group used Izhikevich neurons to test the working memory of a neuron-astrocyte network as well¹⁴⁷. One notable feature of these neuron-astrocyte networks is that the different cell types can operate on different time scales. Neurons can learn the stimulus urgently while astrocytes can embed the stimulus from them at a slower, sustained pace.

To see a more general idea of the impact that astrocytes have on network activity, one group built a network with tripartite synapses and compared the average spiking frequency with and without the astrocytes¹⁴⁸. Rather than just grossly increasing the average network activity, the introduction of astrocytes created stable states for the network to settle into and transition between. This phenomenon of state transitions could prove very useful, especially with issues like binary decisions or catastrophic forgetting.

1.13.2 Classification Tasks

To benefit from astrocytes in more tangible tasks, simple machine learning challenges were a good start. Extending the previously described model¹⁴⁸, the network added a layer for an astrocyte to integrate and influence neural activity^{149,150}. Using Hodgkin-Huxley neurons, this network was tasked to recreate the stimulus image, similar to an autoencoder. With the astrocyte included in the network, the mean squared error of each generated image appeared to be much lower. This demonstrates that astrocytes could both assist with the classification task and also with the direction of creating biological mimics of autoencoders.

Bridging the gap between modern machine learning and these astrocyte networks, one group designed a neuron-astrocyte liquid state machine to tackle common machine learning challenges¹⁵¹. Using LIF neurons and astrocytes, the authors created a standard three-layer liquid state machine, but the hidden layer has an additional astrocyte modulatory unit. This astrocyte both integrated activity from the input layer and controlled the STDP depression learning rate to regulate activity in the hidden layer. Overall, the neuron-astrocyte liquid state machine achieved classification accuracies of 97.61% on MNIST, 97.51% on n-MNIST, and 85.84% on Fashion-MNIST.

With a similar goal in mind, a different group showed that a neuron-astrocyte network can perform the same core computation as a transformer¹⁵². Beyond the excited new computational tool, this serves as an interesting study into how astrocytes and neurons may actually work together in the brain. In addition, the authors mention that their proof could serve as a demonstration of how transformers may function within the brain.

To create a particularly difficult challenge for astrocyte networks, one group created a network of Izhikevich neurons connected to an astrocyte layer and tested it with damage in multiple configurations¹⁵³. Through the working memory tasks that were performed, the model with the astrocyte layer was capable of retaining memory up until 60% synaptic impairment, which is 20% higher than the network without the astrocyte layer. Through this, the robustness of the network resulting from the astrocyte layer is elucidated, which is a very desirable feature.

1.13.3 Neuromorphic and Robotics Implementations

Even though the study of tripartite synapses is relatively new, neuromorphic approaches have already begun development. One group designed the spiking Neuronal-Astrocytic Network module for Intel's Loihi neuromorphic chip¹⁵⁴. The authors mention that astrocytes serve a major role in being able to synchronize the neurons at any point if necessary. In addition, networks with astrocytes can integrate information on a more global level than STDP alone could, which opens up opportunities for fine tuning of learning rules. To follow-up on this development, the authors implemented the network in a full robot to control its locomotion¹⁵⁵. With a hexapod robot controlled by their neuron-astrocyte network, it was able to change speeds and adjust its gait as necessary.

1.14 Experimental Evidence for Adaptive Myelination

While glia at large was briefly touched upon in the last two sections, astrocytes were the main focus. However, oligodendrocytes are glia that are also remarkably important to the dynamic function of the brain. Oligodendrocytes are responsible for myelinating axons, which can dynamically change the conduction velocity with which action potentials travel. Changes in conduction velocity could lead to myriad changes due to STDP and other timing-based integrations. Though adaptive myelination by oligodendrocytes is quite a bit slower than the other plasticity rules mentioned in prior sections, it still can have a heavy impact¹⁵⁶.

To demonstrate the ubiquity of adaptive myelination, one study performed magnetic resonance imaging (MRI) on both human and rat subjects while they were challenged with spatial learning tasks¹⁵⁷. Diffusion tensor imaging was done on rat subjects after they had trained for one day in a Morris water maze and myelination in the fornix and hippocampus was observed to be different. The same was done and observed with humans, except with a racing video game to hone their muscle memory.

Using optogenetic tools, another group was able to increase murine premotor cortex activity to trigger oligodendrogenesis and myelination¹⁵⁸. Electron microscopy was used to study the effect of the induced activity. To demonstrate the dependence of myelin plasticity on spiking behavior as well, the authors showed the increase in myelin thickness that resulted from consistent optogenetic stimulation. Since it has activity-dependent plasticity, adaptive myelination could now serve as a useful function for computational models to employ to solve specific challenges.

In another mouse model, the animals received low intensity repetitive transcranial stimulation non-invasively to induce adaptive myelination¹⁵⁹. Interestingly, the transcranial

stimulation actually had no effect on the average length of myelin sheets. In addition, increasing oligodendrogenesis without actually increasing the activity of the brain had no impact on nodal plasticity. It seems that transcranial stimulation may not be as effective at inducing myelination as optogenetics was. Also, it is possible that the location that was stimulated played a role in myelination not changing.

1.15 Computational Models of Adaptive Myelination

1.15.1 Data Modeling

To implement adaptive myelination in computational models, the major hurdle is the time scales. Since myelination occurs over the scale of hours when most other plasticity processes occur in the span of seconds to minutes, some concessions must be made to create computational models incorporating both.

To model the effect that conduction velocity plasticity would have in a network, one group created a network with Kuramoto oscillators as neurons and a unique differential equation for myelination¹⁶⁰. To simulate a large-scale network, the authors used the human connectome. With only myelin accumulation active, the authors noticed that the network spontaneously began to synchronize. This behavior was not observed if the myelination rule was turned off, demonstrating myelination's clear impact on the system. This network also naturally emulated the metabolic drag effect, which is a phenomenon in the brain such that longer axons tend to have slower conduction velocities. To test the resilience of the network, damage was inflicted upon it in the form of lost connections. Following the damage, the network is able to adapt and recover if the plasticity rule is still active, as different axons are able to change their speed to balance the activity out.

In another model that demonstrated the stability created by adaptive myelination, a different group put together a recurrent network of Poisson neurons to monitor the neural activity that resulted from myelination changes¹⁶¹. However, in this model, as the conduction velocities stabilized, they all started to increase together. It is possible that after adaptive myelination stabilizes the activity of the network, it may start to pull the velocities down to an optimal point. The authors noted an important point with adaptive myelination, which was that a lack of tight control on the myelination process would result in lower, uncorrelated activity. One possible boon to myelination being so slow in the brain is that there is immense control based on consistent activity rather than quick changes resulting from spurious activity.

1.15.2 Classification Tasks

Even though adaptive myelination is not widely adopted yet, there are a couple of computational models that have integrated it to solve machine learning problems. In one model, LIF neurons were assembled in a three-layer network with an adaptive myelination learning rule and a synaptic weight update rule¹⁶². With both learning rules working in concert to embed stimulus information, the network achieved a classification accuracy of 97.36% on MNIST. In contrast to prior models, this weight update rule included a homeostatic regulation term, which may demonstrate that adaptive myelination is not dynamic enough alone to approach stability when other forms of plasticity are active as well.

In a different model, there is a two-layer network composed of LIF neurons¹⁶³. The output neurons control the myelination of the network, as the network is based off of a winner-take-all architecture. When challenged with the USPS image dataset, which contains handwritten digits from 0 to 9, the network with 340 output neurons was able to achieve an F1-measure of

0.66 on the test set. That is a large number of output neurons to have when deciding between only 10 classes, but with only one learning rule active, the model could be built upon in the future.

1.16 Experimental Data for Synaptic Pruning as a Learning Rule

To round out the last of the major glia, this section will discuss synaptic pruning, which is the specialty of microglia^{137,164}. Synaptic pruning is a more powerful, but dangerous version of other synaptic plasticity rules. When a synapse isn't contributing to the overall activity or embedding of a network, pruning rules can reduce energy and computation costs by simply removing it. However, if the pruning rule is too loose, the network could quickly become too sparse to maintain activity. Unlike with potentiation and depression, a pruned synapse would require a significant cost to reinstate. However, if used correctly, synaptic pruning can increase the efficiency of networks dramatically by reducing the number of unnecessary computations.

During development, synaptic pruning is essential to properly stabilize the necessary networks as the organism is rapidly learning about the world¹⁶⁵. Similar to inhibitory STDP, synaptic pruning can occur in many different contexts and through different cell types. As such, it is difficult to model in a biologically accurate way since there is no ground rule for it. There are cases where it is executed correctly, such as in the case of brain development, but also instances where it leads to pathology, such as in autism spectrum disorder. While pathological manifestations of all of the rules discussed in this review exist, synaptic pruning is so destructive that this must be considered before it is experimented with or implemented in any model.

1.17 Computational Models of Synaptic Pruning as a Learning Rule

1.17.1 Classification Tasks

Because of the benefits that can be gained in efficiency, synaptic pruning is a learning rule that has drawn great interest in computational models. In one prime example, a network with both synaptic weight and synaptic connection learning rules to replicate experimental results based on spine dynamics¹⁶⁶. To add to their probabilistic pruning rule, the authors included a version of STDP, which contributed greatly to the model matching the spine dynamics from experimental data. As is true for all of the rules in this review, synaptic pruning will often work best when implemented in concert with other rules.

To further test the utility of synaptic pruning, another group implemented a winner-take-all network with LIF neurons to classify random Poisson spike time inputs¹⁶⁷. With binary synapses, the learning rule remains simple, as a synapse will either exist or it won't depending on what the activity has been. If this model's pruning mechanism were applied to non-binary synapses, there could be issues where very strong synapses were removed because they were inactive for an extended period of time. However, whether that is really an issue or not is dependent on the task and the network design itself. In a different model by the same group, synaptic pruning was defined to regulate synapse formation¹⁶⁸. Instead of focusing on pruning and working from there, the model added new synapses to improve classification for specific classes and then removed redundant synapses attached to the neurons with new ones. By determining a quota of how many synapses any neuron can have, the network is forced to constantly make pruning decisions to make room for the new synapses that are generated. This also avoids the problem of sparsity, as the network is only pruning synapses once it has

established a replacement for them. Using this method, the model achieved classification accuracies of 96.4% on MNIST and 88.1% for MNIST-DVS.

To connect realistic pruning rates to computational models, one group observed the pruning rate in murine somatosensory cortical slices and started from there when optimizing¹⁶⁹. In an effort to determine the optimal pruning mechanism, the model iterated four times with constant, increasing, decreasing, and ending pruning rates. While the first three are self-explanatory, the ending pruning rate refers to all of the pruning happening in the final time step, rather than gradually over the course of the simulation. Ultimately, the decreasing-rate method worked the best, which makes sense as the amount of pruning reduced as the network trained more and adjusted. However, this network did not incorporate edge weights or any STDP rules, so the choice of pruning rule may be different with more moving parts in a model.

1.18 Conclusion

Across this entire review, multiple plasticity rules were discussed in great detail, including the neurobiological bases, computational uses, and robotics applications. While many different rules were covered, there are still others that weren't discussed, such as rate-based methods or non-spiking biological neurons and artificial neural networks. While different neuromodulators and structural proteins were briefly mentioned, the impacts of enzyme kinetics and epigenetics were not touched upon even though they can also result in synaptic plasticity. Time scales and parameter choices were briefly talked about in relation to adaptive myelination, but to make a maximally biological network, they would need to be considered as well. To complement all of these rules, a deep dive into network geometry and signaling latencies could prove useful. With spike-timing so dependent on conduction velocities and refractory periods,

those considerations can't be ignored to maximize realism. On the machine learning side, very few challenges were discussed, as many STDP-based methods are still being developed to improve general performance.

While it was difficult to cover all of the topics with the appropriate depth, this review can serve as a compilation of different neurobiological learning rules as well as an aggregation of different implementations of spiking neural networks. The similarities between them are important, but all of the learning rules and networks stand on their own as well. These machine learning models can inform future experiments in neuroscience and the neuroscientific data can inform future models in machine learning.

1.19 Acknowledgements

Chapter 1, in full, is currently being prepared for submission for publication of the material. Morar, Vikash; Silva, Gabriel A. The dissertation author was the primary author of this material.

1.20 References

1. Van Gerven, M. Computational Foundations of Natural Intelligence. *Front. Comput. Neurosci.* **11**, 112 (2017).
2. Richards, B. A., Lillicrap, T. P., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R. P., De Berker, A., Ganguli, S., Gillon, C. J., Hafner, D., Kepecs, A., Kriegeskorte, N., Latham, P., Lindsay, G. W., Miller, K. D., Naud, R., Pack, C. C., Poirazi, P., Roelfsema, P., Sacramento, J., Saxe, A., Scellier, B., Schapiro, A. C., Senn, W., Wayne, G., Yamins, D., Zenke, F., Zylberberg, J., Therien, D. & Kording, K. P. A deep learning framework for neuroscience. *Nat. Neurosci.* **22**, 1761–1770 (2019).
3. Abbott, L. F. & Regehr, W. G. Synaptic computation. *Nature* **431**, 796–803 (2004).

4. Kastellakis, G., Cai, D. J., Mednick, S. C., Silva, A. J. & Poirazi, P. Synaptic clustering within dendrites: An emerging theory of memory formation. *Prog. Neurobiol.* **126**, 19–35 (2015).
5. Morris, R. G. M., Moser, E. I., Riedel, G., Martin, S. J., Sandin, J., Day, M. & O’Carroll, C. Elements of a neurobiological theory of the hippocampus: the role of activity-dependent synaptic plasticity in memory. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **358**, 773–786 (2003).
6. Takeuchi, T., Duzskiewicz, A. J. & Morris, R. G. M. The synaptic plasticity and memory hypothesis: encoding, storage and persistence. *Philos. Trans. R. Soc. B Biol. Sci.* **369**, 20130288 (2014).
7. Martin, S. J. & Morris, R. G. M. New life in an old idea: The synaptic plasticity and memory hypothesis revisited. *Hippocampus* **12**, 609–636 (2002).
8. Südhof, T. C. The cell biology of synapse formation. *J. Cell Biol.* **220**, e202103052 (2021).
9. Nishiyama, J. & Yasuda, R. Biochemical Computation for Spine Structural Plasticity. *Neuron* **87**, 63–75 (2015).
10. Bosch, M., Castro, J., Saneyoshi, T., Matsuno, H., Sur, M. & Hayashi, Y. Structural and Molecular Remodeling of Dendritic Spine Substructures during Long-Term Potentiation. *Neuron* **82**, 444–459 (2014).
11. Xiong, Q., Znamenskiy, P. & Zador, A. M. Selective corticostriatal plasticity during acquisition of an auditory discrimination task. *Nature* **521**, 348–351 (2015).
12. Hayashi-Takagi, A., Yagishita, S., Nakamura, M., Shirai, F., Wu, Y. I., Loshbaugh, A. L., Kuhlman, B., Hahn, K. M. & Kasai, H. Labelling and optical erasure of synaptic memory traces in the motor cortex. *Nature* **525**, 333–338 (2015).
13. Poort, J., Khan, A. G., Pachitariu, M., Nemri, A., Orsolich, I., Krupic, J., Bauza, M., Sahani, M., Keller, G. B., Mrsic-Flogel, T. D. & Hofer, S. B. Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. *Neuron* **86**, 1478–1490 (2015).
14. Goltstein, P. M., Coffey, E. B. J., Roelfsema, P. R. & Pennartz, C. M. A. In Vivo Two-Photon Ca²⁺ Imaging Reveals Selective Reward Effects on Stimulus-Specific Assemblies in Mouse Visual Cortex. *J. Neurosci.* **33**, 11540–11555 (2013).
15. Liu, X., Ramirez, S. & Tonegawa, S. Inception of a false memory by optogenetic manipulation of a hippocampal memory engram. *Philos. Trans. R. Soc. B Biol. Sci.* **369**, 20130142 (2014).

16. Nabavi, S., Fox, R., Proulx, C. D., Lin, J. Y., Tsien, R. Y. & Malinow, R. Engineering a memory with LTD and LTP. *Nature* **511**, 348–352 (2014).
17. Hugues, S., Chessel, A., Lena, I., Marsault, R. & Garcia, R. Prefrontal infusion of PD098059 immediately after fear extinction training blocks extinction-associated prefrontal synaptic plasticity and decreases prefrontal ERK2 phosphorylation. *Synapse* **60**, 280–287 (2006).
18. Simon, W., Hapfelmeier, G., Kochs, E., Zieglgänsberger, W. & Rammes, G. Isoflurane Blocks Synaptic Plasticity in the Mouse Hippocampus. *Anesthesiology* **94**, 1058–1065 (2001).
19. Younts, T. J., Monday, H. R., Dudok, B., Klein, M. E., Jordan, B. A., Katona, I. & Castillo, P. E. Presynaptic Protein Synthesis Is Required for Long-Term Plasticity of GABA Release. *Neuron* **92**, 479–492 (2016).
20. Yasumatsu, N., Matsuzaki, M., Miyazaki, T., Noguchi, J. & Kasai, H. Principles of Long-Term Dynamics of Dendritic Spines. *J. Neurosci.* **28**, 13592–13608 (2008).
21. Attardo, A., Fitzgerald, J. E. & Schnitzer, M. J. Impermanence of dendritic spines in live adult CA1 hippocampus. *Nature* **523**, 592–596 (2015).
22. Hebb, D. O. *The Organization of Behavior: A Neuropsychological Theory*. (Psychology Press, 2005).
23. Feldman, D. E. The Spike-Timing Dependence of Plasticity. *Neuron* **75**, 556–571 (2012).
24. Caporale, N. & Dan, Y. Spike Timing–Dependent Plasticity: A Hebbian Learning Rule. *Annu. Rev. Neurosci.* **31**, 25–46 (2008).
25. Debanne, D. & Inglebert, Y. Spike timing-dependent plasticity and memory. *Curr. Opin. Neurobiol.* **80**, 102707 (2023).
26. Tang, J., Yuan, F., Shen, X., Wang, Z., Rao, M., He, Y., Sun, Y., Li, X., Zhang, W., Li, Y., Gao, B., Qian, H., Bi, G., Song, S., Yang, J. J. & Wu, H. Bridging Biological and Artificial Neural Networks with Emerging Neuromorphic Devices: Fundamentals, Progress, and Challenges. *Adv. Mater.* **31**, 1902761 (2019).
27. Bi, G. & Poo, M. Synaptic Modifications in Cultured Hippocampal Neurons: Dependence on Spike Timing, Synaptic Strength, and Postsynaptic Cell Type. *J. Neurosci.* **18**, 10464–10472 (1998).
28. Bi, G. & Poo, M. Synaptic Modification by Correlated Activity: Hebb’s Postulate Revisited. *Annu. Rev. Neurosci.* **24**, 139–166 (2001).
29. Markram, H., Lübke, J., Frotscher, M. & Sakmann, B. Regulation of Synaptic Efficacy by Coincidence of Postsynaptic APs and EPSPs. *Science* **275**, 213–215 (1997).

30. Cui, Y., Prokin, I., Mendes, A., Berry, H. & Venance, L. Robustness of STDP to spike timing jitter. *Sci. Rep.* **8**, 8139 (2018).
31. Froemke, R. C. & Dan, Y. Spike-timing-dependent synaptic modification induced by natural spike trains. *Nature* **416**, 433–438 (2002).
32. Tazerart, S., Mitchell, D. E., Miranda-Rottmann, S. & Araya, R. A spike-timing-dependent plasticity rule for dendritic spines. *Nat. Commun.* **11**, 4276 (2020).
33. Skilling, Q. M., Eniwaye, B., Clawson, B. C., Shaver, J., Ognjanovski, N., Aton, S. J. & Zochowski, M. Acetylcholine-gated current translates wake neuronal firing rate information into a spike timing-based code in Non-REM sleep, stabilizing neural network dynamics during memory consolidation. *PLOS Comput. Biol.* **17**, e1009424 (2021).
34. Kasabov, N. K. NeuCube: A spiking neural network architecture for mapping, learning and understanding of spatio-temporal brain data. *Neural Netw.* **52**, 62–76 (2014).
35. Izhikevich, E. M. Simple model of spiking neurons. *IEEE Trans. Neural Netw.* **14**, 1569–1572 (2003).
36. Izhikevich, E. M. Polychronization: Computation with Spikes. *Neural Comput.* **18**, 245–282 (2006).
37. Lubenov, E. V. & Siapas, A. G. Decoupling through Synchrony in Neuronal Circuits with Propagation Delays. *Neuron* **58**, 118–131 (2008).
38. Tan, C. H., Cheu, E. Y., Hu, J., Yu, Q. & Tang, H. in *Neural Inf. Process.* (eds. Lu, B.-L., Zhang, L. & Kwok, J.) **7062**, 493–500 (Springer Berlin Heidelberg, 2011).
39. Diehl, P. U. & Cook, M. Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Front. Comput. Neurosci.* **9**, (2015).
40. Lecun, Y. Gradient-Based Learning Applied to Document Recognition. *Proc. IEEE* **86**, (1998).
41. Iakymchuk, T., Rosado-Muñoz, A., Guerrero-Martínez, J. F., Bataller-Mompeán, M. & Francés-Villora, J. V. Simplified spiking neural network architecture and STDP learning algorithm applied to image classification. *EURASIP J. Image Video Process.* **2015**, 4 (2015).
42. Paulun, L., Wendt, A. & Kasabov, N. A Retinotopic Spiking Neural Network System for Accurate Recognition of Moving Objects Using NeuCube and Dynamic Vision Sensors. *Front. Comput. Neurosci.* **12**, 42 (2018).

43. Kheradpisheh, S. R., Ganjtabesh, M., Thorpe, S. J. & Masquelier, T. STDP-based spiking deep convolutional neural networks for object recognition. *Neural Netw.* **99**, 56–67 (2018).
44. Masquelier, T. Unsupervised Learning of Visual Features through Spike Timing Dependent Plasticity. *PLoS Comput. Biol.* **3**, e31 (2007).
45. Ferré, P., Mamalet, F. & Thorpe, S. J. Unsupervised Feature Learning With Winner-Takes-All Based STDP. *Front. Comput. Neurosci.* **12**, 24 (2018).
46. Saunders, D. J., Patel, D., Hazan, H., Siegelmann, H. T. & Kozma, R. Locally connected spiking neural networks for unsupervised feature learning. *Neural Netw.* **119**, 332–340 (2019).
47. Brader, J. M., Senn, W. & Fusi, S. Learning Real-World Stimuli in a Neural Network with Spike-Driven Synaptic Dynamics. *Neural Comput.* **19**, 2881–2912 (2007).
48. Sjöström, P. J., Turrigiano, G. G. & Nelson, S. B. Rate, Timing, and Cooperativity Jointly Determine Cortical Synaptic Plasticity. *Neuron* **32**, 1149–1164 (2001).
49. Madadi Asl, M., Valizadeh, A. & Tass, P. A. Delay-Induced Multistability and Loop Formation in Neuronal Networks with Spike-Timing-Dependent Plasticity. *Sci. Rep.* **8**, 12068 (2018).
50. Nessler, B., Pfeiffer, M., Buesing, L. & Maass, W. Bayesian Computation Emerges in Generic Cortical Microcircuits through Spike-Timing-Dependent Plasticity. *PLoS Comput. Biol.* **9**, e1003037 (2013).
51. Nessler, B., Pfeiffer, M. & Maass, W. Hebbian Learning of Bayes Optimal Decisions.
52. Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T. & Maida, A. S. Deep Learning in Spiking Neural Networks. *Neural Netw.* **111**, 47–63 (2019).
53. Magee, J. C. & Grienberger, C. Synaptic Plasticity Forms and Functions. *Annu. Rev. Neurosci.* **43**, 95–117 (2020).
54. Ponulak, F. & Kasiński, A. Introduction to spiking neural networks: Information processing, learning and applications.
55. Pfeiffer, M. & Pfeil, T. Deep Learning With Spiking Neurons: Opportunities and Challenges. *Front. Neurosci.* **12**, 774 (2018).
56. Garg, N., Balafrej, I., Stewart, T. C., Portal, J.-M., Bocquet, M., Querlioz, D., Drouin, D., Rouat, J., Beilliard, Y. & Alibart, F. Voltage-dependent synaptic plasticity: Unsupervised probabilistic Hebbian plasticity rule based on neurons membrane potential. *Front. Neurosci.* **16**, 983950 (2022).

57. Jimenez-Romero, C., Sousa-Rodrigues, D. & Johnson, J. Designing Behaviour in Bio-inspired Robots Using Associative Topologies of Spiking-Neural-Networks. in *Proc. 9th EAI Int. Conf. Bio-Inspired Inf. Commun. Technol. Former. BIONETICS* (ACM, 2016). doi:10.4108/eai.3-12-2015.2262580
58. Dumesnil, E., Beaulieu, P.-O. & Boukadoum, M. Robotic implementation of classical and Operant Conditioning as a single STDP learning process. in *2016 Int. Jt. Conf. Neural Netw. IJCNN* 5241–5247 (IEEE, 2016). doi:10.1109/IJCNN.2016.7727892
59. Cyr, A., Thériault, F. & Chartier, S. Revisiting the XOR problem: a neurobotic implementation. *Neural Comput. Appl.* **32**, 9965–9973 (2020).
60. Iwadate, K., Suzuki, I., Watanabe, M., Yamamoto, M. & Furukawa, M. in *Soft Comput. Artif. Intell.* (eds. Cho, Y. I. & Matson, E. T.) **270**, 143–151 (Springer International Publishing, 2014).
61. Cyr, A., Boukadoum, M. & Thériault, F. Operant conditioning: a minimal components requirement in artificial spiking neurons designed for bio-inspired robot's controller. *Front. Neurobotics* **8**, (2014).
62. Chao, Y., Augenstein, P., Roennau, A., Dillmann, R. & Xiong, Z. Brain inspired path planning algorithms for drones. *Front. Neurobotics* **17**, 1111861 (2023).
63. Azimirad, V., Sani, M. F. & Ramezanlou, M. T. Unsupervised Learning of Target Attraction for Robots through Spike Timing Dependent Plasticity. *2017 IEEE 4th Int. Conf. Knowl.-Based Eng. Innov. KBEI Dec 22th 2017* 0428–0433 (2017).
64. Feng, H. & Zeng, Y. A brain-inspired robot pain model based on a spiking neural network. *Front. Neurobotics* **16**, 1025338 (2022).
65. Froemke, R. C. Plasticity of Cortical Excitatory-Inhibitory Balance. *Annu. Rev. Neurosci.* **38**, 195–219 (2015).
66. Edelmann, E., Cepeda-Prado, E. & Leßmann, V. Coexistence of Multiple Types of Synaptic Plasticity in Individual Hippocampal CA1 Pyramidal Neurons. *Front. Synaptic Neurosci.* **9**, (2017).
67. Tsetsenis, T., Broussard, J. I. & Dani, J. A. Dopaminergic regulation of hippocampal plasticity, learning, and memory. *Front. Behav. Neurosci.* **16**, 1092420 (2023).
68. Zhang, J.-C., Lau, P.-M. & Bi, G.-Q. Gain in sensitivity and loss in temporal contrast of STDP by dopaminergic modulation at hippocampal synapses. *Proc. Natl. Acad. Sci.* **106**, 13028–13033 (2009).
69. Brzosko, Z., Schultz, W. & Paulsen, O. Retroactive modulation of spike timing-dependent plasticity by dopamine. *eLife* **4**, e09685 (2015).

70. Yagishita, S., Hayashi-Takagi, A., Ellis-Davies, G. C. R., Urakubo, H., Ishii, S. & Kasai, H. A critical time window for dopamine actions on the structural plasticity of dendritic spines. *Science* **345**, 1616–1620 (2014).
71. Lisman, J., Grace, A. A. & Duzel, E. A neoHebbian framework for episodic memory; role of dopamine-dependent late LTP. *Trends Neurosci.* **34**, 536–547 (2011).
72. Bethus, I., Tse, D. & Morris, R. G. M. Dopamine and Memory: Modulation of the Persistence of Memory for Novel Hippocampal NMDA Receptor-Dependent Paired Associates. *J. Neurosci.* **30**, 1610–1618 (2010).
73. Picciotto, M. R., Higley, M. J. & Mineur, Y. S. Acetylcholine as a Neuromodulator: Cholinergic Signaling Shapes Nervous System Function and Behavior. *Neuron* **76**, 116–129 (2012).
74. Hangya, B., Ranade, S. P., Lorenc, M. & Kepecs, A. Central Cholinergic Neurons Are Rapidly Recruited by Reinforcement Feedback. *Cell* **162**, 1155–1168 (2015).
75. Chubykin, A. A., Roach, E. B., Bear, M. F. & Shuler, M. G. H. A Cholinergic Mechanism for Reward Timing within Primary Visual Cortex. *Neuron* **77**, 723–735 (2013).
76. Warburton, E. C., Koder, T., Cho, K., Massey, P. V., Duguid, G., Barker, G. R. I., Aggleton, J. P., Bashir, Z. I. & Brown, M. W. Cholinergic Neurotransmission Is Essential for Perirhinal Cortical Plasticity and Recognition Memory. *Neuron* **38**, 987–996 (2003).
77. Brzosko, Z., Zannone, S., Schultz, W., Clopath, C. & Paulsen, O. Sequential neuromodulation of Hebbian plasticity offers mechanism for effective reward-based navigation. *eLife* **6**, e27756 (2017).
78. Liu, Z., Zhou, J., Li, Y., Hu, F., Lu, Y., Ma, M., Feng, Q., Zhang, J., Wang, D., Zeng, J., Bao, J., Kim, J.-Y., Chen, Z.-F., El Mestikawy, S. & Luo, M. Dorsal Raphe Neurons Signal Reward through 5-HT and Glutamate. *Neuron* **81**, 1360–1374 (2014).
79. Nakamura, K., Matsumoto, M. & Hikosaka, O. Reward-Dependent Modulation of Neuronal Activity in the Primate Dorsal Raphe Nucleus. *J. Neurosci.* **28**, 5331–5343 (2008).
80. He, K., Huertas, M., Hong, S. Z., Tie, X., Hell, J. W., Shouval, H. & Kirkwood, A. Distinct Eligibility Traces for LTP and LTD in Cortical Synapses. *Neuron* **88**, 528–538 (2015).
81. Redondo, R. L. & Morris, R. G. M. Making memories last: the synaptic tagging and capture hypothesis. *Nat. Rev. Neurosci.* **12**, 17–30 (2011).
82. Gerstner, W., Lehmann, M., Liakoni, V., Corneil, D. & Brea, J. Eligibility Traces and Plasticity on Behavioral Time Scales: Experimental Support of NeoHebbian Three-Factor Learning Rules. *Front. Neural Circuits* **12**, 53 (2018).

83. Legenstein, R., Chase, S. M., Schwartz, A. B. & Maass, W. A Reward-Modulated Hebbian Learning Rule Can Explain Experimentally Observed Network Reorganization in a Brain Control Task. *J. Neurosci.* **30**, 8400–8410 (2010).
84. Hao, Y., Huang, X., Dong, M. & Xu, B. A biologically plausible supervised learning method for spiking neural networks using the symmetric STDP rule. *Neural Netw.* **121**, 387–395 (2020).
85. Fang, H., Zeng, Y. & Zhao, F. Brain Inspired Sequences Production by Spiking Neural Networks With Reward-Modulated STDP. *Front. Comput. Neurosci.* **15**, 612041 (2021).
86. Apolinario, M. P. E. & Roy, K. S-TLLR: STDP-inspired Temporal Local Learning Rule for Spiking Neural Networks. Preprint at <http://arxiv.org/abs/2306.15220> (2023)
87. Hoerzer, G. M., Legenstein, R. & Maass, W. Emergence of Complex Computational Structures From Chaotic Neural Networks Through Reward-Modulated Hebbian Learning. *Cereb. Cortex* **24**, 677–690 (2014).
88. Florian, R. V. Reinforcement Learning Through Modulation of Spike-Timing-Dependent Synaptic Plasticity. *Neural Comput.* **19**, 1468–1502 (2007).
89. Evans, R. Reinforcement Learning in a Neurally Controlled Robot Using Dopamine Modulated STDP. Preprint at <http://arxiv.org/abs/1502.06096> (2015)
90. Shim, M. S. & Li, P. Biologically inspired reinforcement learning for mobile robot collision avoidance. in *2017 Int. Jt. Conf. Neural Netw. IJCNN* 3098–3105 (IEEE, 2017). doi:10.1109/IJCNN.2017.7966242
91. Bing, Z., Meschede, C., Huang, K., Chen, G., Rohrbein, F., Akl, M. & Knoll, A. End to End Learning of Spiking Neural Network Based on R-STDP for a Lane Keeping Vehicle. in *2018 IEEE Int. Conf. Robot. Autom. ICRA* 4725–4732 (IEEE, 2018). doi:10.1109/ICRA.2018.8460482
92. Giraldo, N. S., Isaza, S. & Velásquez, R. A. Sailboat navigation control system based on spiking neural networks. *Control Theory Technol.* **21**, 489–504 (2023).
93. Clawson, T. S., Ferrari, S., Fuller, S. B. & Wood, R. J. Spiking neural network (SNN) control of a flapping insect-scale robot. in *2016 IEEE 55th Conf. Decis. Control CDC* 3381–3388 (IEEE, 2016). doi:10.1109/CDC.2016.7798778
94. Wang, H.-X., Gerkin, R. C., Nauen, D. W. & Bi, G.-Q. Coactivation and timing-dependent integration of synaptic potentiation and depression. *Nat. Neurosci.* **8**, 187–193 (2005).
95. Froemke, R. C. Temporal modulation of spike-timing-dependent plasticity. *Front. Synaptic Neurosci.* (2010). doi:10.3389/fnsyn.2010.00019

96. Pfister, J.-P. & Gerstner, W. Triplets of Spikes in a Model of Spike Timing-Dependent Plasticity. *J. Neurosci.* **26**, 9673–9682 (2006).
97. Bienenstock, L., Cooper, N. & Munro, W. THEORY FOR THE DEVELOPMENT OF NEURON SELECTIVITY: ORIENTATION SPECIFICITY AND BINOCULAR INTERACTION IN VISUAL CORTEX.
98. Wilmes, K. A. & Clopath, C. Dendrites help mitigate the plasticity-stability dilemma. *Sci. Rep.* **13**, 6543 (2023).
99. Krunglevicius, D. Modified STDP Triplet Rule Significantly Increases Neuron Training Stability in the Learning of Spatial Patterns. *Adv. Artif. Neural Syst.* **2016**, 1–12 (2016).
100. Wang, Z., Zeng, T., Ren, Y., Lin, Y., Xu, H., Zhao, X., Liu, Y. & Ielmini, D. Toward a generalized Bienenstock-Cooper-Munro rule for spatiotemporal learning via triplet-STDP in memristive devices. *Nat. Commun.* **11**, 1510 (2020).
101. Abbott, L. F. & Nelson, S. B. Synaptic plasticity: taming the beast. *Nat. Neurosci.* **3**, 1178–1183 (2000).
102. Froemke, R. C., Poo, M. & Dan, Y. Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* **434**, 221–225 (2005).
103. Spruston, N. Pyramidal neurons: dendritic structure and synaptic integration. *Nat. Rev. Neurosci.* **9**, 206–221 (2008).
104. Sjöström, P. J. & Häusser, M. A Cooperative Switch Determines the Sign of Synaptic Plasticity in Distal Dendrites of Neocortical Pyramidal Neurons. *Neuron* **51**, 227–238 (2006).
105. Letzkus, J. J., Kampa, B. M. & Stuart, G. J. Learning Rules for Spike Timing-Dependent Plasticity Depend on Dendritic Synapse Location. *J. Neurosci.* **26**, 10420–10429 (2006).
106. Koch, G., Ponzo, V., Di Lorenzo, F., Caltagirone, C. & Veniero, D. Hebbian and Anti-Hebbian Spike-Timing-Dependent Plasticity of Human Cortico-Cortical Connections. *J. Neurosci.* **33**, 9725–9733 (2013).
107. Burbank, K. S. & Kreiman, G. Depression-Biased Reverse Plasticity Rule Is Required for Stable Learning at Top-Down Connections. *PLoS Comput. Biol.* **8**, e1002393 (2012).
108. Burbank, K. S. Mirrored STDP Implements Autoencoder Learning in a Network of Spiking Neurons. *PLoS Comput. Biol.* **11**, e1004566 (2015).

109. Mozafari, M., Kheradpisheh, S. R., Masquelier, T., Nowzari-Dalini, A. & Ganjtabesh, M. First-Spike-Based Visual Categorization Using Reward-Modulated STDP. *IEEE Trans. Neural Netw. Learn. Syst.* **29**, 6178–6190 (2018).
110. Tavanaei, A. & Maida, A. S. BP-STDP: Approximating Backpropagation using Spike Timing Dependent Plasticity. Preprint at <http://arxiv.org/abs/1711.04214> (2018)
111. Saranirad, V., Dora, S., McGinnity, T. M. & Coyle, D. Assembly-based STDP: A New Learning Rule for Spiking Neural Networks Inspired by Biological Assemblies. in *2022 Int. Jt. Conf. Neural Netw. IJCNN* 1–7 (IEEE, 2022). doi:10.1109/IJCNN55064.2022.9891925
112. Neftci, E., Das, S., Pedroni, B., Kreuz-Delgado, K. & Cauwenberghs, G. Event-driven contrastive divergence for spiking neuromorphic systems. *Front. Neurosci.* **7**, (2014).
113. Schuemann, A., Klawiter, A., Bonhoeffer, T. & Wierenga, C. J. Structural plasticity of GABAergic axons is regulated by network activity and GABAA receptor activation. *Front. Neural Circuits* **7**, (2013).
114. Kodangattil, J. N., Dacher, M., Authement, M. E. & Nugent, F. S. Spike timing-dependent plasticity at GABAergic synapses in the ventral tegmental area. *J. Physiol.* **591**, 4699–4710 (2013).
115. Udakis, M., Pedrosa, V., Chamberlain, S. E. L., Clopath, C. & Mellor, J. R. Interneuron-specific plasticity at parvalbumin and somatostatin inhibitory synapses onto CA1 pyramidal neurons shapes hippocampal output. *Nat. Commun.* **11**, 4395 (2020).
116. Hennequin, G., Agnes, E. J. & Vogels, T. P. Inhibitory Plasticity: Balance, Control, and Codependence. *Annu. Rev. Neurosci.* **40**, 557–579 (2017).
117. Pelkey, K. A., Chittajallu, R., Craig, M. T., Tricoire, L., Wester, J. C. & McBain, C. J. Hippocampal GABAergic Inhibitory Interneurons. *Physiol. Rev.* **97**, 1619–1747 (2017).
118. Yap, E.-L., Pettit, N. L., Davis, C. P., Nagy, M. A., Harmin, D. A., Golden, E., Dagliyan, O., Lin, C., Rudolph, S., Sharma, N., Griffith, E. C., Harvey, C. D. & Greenberg, M. E. Bidirectional perisomatic inhibitory plasticity of a Fos neuronal network. *Nature* **590**, 115–121 (2021).
119. Vickers, E. D., Clark, C., Osypenko, D., Fratzl, A., Kochubey, O., Bettler, B. & Schneggenburger, R. Parvalbumin-Interneuron Output Synapses Show Spike-Timing-Dependent Plasticity that Contributes to Auditory Map Remodeling. *Neuron* **99**, 720-735.e6 (2018).
120. Haas, J. S., Nowotny, T. & Abarbanel, H. D. I. Spike-Timing-Dependent Plasticity of Inhibitory Synapses in the Entorhinal Cortex. *J. Neurophysiol.* **96**, 3305–3313 (2006).

121. D'amour, J. A. & Froemke, R. C. Inhibitory and Excitatory Spike-Timing-Dependent Plasticity in the Auditory Cortex. *Neuron* **86**, 514–528 (2015).
122. Field, R. E., D'amour, J. A., Tremblay, R., Miehl, C., Rudy, B., Gjorgjieva, J. & Froemke, R. C. Heterosynaptic Plasticity Determines the Set Point for Cortical Excitatory-Inhibitory Balance. *Neuron* **106**, 842-854.e4 (2020).
123. Bourne, J. N. & Harris, K. M. Coordination of size and number of excitatory and inhibitory synapses results in a balanced structural plasticity along mature hippocampal CA1 dendrites during LTP. *Hippocampus* **21**, 354–373 (2011).
124. Castillo, P. E., Chiu, C. Q. & Carroll, R. C. Long-term plasticity at inhibitory synapses. *Curr. Opin. Neurobiol.* **21**, 328–338 (2011).
125. Winters, B. D. & Golding, N. L. Glycinergic Inhibitory Plasticity in Binaural Neurons Is Cumulative and Gated by Developmental Changes in Action Potential Backpropagation. *Neuron* **98**, 166-178.e2 (2018).
126. Wang, L. & Maffei, A. Inhibitory Plasticity Dictates the Sign of Plasticity at Excitatory Synapses. *J. Neurosci.* **34**, 1083–1093 (2014).
127. Vogels, T. P., Froemke, R. C., Doyon, N., Gilson, M., Haas, J. S., Liu, R., Maffei, A., Miller, P., Wierenga, C. J., Woodin, M. A., Zenke, F. & Sprekeler, H. Inhibitory synaptic plasticity: spike timing-dependence and putative network function. *Front. Neural Circuits* **7**, (2013).
128. Kaleb, K., Pedrosa, V. & Clopath, C. Network-centered homeostasis through inhibition maintains hippocampal spatial map and cortical circuit function. *Cell Rep.* **36**, 109577 (2021).
129. Wilmes, K. A. & Clopath, C. Inhibitory microcircuits for top-down plasticity of sensory representations. *Nat. Commun.* **10**, 5055 (2019).
130. Vogels, T. P., Sprekeler, H., Zenke, F., Clopath, C. & Gerstner, W. Inhibitory Plasticity Balances Excitation and Inhibition in Sensory Pathways and Memory Networks. *Science* **334**, 1569–1573 (2011).
131. Aljadeff, J., D'amour, J., Field, R. E., Froemke, R. C. & Clopath, C. Cortical credit assignment by Hebbian, neuromodulatory and inhibitory plasticity. Preprint at <http://arxiv.org/abs/1911.00307> (2019)
132. Payeur, A., Guerguiev, J., Zenke, F., Richards, B. A. & Naud, R. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *Nat. Neurosci.* **24**, 1010–1019 (2021).
133. Sacramento, J., Costa, R. P., Bengio, Y. & Senn, W. Dendritic cortical microcircuits approximate the backpropagation algorithm.

134. Srinivasa, N. & Cho, Y. Unsupervised discrimination of patterns in spiking neural networks with excitatory and inhibitory synaptic plasticity. *Front. Comput. Neurosci.* **8**, (2014).
135. Sporea, I. & Grüning, A. Supervised Learning in Multilayer Spiking Neural Networks. *Neural Comput.* **25**, 473–509 (2013).
136. Amirshahi, A. & Hashemi, M. ECG Classification Algorithm Based on STDP and R-STDP Neural Networks for Real-Time Monitoring on Ultra Low-Power Personal Wearable Devices. *IEEE Trans. Biomed. Circuits Syst.* **13**, 1483–1493 (2019).
137. Sancho, L., Contreras, M. & Allen, N. J. Glia as sculptors of synaptic plasticity. *Neurosci. Res.* **167**, 17–29 (2021).
138. Perea, G., Navarrete, M. & Araque, A. Tripartite synapses: astrocytes process and control synaptic information. *Trends Neurosci.* **32**, 421–431 (2009).
139. Hösli, L., Binini, N., Ferrari, K. D., Thieren, L., Looser, Z. J., Zuend, M., Zanker, H. S., Berry, S., Holub, M., Möbius, W., Ruhwedel, T., Nave, K.-A., Giaume, C., Weber, B. & Saab, A. S. Decoupling astrocytes in adult mice impairs synaptic plasticity and spatial learning. *Cell Rep.* **38**, 110484 (2022).
140. Araque, A., Carmignoto, G., Haydon, P. G., Oliet, S. H. R., Robitaille, R. & Volterra, A. Gliotransmitters Travel in Time and Space. *Neuron* **81**, 728–739 (2014).
141. Kastanenka, K. V., Moreno-Bote, R., De Pittà, M., Perea, G., Eraso-Pichot, A., Masgrau, R., Poskanzer, K. E. & Galea, E. A roadmap to integrate astrocytes into Systems Neuroscience. *Glia* **68**, 5–26 (2020).
142. Manninen, T., Aćimović, J. & Linne, M.-L. Analysis of Network Models with Neuron-Astrocyte Interactions. *Neuroinformatics* **21**, 375–406 (2023).
143. Alvarez-Gonzalez, S., Cedron, F., Pazos, A. & Porto-Pazos, A. B. Artificial glial cells in artificial neuronal networks: a systematic review. *Artif. Intell. Rev.* **56**, 2651–2666 (2023).
144. Gordleeva, S. Yu., Ermolaeva, A. V., Kastalskiy, I. A. & Kazantsev, V. B. Astrocyte as Spatiotemporal Integrating Detector of Neuronal Activity. *Front. Physiol.* **10**, 294 (2019).
145. De Pittà, M., Volman, V., Berry, H. & Ben-Jacob, E. A Tale of Two Stories: Astrocyte Regulation of Synaptic Depression and Facilitation. *PLoS Comput. Biol.* **7**, e1002293 (2011).
146. De Pittà, M. & Brunel, N. *Multiple forms of working memory emerge from synapse-astrocyte interactions*. (Neuroscience, 2021). doi:10.1101/2021.03.25.436819

147. Gordleeva, S. Yu., Tsybina, Y. A., Krivonosov, M. I., Ivanchenko, M. V., Zaikin, A. A., Kazantsev, V. B. & Gorban, A. N. Modeling Working Memory in a Spiking Neuron Network Accompanied by Astrocytes. *Front. Cell. Neurosci.* **15**, 631485 (2021).
148. Gordleeva, S. Y., Stasenko, S. V., Semyanov, A. V., Dityatev, A. E. & Kazantsev, V. B. Bi-directional astrocytic regulation of neuronal activity within a network. *Front. Comput. Neurosci.* **6**, (2012).
149. Stasenko, S. V. & Kazantsev, V. B. Dynamic Image Representation in a Spiking Neural Network Supplied by Astrocytes. *Mathematics* **11**, 561 (2023).
150. Stasenko, S. & Kazantsev, V. in *Adv. Neural Comput. Mach. Learn. Cogn. Res. VI* (eds. Kryzhanovsky, B., Dunin-Barkowski, W., Redko, V. & Tiumentsev, Y.) **1064**, 200–206 (Springer International Publishing, 2023).
151. Ivanov, V. A. & Michmizos, K. P. Increasing Liquid State Machine Performance with Edge-of-Chaos Dynamics Organized by Astrocyte-modulated Plasticity.
152. Kozachkov, L., Kastanenko, K. V. & Krotov, D. Building transformers from neurons and astrocytes. *Proc. Natl. Acad. Sci.* **120**, e2219150120 (2023).
153. Naghieh, P., Delavar, A., Amiri, M. & Peremans, H. Astrocyte’s self-repairing characteristics improve working memory in spiking neuronal networks. *iScience* **26**, 108241 (2023).
154. Tang, G., Polykretis, I. E., Ivanov, V. A., Shah, A. & Michmizos, K. P. Introducing Astrocytes on a Neuromorphic Processor: Synchronization, Local Plasticity and Edge of Chaos. in *Proc. 7th Annu. Neuro-Inspired Comput. Elem. Workshop* 1–9 (ACM, 2019). doi:10.1145/3320288.3320302
155. Polykretis, I., Tang, G. & Michmizos, K. P. An Astrocyte-Modulated Neuromorphic Central Pattern Generator for Hexapod Robot Locomotion on Intel’s Loihi. in *Int. Conf. Neuromorphic Syst. 2020* 1–9 (ACM, 2020). doi:10.1145/3407197.3407205
156. Bechler, M. E., Swire, M. & French-Constant, C. Intrinsic and adaptive myelination—A sequential mechanism for smart wiring in the brain. *Dev. Neurobiol.* **78**, 68–79 (2018).
157. Hofstetter, S., Tavor, I., Tzur Moryosef, S. & Assaf, Y. Short-Term Learning Induces White Matter Plasticity in the Fornix. *J. Neurosci.* **33**, 12844–12850 (2013).
158. Gibson, E. M., Purger, D., Mount, C. W., Goldstein, A. K., Lin, G. L., Wood, L. S., Inema, I., Miller, S. E., Bieri, G., Zuchero, J. B., Barres, B. A., Woo, P. J., Vogel, H. & Monje, M. Neuronal Activity Promotes Oligodendrogenesis and Adaptive Myelination in the Mammalian Brain. *Science* **344**, 1252304 (2014).

159. Cullen, C. L., Pepper, R. E., Clutterbuck, M. T., Pitman, K. A., Oorschot, V., Auderset, L., Tang, A. D., Ramm, G., Emery, B., Rodger, J., Jolivet, R. B. & Young, K. M. Periaxonal and nodal plasticities modulate action potential conduction in the adult mouse brain. *Cell Rep.* **34**, 108641 (2021).
160. Noori, R., Park, D., Griffiths, J. D., Bells, S., Frankland, P. W., Mabbott, D. & Lefebvre, J. Activity-dependent myelination: A glial mechanism of oscillatory self-organization in large-scale brain networks. *Proc. Natl. Acad. Sci.* **117**, 13227–13237 (2020).
161. Talidou, A., Frankland, P. W., Mabbott, D. & Lefebvre, J. Homeostatic coordination and up-regulation of neural activity by activity-dependent myelination. *Nat. Comput. Sci.* **2**, 665–676 (2022).
162. Hong, C., Wei, X., Wang, J., Deng, B., Yu, H. & Che, Y. Training Spiking Neural Networks for Cognitive Tasks: A Versatile Framework Compatible With Various Temporal Codes. *IEEE Trans. Neural Netw. Learn. Syst.* **31**, 1285–1296 (2020).
163. Chaplinskaia, N. V. & Bazenkov, N. I. Axonal Myelination as a Mechanism for Unsupervised Learning in Spiking Neural Networks.
164. Colonna, M. & Butovsky, O. Microglia Function in the Central Nervous System During Health and Neurodegeneration. *Annu. Rev. Immunol.* **35**, 441–468 (2017).
165. Piochon, C., Kano, M. & Hansel, C. LTD-like molecular pathways in developmental synaptic pruning. *Nat. Neurosci.* **19**, 1299–1310 (2016).
166. Hiratani, N. & Fukai, T. Hebbian Wiring Plasticity Generates Efficient Network Structures for Robust Inference with Synaptic Weight Plasticity. *Front. Neural Circuits* **10**, (2016).
167. Roy, S. & Basu, A. An Online Unsupervised Structural Plasticity Algorithm for Spiking Neural Networks. *IEEE Trans. Neural Netw. Learn. Syst.* **28**, 900–910 (2017).
168. Hussain, S. & Basu, A. Multiclass Classification by Adaptive Network of Dendritic Neurons with Binary Synapses Using Structural Plasticity. *Front. Neurosci.* **10**, (2016).
169. Navlakha, S., Barth, A. L. & Bar-Joseph, Z. Decreasing-Rate Pruning Optimizes the Construction of Efficient and Robust Distributed Networks. *PLOS Comput. Biol.* **11**, e1004347 (2015).

Chapter 2 A Computational Model for Storing and Embedding Data in Dynamically Evolving Network Synaptic Structures

2.1 Abstract

Spike-timing dependent plasticity (STDP) is widely accepted as a mechanism through which the brain can learn information. Basing synaptic changes on the timing between spikes enhances contributing edges within a network. While STDP rules control the evolution of networks, most research focuses on spiking rates or specific activation paths when evaluating learned information. While imaging studies demonstrate physical changes to synapses due to STDP, these changes have not been interrogated based on their embedding capacity of a stimulus. Here, we show that networks with biological features can embed stimulus information into their synaptic weights. We use a k-nearest neighbor algorithm on the synaptic weights of independent networks to identify their stimulus with high accuracy based on local neighborhoods. While spike rates and timings remain useful, structural embeddings represent a new way to integrate information within a biological network. Our results demonstrate that there may be value in observing these changes directly. Beyond computational applications, this analysis may also inform investigation into neuroscience. Research is underway on the potential of astrocytes to integrate synapses in the brain and communicate that information elsewhere. In addition, observations of these synaptic embeddings may lead to novel therapies for memory disorders, such as transient epileptic amnesia.

2.2 Introduction

The current state of the art in machine learning (ML) fundamentally relies on variations of gradient descent algorithms that minimize appropriately constructed cost or loss functions [1, 2]. These methods take advantage of the stochastic convergence of weights in artificial neural

networks in order to capture non-trivial latent statistical associations that encode inputs and relationships between data. A critical component of this process is the need to train ML systems in either supervised or unsupervised ways by exposing them to large numbers of training examples.

In many real-life scenarios, training requirements pose significant challenges. There may be insufficient data (or not enough high-quality data) for training. Or tasks such as learning and classification may need to be done ‘on the fly’ in near real-time to support just-in-time inference or decision-making. While some existing ML can perform near real-time learning, it still requires expensive pre-trained models or access to continued real-time streams of sufficient high-quality data [3, 4]. Current large data, large compute, and large model trends in ML cannot continue to scale indefinitely. And at least under some conditions, there is a real need for more data, compute, and energy-efficient paradigms for ML. The development of new algorithms that can support this could emerge from models of neural computation in the biological brain. The human brain, in particular, has the ability to achieve real-time inference and decision-making in the face of very limited data that, while never perfect, is often sufficient for real-world problem-solving in a way that current ML still cannot match.

In this work, we took inspiration from foundational neurophysiology and explored the feasibility of learning at the scale of synaptic weights in a geometric network model of neurobiological spatial and temporal summation. This model can efficiently encode arbitrary inputs in an unsupervised way without prior training or exposure to any a priori data. There is no readout layer to our model, but rather, inputs are encoded by the dynamics of the networks themselves, which are then represented in a metric phase space on which downstream statistics and machine learning can be applied. Furthermore, the model lends itself to neurobiological

simulations and numerical experiments capable of exploring the computational properties of the underlying neurophysiology itself — specifically, the data embedding space supported by adaptive weight changes.

To achieve this, the model we have developed takes advantage of inherent signaling latencies produced by the spatial and structural geometry of networks when combined with finite signaling speeds (conduction velocities) and node (neuronal) refractory periods. We show how spike-timing dependent plasticity (STDP) learning rules evolve the dynamics of the network as a function of vectors that encode inputs. Because the dynamics can be captured in a metric space, distances between points in this space capture inherent relationships (i.e., similarities) between different inputs and data without having to train the network.

The dependency of the model on learning by synaptic plasticity taking place on the weights via STDP is critical [5, 6]. The biological brain requires state changes to transfer information [7–12], and spike-timing dependent plasticity rules seem to be critical for implementing such state changes. In both excitatory and inhibitory neurons, potentiation and depression are essential to adapting networks based on the stimulus presented to them [13, 14]. However, excitatory and inhibitory neurons need different learning rules for potentiation and depression, as their synapses strengthen from opposite postsynaptic neuron outcomes [15], a fact we took advantage of in the model here.

Specifically, we used the leaky integrate-and-fire neuron model [16] with refractory periods [17] and axonal delays to incorporate geometry into the network [18]. These axonal delays were based on measurements of murine hippocampal axon lengths and conduction velocities [19, 20]. We implemented inhibition within the network at a percentage measured in

rat hippocampus [21]. We incorporated different STDP rules for excitatory [22] versus inhibitory edges [14] based on murine hippocampal data.

While potentiation is associated with memory formation and learning, recent evidence suggests depression may be equally important, particularly for tasks involving object recognition [23–28]. While depression is often used to stabilize weights after potentiation, depression is also used by the brain to intentionally reduce weights to encode information [29, 30]. Recent experiments also suggest that STDP may embed stimuli into the structure of the network itself [31, 32]. STDP can augment AMPA receptor expression on the post-synaptic dendrite in response to a stimulus [6, 33]. Potentiation and depression are appealing candidates for memory formation via changes at the dendrites [34–38].




Figure 2.1 Overview of the network model and analysis methods. (A) We used the leaky-integrate-and-fire model of a neuron to construct networks with a time delay (temporal latency) and weights for each edge. MNIST images consisting of 784 pixels were fed into the input layer with each input neuron corresponding to one pixel. Each input layer neuron connected to each recurrent layer neuron with a different delay. As the recurrent layer neurons activate, they send signals to each other with varying delays for signals propagating within the layer. As signals travel within the recurrent layer, a biologically-inspired STDP rule modulated the weights of each edge. (B) For our simulations, we randomly initialized many identical networks before stimulating them with different MNIST images to run in parallel. As the simulations progressed, the STDP rules resulted in similar modifications between networks stimulated by similar images. The weight values of each edge were concatenated into a weight vector. Weight vectors from similar stimuli traveled along similar trajectories. (C) During our simulation, we ran concatenated weight vectors through a k-nearest neighbors algorithm. This allowed us to classify vectors from different simulations with unknown stimuli based on their five nearest neighbors to a known stimuli representing a label.

2.3 Methods

2.3.1 Neuron Model

Figure 2.1 provides an overview of our model and approach. We used an extension of the leaky integrate-and-fire neuron model [16] in combination with a network model of neural signaling that explicitly takes into account axon geometry and path lengths that when combined with conduction velocities produce signaling dynamic latencies [18, 39]. Each neuron accumulates membrane potential from arriving signals while its overall potential decays away (leaks) as a function of time. The interplay between the frequency of fractionally contributing amplitudes from arriving signals with a decaying potential determines when and how that neuron

fires. Using the subscript i representing the postsynaptic neuron and subscript j representing any presynaptic neuron that has sent a signal, the membrane potential can be expressed as:

$$V_i(t) = \sum_j \gamma(w_{ji}, t_j) \quad (1)$$

The membrane potential V_i equals the sum of the weights of each incoming signal w_{ji} after they run through the decay function γ based on their respective arrival times t_j . Once the threshold potential is reached, the neuron fires along all its edges and then becomes refractory for a period of time before its membrane potential resets [17].

The weights of each edge are modulated using spike-timing dependent plasticity (STDP) rules derived from biological neurons [40]. Because excitatory and inhibitory edges need to be updated differently in response to network activity, a different STDP rule was used for each (Figure 2.2; see [14, 22]). For excitatory edges, potentiation occurs when a postsynaptic neuron fires soon after a signal arrives and depression occurs when the postsynaptic neuron is refractory when a signal arrives. For inhibitory edges, potentiation (i.e., making the edge more inhibitory) occurs when a postsynaptic neuron does not fire soon after a signal arrives, and depression occurs when a postsynaptic neuron fires near the signal arrival.

In addition to the arrival and departure times of signals, variables for the amplitude of the weight change and the time constants for scaling the temporal effect of weight changes play important roles. These variables provide direct control over how large each weight change can be and how dependent or sensitive each weight change is to the time difference between the arriving signal relative to the firing of the neuron.

For excitatory edges, the STDP rule we used can be expressed as:

$$\begin{aligned} \Delta w &= A^- e^{-\frac{\Delta t}{\tau^-}} \text{ if } \Delta t < 0 \\ \Delta w &= A^+ e^{\frac{\Delta t}{\tau^+}} \text{ if } \Delta t > 0 \end{aligned} \quad (2)$$

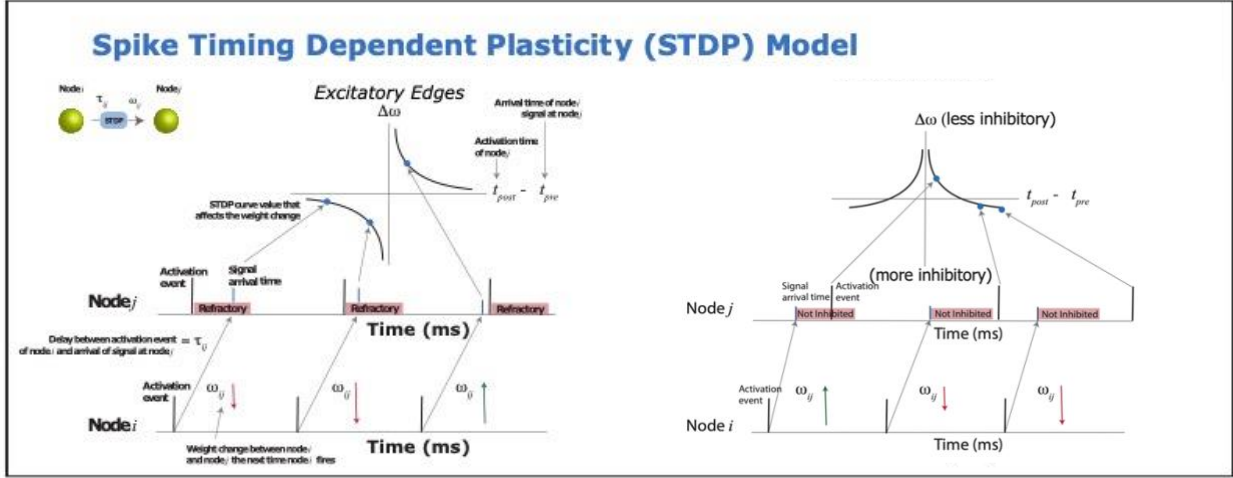


Figure 2.2 Biologically-inspired STDP rules we used in our networks. If an excitatory signal from a presynaptic neuron arrives at a postsynaptic neuron before it fires, the weight of the edge between them increases. If the signal arrives at the postsynaptic neuron while it is refractory, the weight of that edge decreases. Inhibitory edges are modified with a different biologically-inspired STDP rule. If the postsynaptic neuron fires close to when the presynaptic signal arrives, the weight of the inhibitory edge decreases. If a sufficiently long time passes after the presynaptic signal arrives without a postsynaptic activation, then the magnitude of the inhibitory edge increases.

In this expression, Δt is the time of postsynaptic neuron firing minus the time of presynaptic neuron firing. A^+ and A^- represent the amplitudes for potentiation and depression, respectively. τ^+ and τ^- are the STDP time constants as previously described. Δw represents how much the edge weight will change in the context of STDP-mediated potentiation and depression.

For inhibitory edges, weight changes are given by:

$$\Delta w = \eta A e^{-\Delta t/\tau} - \eta \alpha \quad (3)$$

Similar to equation 2, Δt is the time of postsynaptic neuron firing minus the time of presynaptic neuron firing. In contrast, though, the variable A only needs to take on one value for inhibitory edges, as potentiation and depression are weighted equally (i.e., they have equal amplitudes). Similarly, there is only one τ time constant, as the STDP curve is symmetric along the y-axis. η is a variable unique to inhibitory edges. It controls the learning rate of the edges by scaling each

weight change. α is a small value that is subtracted from the weight when there is a presynaptic firing. This allows for Δw to be negative when Δt has a large magnitude in either direction. Also, for clarification, a negative Δw would result in a weaker inhibitory edge.

The learning parameters for our STDP rules were taken from data fits to physiological murine neuron experiments [41]. Based on our STDP rule for excitatory edges, the choice of parameters reflected the amplitude of depression changes being greater than the amplitude of potentiation changes [5, 42]. This allowed for depression to control the activity of the network to avoid signals propagating indefinitely. For the STDP rule for inhibitory edges, the parameters weighted potentiation and depression equally, as there was no reason to augment inhibitory edges differently based on whether they were successful at reducing network activity or not [43].

The geometric embedding of our model is quite consequential here. Each edge having a specific length varied the amount of time it took for signals to reach their target neurons [44]. All signals traveled at the same speed (i.e., a constant conduction velocity), so the offset in their arrival times was based solely on the edge path lengths. The latencies (temporal delays) along the edges were taken from a normal distribution around $730\mu\text{m}$ [19], with a conduction velocity of 1.68 m/s [20].

The initial weights were chosen from a uniform distribution, such that 80% of the edges were excitatory and ranged from the threshold potential down to just above 0, and 20% of the edges were inhibitory and ranged from the negative magnitude of the threshold up to just below 0. This resulted in single neurons having both inhibitory and excitatory outgoing edges, rather than being explicitly just inhibitory or excitatory. Consistent with rat hippocampus [21], 20% of the edges (synapses) were inhibitory. Neurons do not normally have both excitatory and inhibitory outgoing edges due to Dale's law, so we assumed inhibitory edges contained length-

less inhibitory interneurons to simplify the model [45]. These inhibitory interneurons needed axon lengths of 0 to keep the average latency of inhibitory edges the same as the average latency of excitatory edges. Excitatory and inhibitory edges used the same ranges for their weights, as this equalization has been observed in murine pyramidal neurons [46]. To improve computational efficiency, we used an event-driven simulator instead of integrating the neurons over time [47]. Following this approach, computations were done every time a signal arrives at a neuron. This allowed for refractory periods, membrane potential summation, and membrane leak to be calculated only when necessary.

2.3.2 Network Model

We used a graph abstraction to model the structure of the network. A directed graph G consists of vertices $v \in V$ which may be connected by directed edges $e \in E$. The dynamics on the graph are analogous to that of a biological neural network, where each vertex represents a neuron and each edge represents an axon. We abstracted a number of salient features of the biological neural network, including temporal delays associated with propagating signals (action potentials) along edges (temporal latencies) and edge weights, and directly encode them in the structure of our network. Each edge is endowed with a delay d_{ij} whose subscripts i and j represent the initial and terminal vertices, respectively. The edge delay represents the duration of the traversal of an action potential along the axon of a neuron. The value of the edge delay accounts for the conduction velocity of the action potential and the physical distance the action potentials must traverse between neurons. Each edge was also included a weight w_{ij} , which represents the synaptic weight and was modified by the STDP rules and the neuron model. In Section 2.3.1, we describe in detail how values for d_{ij} and w_{ij} were selected.

Table 2.1 Each of the parameters chosen for the latencies in the network and the STDP rules was taken from murine hippocampus experimental data. The only exception is the edge weight range and threshold potential, which were chosen to make the computations manageable.

Refractory Period [17]	3.9 ms
Resting Potential	0.0
Threshold Potential	2.0
Upper Bound for Edge Weight	2.0
Lower Bound for Edge Weight	-2.0
Mean of Edge Length Gaussian [19]	730 μm
Standard Deviation of Edge Length Gaussian	219 μm
Speed of Signals Across Edges [20]	1.68 m/s
Membrane Potential Leak Time Constant [41]	0.0093
Excitatory STDP Potentiation Amplitude [41]	0.00008
Excitatory STDP Potential Time Constant [41]	7
Excitatory STDP Depression Amplitude [41]	0.00014
Excitatory STDP Depression Time Constant [41]	10
Inhibitory STDP Potentiation Amplitude [22]	0.2
Inhibitory STDP Potentiation Time Constant [22]	20
Inhibitory STDP Depression Amplitude [22]	0.2
Inhibitory STDP Depression Time Constant [22]	20

Our network consists of two distinct layers of vertices: the input layer and the recurrent layer. Importantly, as discussed in detail below, in contrast to essentially all other artificial neural network models, including other recurrent models, our model does not have an output layer. Each layer is distinguished by its connectivity, number of vertices, and initial values. Vertices of the input layer have edges connected in a feed-forward manner, such that each input vertex is connected to all recurrent layer vertices. Each vertex of the recurrent layer is connected to every other recurrent layer vertex, which results in a complete subgraph $G^{\text{hid}} \subset G$.

2.3.3 Inputs

To stimulate our network, we use images from the MNIST data set from the National Institute of Standards and Technology [48]. This data set contains images of handwritten digits between 0 and 9. Each image consists of 784 pixels in a 28 by 28 grid (matrix). We mapped each pixel to one node in the input layer of the network. The input layer is a concatenated vector equal

in size to the pixel matrix representing the MNIST digits. Every non-zero pixel causes the corresponding input node to fire when the simulation starts. Pixels with a maximal intensity of 255 were treated the same as pixels with an intensity of 1. All appropriate input nodes fire instantaneously at time 0, but their signals reach connected target neurons within the recurrent layer at different times depending on each edge's distance and conduction velocity, i.e., as a function of the latencies on the edges. The network is stimulated only once by this process. The resultant dynamics on the network are then allowed to run and evolve undisturbed.

2.3.4 Simulation Observables

We analyzed the state-space trajectory of edge weights to determine whether images from the same class resulted in similar dynamics. Initializing the network with a sample MNIST image caused a subset of the input layer vertices to activate. When a vertex was activated, it generated a signal on each of its outgoing edges. These signals either caused downstream vertices to activate or had no effect if the receiving vertex was refractory. Edge synaptic weights were modified by the STDP rules described in Section 2.3.1. The changes in synaptic weight encoded the dynamics on the network induced by the stimulus.

The vector of synaptic weights was written as $w_{ij}^k(t)$, where ij identifies the edge for the k^{th} MNIST simulation. To calculate the classification accuracy, we sampled the vector of synaptic weights every 100 milliseconds. We did not consider the synaptic weights of the edges between the input layer and the recurrent layer, but only the synaptic weights of the edges within the recurrent layer itself. This was because we focused on synaptic changes caused by activations within the recurrent layer, rather than directly from the image itself. As the input is only fed in once, there would only be one STDP change at most per edge between the input and recurrent

layer. Therefore, for n recurrent layer vertices, the size of the synaptic weight vector at each sampling time is $n(n - 1)$.

For the classification task, we compared the vectors of synaptic weights at each sampling time. The feature vectors of K MNIST simulations were written in the following matrix notation:

$$\bar{A}(t) = [w(t)^1 \ w(t)^2 \ \dots \ w(t)^K] \quad (4)$$

$\bar{A}(t)$ tracks the evolution of the feature vector over time.

2.3.5 Parallel Embedding of Inputs

Each input to the network was treated independently. When we presented the network with an image, each input layer neuron was activated only once when its corresponding pixel had a value > 0 . We did not present periodic, constant, or random rate-based inputs to the network. Upon stimulation with an image, we recorded the dynamics over the course of 2 seconds. From each simulation, we analyzed the state-space trajectory of edge weights. Since each simulation was independent of the others, we were able to simulate many networks in parallel. As such, there is no sequential training that was required. Each parallel simulation began with the same initial conditions. The number of parallel simulations is only constrained by computational resources.

Using this approach, we analyzed 10,000 MNIST images and compared vectors of synaptic weights between them. With only one stimulation of the relevant input neurons that encoded the input image, the resultant evolving dynamics within the recurrent layer were able to encode the input without any notion of training the network. Inputs are inherently captured by the very nature of the model itself. Instead of training individual networks on specific classes, we compared the resultant dynamics and synaptic weight vectors across different inputs to determine which class they are most similar to. In effect, the state space of dynamic encodings acts as a

metric space on which downstream statistical - or other machine learning - analyses on the encoded inputs can be performed.

2.3.6 KNN for Classification Accuracy

In the work we did in this paper, once we obtained the vectors of the synaptic weights from the many simulations, we ran them through a weighted k-nearest neighbors (w- KNN) algorithm to determine their localities [49]. Using w-KNN provided us with a simple way to assess the synaptic weight vectors based on their Manhattan distances to discern if they are of the same class or not. Using this w-KNN, we classified each vector in the test set based on their distances from the 9,000 vectors in the training set. The w-KNN weighting of each neighbor comes from the inverse of the distance between it and the weight vector being evaluated; whichever MNIST digit has the highest total weight based on the 5-nearest neighbors is the one that is chosen for the classification. The 10,000 total vectors were sorted into the training and test sets randomly and used for classification over 10 iterations to avoid any sampling bias. In addition, because there is approximately the same amount of simulations per digit in our population of 10,000 vectors, imbalances should not create biases in modes when many neighbors are considered. Importantly, even though the neighbors are pulled from a “training set”, there is no actual training in this system, as all of the simulations were performed independently.

2.3.7 UMAP for Visualization

In order to visualize the clustering of these synaptic weight vectors, we used the UMAP dimension reduction algorithm [50]. We took edge weight vectors from 10,000 simulations and ran them through UMAP to see how separated each digit is from the others. UMAP employs k-nearest neighbors concepts with a unique local metric space implementation to effectively reduce

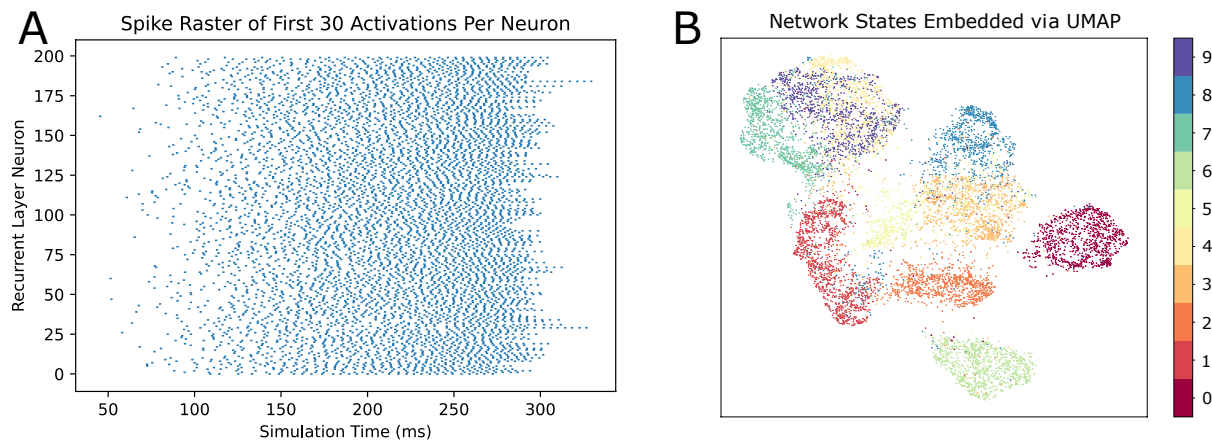


Figure 2.3 Qualitative characteristics of the model. (A) Spike raster for the first 30 activations for each neuron in the network. The activity of the network starts slowly due to the delays from the input nodes, but quickly increases. In addition, even with just 200 recurrent nodes, there are approximately as many as 20,000 activations per second in the simulation. As such, the activation dynamics are constantly changing the weights of the system. (B) Computed UMAP projection to reduce the dimensionality of the edge weight vectors and the resultant separation of classes. Edge weight vectors from stimuli corresponding to the same digit generally cluster together in the projection and are visually separable. This implies that the resultant network dynamics to within class inputs are more similar to each other than to vectors from other digits.

the dimensionality of large data. By using this, we obtained a two-dimensional plot of all of the vectors labeled by their input digit. This revealed how similar or distinct the state-spaces of vectors of each digit can be based on their proximity to each other.

2.4 Results

2.4.1 Characterization of Network Activity and Embedding Quality

We first explored and characterized the dynamic activity of network. We recorded the time of the first 30 activations for each recurrent layer neuron in a 200-neuron network. The activity is not identical across all neurons, but most neurons had fired 30 times by 300 ms post-stimulation (Figure 2.3A). Maximal information embedding was achieved at 300 ms post-stimulation, which corresponded with peak classification accuracy (Figure 2.4). This implies that the network is not just embedding a few key activations that reflect the stimulus directly, but

rather aggregating thousands of activations into the dynamical structure of the network, with each activation producing hundreds of synaptic weight changes via the STDP rules.

To better understand the behavior of networks stimulated by different images of the same class, we took edge weight vectors at different simulation times and ran them through a UMAP dimension reduction. The greatest separation occurred at 300 ms post-simulation (Figure 2.3B). Notably, even though the networks were stimulated independently, they tended to cluster together based on the class of the image used. To generate Figure 2.3, only the edge weight vectors were used, with no additional information about spiking activity or stimulus patterns. In other words, within the structural embedding itself, there was information about the stimulus. Further, we did not include edge weights between the input and recurrent layer, since these would not reflect the activity of the recurrent layer. Those signals would only travel along edges due to image pixels, so those edges would only be affected by STDP once. Since the goal was to study how well network dynamics can embed information through thousands of signals modulated by STDP, the input layer to recurrent layer edges were not included.




Figure 2.4 Summary of KNN classification accuracies over time. We explored how different parameter values impacted performance. (A) We tested various recurrent layer sizes to balance the classification accuracy with simulation efficiency. The peak classification accuracy improved as the size of the network grew, reaching a constant at 300 ms of simulation time. Observed increases in classification accuracy performance in networks larger than 200 nodes were not significant at peak classification accuracy, but rather affected decreases in accuracies post-peak over time. (B) We tested different ratios of excitatory to inhibitory edges. Some percentages of inhibitory edges, such as 70% and higher, were so negative that they silenced the network prematurely. When comparing the peak classification accuracies across all percentages, we observed that 20% resulted in the highest peak accuracy. Interestingly, some inhibitory percentages have peak accuracies at simulation times later than 300 ms, which is a phenomenon that was not observed with many other parameter configurations.

We then used a k-nearest neighbors algorithm to classify networks with unknown stimuli based on their edge weight neighbors. Because the information capacity of the edge weight vector is bound to its dimension, we decided to experiment with the number of neurons (and therefore number of edges) within the recurrent layer. As the network size grew and the number of edges increased, the peak classification accuracy determined using a k-nearest neighbors algorithm also increased, but with diminishing returns (Figure 2.4). With 200 neurons in the recurrent layer, the peak classification accuracy was $96.5 \pm 0.5\%$. With networks larger than that, peak classification marginally increased while substantially enlarging the edge weight vector space. With n recurrent layer neurons, an edge weight vector would contain $n(n - 1)$ edges. Though a recurrent layer of 300 neurons has slightly higher accuracy, the computational costs are disproportionately larger, effectively reducing the return on investment of increased network size.

For all network sizes, the peak classification accuracy was reached at 300 ms post-stimulation, which is an interesting phenomenon independent of the number of neurons. Interestingly, as the network size increases, the embedding of the stimulus seems to be retained longer within the network. With recurrent layers as large as 300 or 400 neurons (89,700 or 159,600 edges), a decrease in classification accuracy was not measurable within the 2-second time window we were using.

Lastly, we explored the effects of changing different ratios of excitatory to inhibitory edges. For a network with a 200-neuron recurrent layer, the peak classification accuracy of $96.5 \pm 0.5\%$ was achieved with an inhibitory edge fraction of 10-30% (Figure 2.4). This ratio matches biological estimates of excitatory to inhibitory ratios quite accurately [21, 51–53].

Interestingly, with 50% of the edges being inhibitory, the network neither becomes quiescent nor loses the embedding of the stimulus over time. While this fraction did not achieve the same peak classification accuracy as lower inhibitory fractions, this maintained embedding could be useful in some biological systems.

2.4.2 Extending Simulations Beyond Experimental Plausibility

We then evaluated the relative independent contribution of potentiation and depression to the network embeddings by running simulations with each separately. Interestingly, turning off potentiation had a minimal effect on the peak classification accuracy (Figure 2.5). With just depression gradually decreasing the edge weights, the network could still embed the stimulus with the same classification accuracy. But when depression was removed, potentiation alone was incapable of maintaining the embedding quality of the stimulus. With depression and with or without potentiation, a classification accuracy of $96.36 \pm 0.47\%$ was achieved with a recurrent layer of 200 neurons. In contrast, without depression, the classification accuracy fell to $87.24 \pm 0.78\%$ for the same sized network.




Figure 2.5 Testing the individual contributions of potentiation and depression on edge weight changes and classification accuracies. Depression alone was sufficient to maintain high classification accuracy, but potentiation by itself reduced the peak accuracy substantially. We then switched the parameter values for each learning rule to see if they were the cause, but we observed the same phenomena regardless of which parameter values were used for each learning rule. Edge weight depression seems to be necessary for stimulus embedding in this system.

To disqualify the effect of our chosen parameter values themselves (which were derived experimentally), we switched the parameters used for depression and potentiation. The original parameters produced depression with higher amplitude than potentiation. But even after this was switched, depression still had a larger impact on classification accuracy (Figure 2.5). In any paradigm without depression, the accuracy would drop from 96% to 88%. On top of that, paradigms without potentiation maintained 96% accuracy as long as depression was active.




Figure 2.6 Increased and decreased refractory periods of neurons within the recurrent layer suggests a quantifiable reason for the scale of biological refractory periods. Regardless of how large or small the refractory period is, with both potentiation and depression changing weights, the classification accuracy remains consistent. However, with just depression affecting the weights, longer refractory periods result in better classification accuracy. Inversely, with just potentiation, shorter refractory periods lead to better classification accuracy. There seems to be an inflection point at around 1 ms where depression alone begins to embed information better than potentiation alone.

To further investigate the balance between potentiation and depression, we next changed the refractory period of the neurons in the recurrent layer. As the frequency of potentiation to depression events depends on the length of the refractory periods, it follows that this would have a large impact on which process has more influence. For all the refractory periods we tested, from 100 ms to 10 ms, having both potentiation and depression changing edge weights maintained maximum classification accuracy in our system as seen in Figure 2.6. It was only

when we investigated potentiation or depression alone that differences became noticeable. With shorter refractory periods, just potentiating changes performed better than just depressive changes. However, once refractory periods grew longer and more depressive events took place, depression alone performed better than potentiation alone. Across all tested refractory periods, having just one or the other was sufficient to reach the maximum classification accuracy obtained by having them both together. The inflection point where depression becomes more impactful than potentiation seemed to be with a refractory period of around 1 ms.

2.4.3 Comparing Edge Weight Vectors to Image Stimuli

Given the high classification accuracy of the embeddings using KNN across simulations, we investigated more closely the raw images themselves. We plotted the 9,000 nearest neighbors for one particular MNIST image based on pixel value distance (Figure 2.7). Then, over the course of a simulation, we plotted the 9,000 nearest neighbors of that same image based on the edge weight vector distance. Taking the running mean of each group of 100 edge weight vector distances, we observed similar patterns start to form and disappear over the course of a typical simulation. Notably, the edge weight vector distances seem to correlate with the pixel value distances between 200 and 400 ms of simulation, when peak classification accuracy occurs. After 400 ms, the edge weight vector distances lose correlation with the pixel value distances.

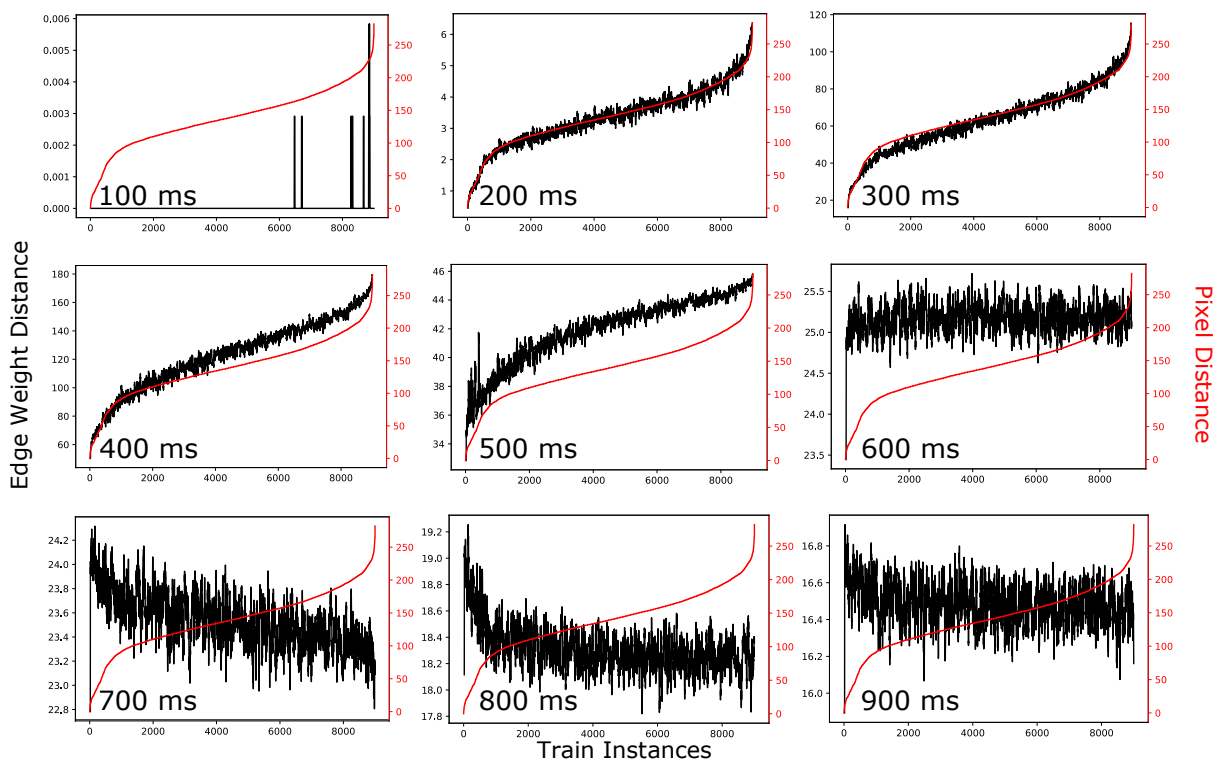


Figure 2.7 Distances of nearest neighbors to a single image based on their raw pixel values sorted to be ascending (red lines). Over the course of the simulation, we used this ordering to see how it compared to distances of the same image based on the network edge weights while preserving the order from the raw pixel values. The running mean of the following nearest 100 image edge weight distances is used to produce the black traces. We can see that around 300 ms, where the embedding is at its best, the edge weight distances seem to correlate nicely with the red raw pixel distances. However, this correlation is transient and is not reached before 200 ms of simulation time nor preserved after 400 ms.

2.4.4 Comparison Against a Standard Neural Network

To assess how our approach compared to a conventional neural network, we constructed a three-layer sequential model in TensorFlow. The first layer contained 28 by 28 nodes, equivalent to the 784 nodes in our model. A second hidden layer contained a varying number of nodes to serve as an analogue for our recurrent layer. Lastly, a third output layer contained 10 nodes. One-shot learning was implemented on the ANN by training it until the model over-fitted the training data, which consisted of one randomly chosen image from each class. For the ANN,

we used the Adam optimizer, the Cross Entropy loss function, and a ReLu node activation model.

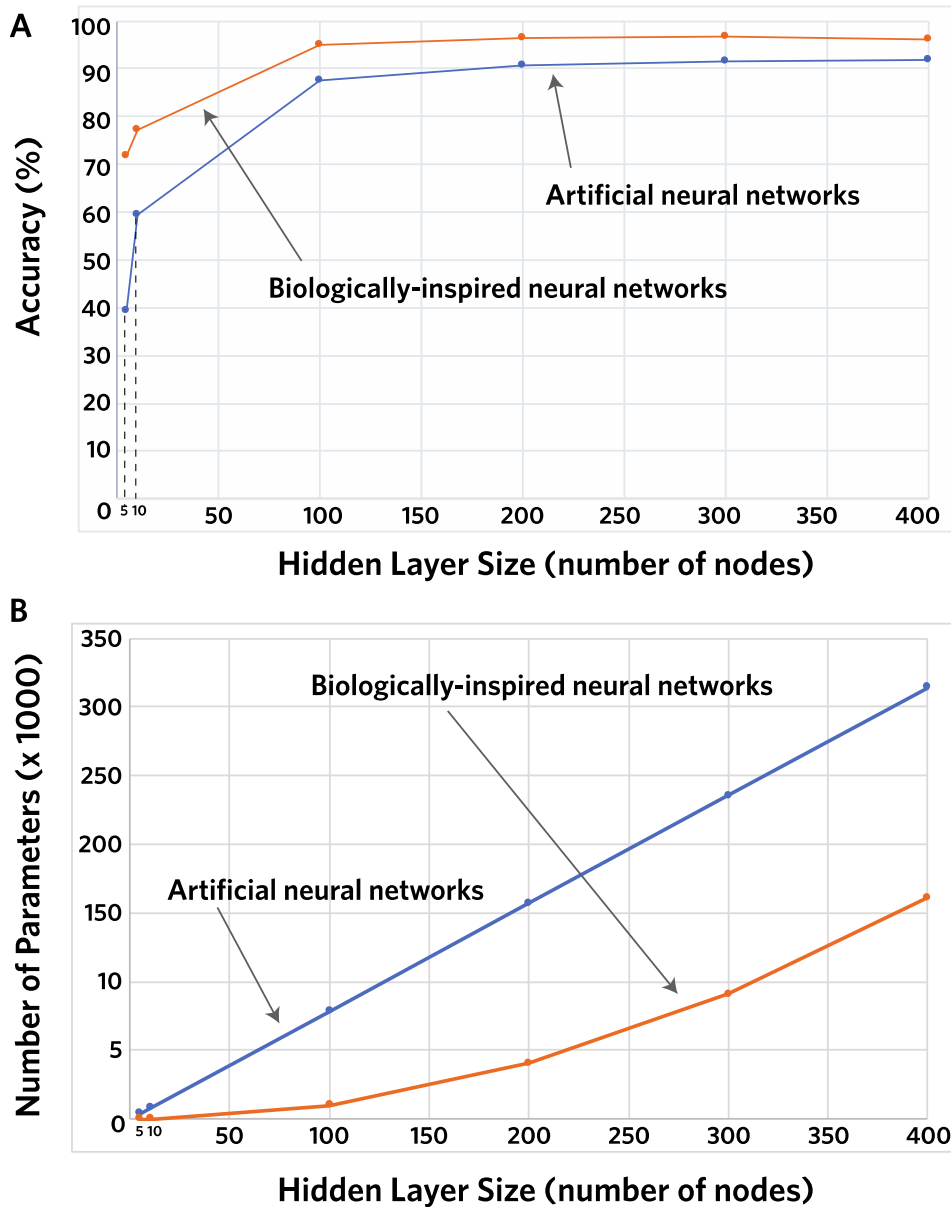


Figure 2.8 (A) We compared the classification accuracy performance on an MNIST data set between our biologically-inspired neural network and an artificial neural network using a similar edge weight KNN method. We see that our network achieves better classification accuracy regardless of the number of neurons within the recurrent/hidden layer. (B) When looking at the number of parameters adjusted in networks with different recurrent/hidden layer sizes, we see that our network scales up with fewer parameters than an artificial neural network. While the number of parameters begins similarly, the artificial neural network parameters balloon as more hidden nodes are added.

After training the ANN, an embedding space of 9,000 was generated using 9,000 images from the test set. Each dimension of the embedding space corresponded to the activation value of a hidden layer node. With this embedding space, different sets of 1,000 test images were assessed for their classification accuracies. For each set, we used a cosine distance-based w-KNN. For both the ANN and our biologically-inspired neural network, each KNN was run 10 times with the training and testing sets reshuffled each time to avoid any unwanted sampling bias. Given these parameters, the ANN consistently performed worse on one-shot classification accuracy regardless of the recurrent/hidden layer size tested (Figure 2.8). In addition, it consistently required more parameters than our model for all tested hidden layer sizes.

2.5 Discussion

2.5.1 Neurobiological Relevance of Our Results

The embeddings we observe in this work are represented as compositions of large combinations of weight changes. Rather than a few highly specific activations triggered by a stimulus, thousands of activations are integrated to produce the network state that ultimately encodes the input data. Without any pre-training, the base network stimulated independently with two images of the same class will likely end up with two similar network states 300 ms into the simulation. Further, images with various orientations across the same class share similar enough patterns to trigger similar weight changes overall. Figure 2.3B supports this interpretation, as states of networks of the same classes aggregate at particular times when simulated independently of each other. It is possible that some networks in the brain may work similarly, particularly in the perirhinal cortex, where object recognition is accomplished by networks that have strong depression, as our networks do [23]. In the brain, similar stimuli may result in similar structural states for object identification.

Broadly speaking, our results match mechanistic neurobiological expectations, and certain predicted implications of such mechanisms and processes. For example, other work suggests that neurobiological networks need to be sufficiently large in order to embed a stimulus [54]. And yet other another study suggested that inhibition necessarily plays an essential role in information embedding [14]. In our work here, we extend these results by showing the effectiveness of leveraging neurophysiological STDP rules to encode structural embeddings in networks. This is relevant to proposed ideas about structural changes being a part of the process of embedding memories [34]. In fact, at least within the context of our simplified model we show that such structural embeddings are sufficient to encode a memory-like events at a synaptic scale.

Our simulations also allow exploring the trade-offs between stimulus embedding and network size. As we increase the size of our networks beyond 200 recurrent neurons, the classification accuracy improves only marginally while increasing the computation time dramatically. This could be a relevant consideration in the context of physical resources in the brain, such as neurotransmitters, neuromodulators and metabolic ATP considerations. It could be inefficient to rely on the activation of larger networks to complete a task because the embedding capacity may improve only marginally at a significant resource cost.

Other studies have seemingly observed this phenomenon in working memory; training increases the number of neurons initially stimulated, but it does not increase the number of neurons recruited to the assembly in general [55]. While raising the number of neurons that are initially stimulated increases ATP consumption, this increase is less than that by enlarging the assembly itself. Therefore, there is a more efficient trade-off between increasing resource usage and computational capacity in these neuron assemblies.

The importance of appropriate ratios of excitation to inhibition was also reflected in our model and simulations. While a number of experimental studies have identified the percentages of inhibitory synapses in the brain [21, 51–53], our simulations allowed us to test the extreme ranges of this ratio. We measured the best classification accuracy with inhibition percentages between 10-30%, which matches what has been measured experimentally. These results support the hypothesis that the brain may have selectively evolved networks these inhibitory percentages in order to optimally embed information.

However, one exception that is not easily explainable is why 50% inhibition produces a classification accuracy that did not decrease significantly after hundreds of milliseconds of simulation. One possible function for a network with this inhibitory percentage is the acquisition and maintenance of fear memory through inhibitory neurons in the intercalated cell masses (ITCs) of the amygdala [56]. These masses are densely inhibitory and connect the amygdala to surrounding structures. In our model, all inhibitory neurons follow the same activation and plasticity rules rather than considering variability. But *in vivo* experiments have shown that different populations of inhibitory interneurons work independently of each other [57, 58]. Perhaps the phenomena we observed with 50% inhibition resulted from an oversimplification of the different populations (and associated dynamics) in biological inhibition.

2.5.2 Relative Contributions of Potentiation and Depression to Information Embedding

Our results suggest that within the range of neurobiological refractory periods, depression alone embeds a stimulus better than potentiation alone. Experimentally, it has been previously shown in a mouse model that potentiation by itself is less effective at embedding working memory than when both potentiation and depression are active together [59]. Based on the different refractory periods we tested, we saw that potentiation and depression independently

each consistently embed information at different efficiencies. With shorter refractory periods, potentiation works better alone than depression while with longer refractory periods, depression works better alone than potentiation. This observation is intuitive, as fewer signals arrive at refractory post-synaptic neurons with a shorter refractory period and, therefore, those signals could contribute to future activations.

Another consideration for an optimal refractory period is the relative strengths of potentiation and depression. Experimental data has shown depression to have a larger magnitude of change than potentiation [5, 13, 60]. However, this may just be a consequence of the refractory period, as our data shows that depression has stronger embedding capabilities than potentiation with a biological refractory period regardless of plasticity parameter values. As such, the magnitudes of depressive changes may seem larger experimentally because the embedding capability of depression is better at that refractory period. Quantitative experimental observations may be influenced by the efficacy of depression at embedding information with a biological refractory period. In addition, since our model shows the importance of depressive events for embedding, we speculate that the biological magnitude is larger for depression than potentiation to increase the fidelity and range of the depressive weight changes. This is a could be beneficial because it would allow for embedding through unique, subtle depressive changes to weights.

An analysis of the synaptic weight space of our model and networks agrees well with studies that have looked at the state spaces of memory [61, 62]. It has been proposed that networks in the brain may alternate between persistent states to store memories. These states could possibly consist of synaptic weights, synaptic connections, membrane voltages, neuronal activation patterns, or other biochemical variables. In our model, we directly used synaptic

weights to define the different states that networks transition between as they are stimulated. In particular, we show that networks stimulated by images of the same class reach similar states after 300 ms of simulation.

2.5.3 A Potential Computational Role for Astrocytes in Encoding and Storing Memories

One natural question that arises from the synaptic weight state space is its potential capacity. We show that there is a useful amount of separation with ten classes, but further work is needed to determine if there is sufficient separation resolution with a greater number of classes. In addition, these synaptic weight states do not seem to persist over time, so the dynamics that produce such dynamic transience could also be further studied. One intriguing idea would be to incorporate a glial component to the network, which could stabilize states to persist for longer. Experimental studies have shown that glial cells play a key (though not yet well understood) role in hippocampal structural maintenance [63, 64]. Incorporating glial cell support could allow the weight changes in our network to preserve information for longer and delay our observed loss in accuracy.

In addition to homeostatic support functions, glial cells, and in particular astrocytes, may function as downstream readers of synaptic weights in the hippocampus and other parts of the brain. It is possible these cells could play a role reminiscent of our KNN algorithm. In particular, in the hippocampus a single astrocyte can extend processes to numerous synapses that, from a cellular mechanistic perspective at least, can monitor and actively participate in neuronal synaptic signaling, thus potentially contributing to the activity and plasticity that takes place [65]. While the astrocytes may not keep track of the absolute synaptic strengths themselves, their internal calcium gradients change with potentiation and depression of the synapses [66, 67]. As such, their internal states integrate the overall plasticity changes over thousands of synapses to

monitor the population as a whole [68]. As astrocytes can communicate with each other through intercellular calcium waves and gap junctions, they can transfer information across the astrocyte syncytium to other astrocytes [69, 70]. In effect, this system has the potential to 'short circuit' the neuronal network, thereby putatively affecting plasticity [71].

Taken all together, astrocytes, could be reading global synaptic plasticity changes in a network caused by a stimulus, possibly allowing them to embed the stimulus and transfer it to another downstream network to facilitate the flow of information. Our KNN algorithm is similar in the sense that it is sensitive to the edge weight changes within the networks rather than their absolute weights, which are largely influenced by the initial state of the network. The synaptic plasticity information astrocytes integrate could be the biological analogue of our edge weight vectors and could serve as a plausible mechanism through which information is transferred in the hippocampus and possibly other parts of the brain.

Furthermore, astrocytes as readers of synaptic weights may explain the cause of transient epileptic amnesia [72]. When considering the paradigm of rate coding, seizures represent a unique firing rate for various regions of the brain. Most patients suffer amnesia during an episode and often during the postictal state [73]. This would imply that there is more to memory formation than firing rates alone, as epileptic firing rates would not be confounded by another event. This is where our hypothesis fits in. The chaotic network activity of seizures may cause STDP changes that do not properly embed any particular stimulus. As such, when the brain tries to recall memories from the event, it may not be able to accurately process what the synaptic changes represent. Morphological changes have been observed in the synapses of epileptic brain regions [74], so there is a precedent that seizures abnormally change synaptic plasticity.

In addition, there may be dysfunction at the level of astrocytes as well. As astrocytes are critically important for removing excess neurotransmitters from synapses as part of their homeostatic functions, seizures due to insufficient removal may be caused by astrocyte dysfunction [75]. Any condition that results in a disruption of astrocytic homeostasis at the synapse, could conceivably also affect how astrocytes read synapses. While there have been hypotheses for how amnesia can result from rate coding, such as misaligned gamma and theta wave synchronization [74], a consideration of the critical roles astrocytes could be playing in storing long term memories reflects a potential significant computational function for these cells of which we still know very little about.

2.5.4 Future Neurobiological Directions and Implications

In the future, the introduction of noise into the system could be studied as well, as experimental studies have shown STDP to struggle when noise via jitter is added to the spike timing [76]. Inversely, adding minimal noise may be beneficial to allow subthreshold signals to impact network dynamics [77]. Further, our threshold potential was chosen such that an average of three excitatory signals would be required to reach it. Follow-up experiments could adjust the threshold potential to observe how increasing and decreasing the number of activations affects the STDP embedding of the stimulus. Lastly, another future direction could involve studying the effect of multiple sequential stimuli on the embedding. If the synaptic weights are capable of reflecting one stimulus, what happens when two are presented? Do the weights reflect a composition of the inputs?

When looking at neurons within the brain, our synaptic weight spaces offer valuable insight. While rate coding, postsynaptic potential magnitudes, and myelin thicknesses are primarily considered as evidence of learning [7–9], structural changes, in particular, can be

meaningful as well. Our methodology demonstrates that the stimulus can be identified by interrogating the structural state of a network following stimulation. In addition to astrocytes potentially monitoring synapses, the structural state may also strongly influence the overall rate and magnitudes of activations within the network. As such, the structural state may correlate well with current observations of activation rates and postsynaptic potential magnitudes. As brain slice imaging techniques continue to improve, subtle structural differences associated with different learned patterns may also become observable experimentally.

When designing our network, we used a single recurrent layer to avoid any brain region-specific architecture. To generalize functions and interactions of various brain regions, we minimized topological variability by connecting each input neuron to every recurrent layer neuron. However, various regions of the hippocampus and parahippocampal cortex are recurrent networks, so our model could represent an apt analogue for them [78, 79]. In addition, we maintained the same initial conditions for simulations run in parallel (axon lengths, starting weights) to minimize variability. To properly evaluate our edge weight and delay distributions, we tested 11 randomly chosen initial conditions and observed a classification accuracy of $96.53 \pm 0.06\%$ with a 200-neuron recurrent layer. With these controls in place, the experimental data we gathered resulted from changes to the refractory periods, network sizes, inhibitory percentages, and STDP rules. Regarding this design philosophy, further work can be done to investigate how different network architectures are impacted by varied refractory periods and potentiation/depression isolation. Furthermore, allowing for variability in initialization or even sequential stimulation could lead to new insights into the homeostatic capabilities of the networks to recalibrate and resume learning. To add more neuroscientific features to the network, we could also add dynamic dendritic spine calcium concentrations, as a recent model

observed their effect on STDP strength [80]. In particular, in high-activity networks such as ours, the magnitude of STDP changes could be quite variable depending on the calcium transients.

Experimental techniques capable of observing the graded breakdown of synaptic embeddings could produce diagnostic tools for neural dysfunctions and new understandings of computation in the brain. Research is currently being done to quantify how these synaptic embeddings relate to information storage [81], but it is still unclear how to best monitor them. In our model, embeddings are intimately tied to the timing of signals between neurons. Variability of a network's geometry and signaling parameters affect timing, and volatility in either can cause failure in embedding robustness. Changes in a network's geometry arise from available neural circuits due to disruptions to oscillatory neural patterns, and neuromodulatory signals [82–84]. Dysfunction affecting the conduction velocity of axons, whether due to myelination or ion channel irregularities, could negatively impact memory storage due to inconsistent signaling parameters affecting the signal timings. Further, if there is limited neurotransmitter availability, either due to poor astrocyte recycling or insufficient production, the resulting variability in spike amplitudes will affect signal timing, thus the embedding, and may lead to memory problems. Another potential source of issues with memory could be connectivity changes due to synaptic pruning or neurotropy, which could change the paths that signals travel along in response to the same stimulus. Lastly, conditions constraining the sizes of neuron assemblies could result in poor memory due to an insufficient number of edges on which information can be embedded. Biological experiments focused on any of these issues could elucidate mechanisms of memory loss beyond just synapse destruction due to neuroinflammation.

2.5.5 Edge Weights As A Machine Learning Methodology

After 300 ms of simulation time of our model in response to a specific input vector, the nearest neighbors to a specific vector closely resemble the nearest neighbors of the input image itself. This observation highlights the functional potential of our model. With STDP, neuron (node) refraction, inhibitory edges, and edge latencies, the network is able to encode the stimulus after 300 ms. These results suggest that adding neurobiological features to canonical artificial neural networks could enhance their function for applications that, in particular, necessitate encoding or embedding sparse or incomplete data that is not sufficient to train a traditional ANN. By allowing the networks to dynamically evolve in response to the unique stimulation pattern encoded in the input vector, our model is able to take advantage of the resultant structural weight states to classify the input in an unsupervised way with no prior training or exposure to a priori data. This reduces the number of parameters when compared to a deeper neural network [85]. More broadly, our work fits into a body of work that is exploring how neurophysiological and biological principles can advance machine learning from the current state of the art [86–90].

Since each simulation is run independently of the others, there is no “training phase” in our system. Each individual network embeds information in parallel via one-shot learning [85]. The network activity of our model generates a dynamic embedding of the stimulus through the synaptic weight changes, which can be subsequently used for classification. Our dynamic embedding can be considered as a functional transformation of the stimulus. When we ran the KNN on the pixels of the input images we attained a classification accuracy of 94.5%, whereas, the KNN on our dynamic embedding resulted in a classification accuracy of 96.5%. Given that we empirically observed an increase in classification accuracy as a result of the dynamic embedding, we postulate that the resultant transformations are non-trivial.

Further work should explore and develop the machine learning potential of this system as well as other biologically-inspired systems. Although we only tested our network on the MNIST data set, which is no longer a hurdle for machine learning systems, we plan to test more challenging data sets next. While the absolute classification accuracy of $96.5\pm 0.5\%$ is not as high as state-of-the-art neural networks achieve, it was achieved with one-shot learning. Critically, the intent of our model is not to compete with classical machine learning, but rather, to develop a neurobiologically derived machine learning paradigm that is capable of operating under conditions that still challenge even the most advanced classical system.

2.6 Acknowledgements

We thank the Lawrence Livermore National Laboratory for providing us with extensive computational resources for our experiments. And to our colleagues at Microsoft Research for many fruitful discussions and earlier related work.

Chapter 2, in full, has been submitted for publication of the material as it may appear in Nature Scientific Reports, 2023, George, Vivek K.; Morar, Vikash; Silva, Gabriel A., Nature Portfolio, 2023. The dissertation author was one of two primary researchers and authors of this paper.

2.7 References

- [1] Rusell, S. & Norvig, P. Artificial Intelligence: A Modern Approach, 4th edition (Pearson, London, 2000).
- [2] Murphy, K. Machine Learning: A Probabilistic Perspective (MIT Press, Cambridge, MA, 2012).
- [3] Nishihara, R., Moritz, P., Wang, S., Tumanov, A., Paul, W., Schleier-Smith, J., Liaw, R., Niknami, M., Jordan, M. & Stoica, I. Real-time machine learning: The missing pieces. arXiv 1703.03924 (2017).
- [4] Mintz, Y. & Brodie, R. Introduction to artificial intelligence in medicine. Minimally Invasive Therapy and Allied Technologies doi: 10.1080/13645706.2019.1575882 (2019).

- [5] Abbott, L. & Nelson, S. B. Synaptic plasticity: taming the beast. *Nature Neuroscience* 3, 1178–1183 (2000).
- [6] Humeau, Y. & Choquet, D. The next generation of approaches to investigate the link between synaptic plasticity and learning. *Nature Neuroscience* 22, 1536–1543 (2019).
- [7] Skilling, Q. M., Eniwaye, B., Clawson B. C., Shaver, J., Ognjanovski, N., Aton, S. J. & Zochowski, M. Acetylcholine-gated current translates wake neuronal firing rate information into a spike timing-based code in non-rem sleep, stabilizing neural network dynamics during memory consolidation. *PLoS Computational Biology* 17, 9 (2021).
- [8] Froemke, R. C., Debanne, D. & Bi, G.-Q. Temporal modulation of spike-timing-dependent plasticity. *Frontiers in Synaptic Neuroscience* 2, 19 (2010).
- [9] Sancho, L., Contreras, M. & Allen, N. J. Glia as sculptors of synaptic plasticity. *Neuroscience Research* 167, 17–29 (2021).
- [10] Gibson, E. M., Purger, D., Mount, C. W., Goldstein, A. K., Lin, G. L., Wood, L. S., Inema, I., Miller, S. E., Bieri, G., Zuchero, J. B., Barres, B. A., Woo, P. J., Vogel, H. & Monje, M. Neuronal activity promotes oligodendrogenesis and adaptive myelination in the mammalian brain. *Science* 344 (2014).
- [11] Bechler, M. E., Swire, M. & Ffrench-Constant, C. Intrinsic and adaptive myelination—a sequential mechanism for smart wiring in the brain. *Developmental Neurobiology* 78, 68–79 (2018).
- [12] Bliss, T. V. & Cooke, S. F. Long-term potentiation and long-term depression: a clinical perspective. *CLINICS* 66, 3–17 (2011).
- [13] Feldman, D. E. The spike-timing dependence of plasticity. *Neuron* 4, 556–571 (2012).
- [14] Vogels, T., Sprekeler, H., Zenke, F., Clopath, C. & W, G. Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* 334, 1569–1573 (2011).
- [15] Vogels, T. P., Froemke, R. C., Doyon, N., Gilson, M., Haas, J. S., Liu, R., Maffei, A., Miller, P., Wierenga, C. J., Woodin, M. A., Zenke, F. & Sprekeler, H. Inhibitory synaptic plasticity: spike timing-dependence and putative network function. *Frontiers in Neural Circuits* 7, 119 (2013).
- [16] Abbott, L. Lapique’s introduction of the integrate-and-fire model neuron (1907). *Brain Research Bulletin* 50, 303–304 (1999).
- [17] Raastad, M. & Shepherd, G. M. Single-axon action potentials in the rat hippocampal cortex. *The Journal of Physiology* 548, 745–752 (2003).

- [18] Silva, G. A. The effect of signaling latencies and node refractory states on the dynamics of networks. *Neural Computation* 31, 2492–2522 (2019).
- [19] Spruston, N. Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience* 9, 206–221 (2008).
- [20] Henze, D. A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K. D. & Buzsáki, G. Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *Journal of Neurophysiology* 84, 390–400 (2000).
- [21] Gulyás, A., Megias, M., Emri, Z. & Freund, T. Total number and ratio of excitatory and inhibitory synapses converged onto single interneurons of different types in the ca1 area of the rat hippocampus. *Journal of Neuroscience* 19, 10082–10097 (1999).
- [22] Kleberg, F. I., Fukai, T. & Gilson, M. Excitatory and inhibitory stdp jointly tune feedforward neural circuits to selectively propagate correlated spiking activity. *Frontiers in Computational Neuroscience* 8, 53 (2014).
- [23] Collingridge, G. L., Peineau, S., Howland, J. G. & Wang, Y. T. Long-term depression in the cns. *Nature Reviews Neuroscience* 11, 459–473 (2010).
- [24] Griffiths, S., Scott, H., Glover, C., Bienemann, A., Ghorbel, M. T., Uney, J., Brown, M. W., Warburton, E. C. & Bashir, Z. I. Expression of long-term depression underlies visual recognition memory. *Neuron* 58, 159–161 (2008).
- [25] Ge, Y., Dong, Z., Bagot, R. C., Howland, J. G., Philips, A. G., Wong, T. P. & Wang, Y. T. Hippocampal long-term depression is required for the consolidations of spatial memory. *PNAS* 107, 16697–16702 (2010).
- [26] Ito, M. Cerebellar long-term depression: Characterization, signal transduction, and functional roles. *Physiological Reviews* 81, 1143–1195 (2001).
- [27] Braunewell, K.-H. & Manahan-Vaughan, D. Long-term depression: A cellular basis for learning? *Reviews in the Neurosciences* 12, 121–140 (2001).
- [28] Piochon, C., Kano, M. & Hansel, C. Ltp-like molecular pathways in developmental synaptic pruning. *Nature Neuroscience* 19, 1299–1310 (2016).
- [29] Rolls, E. T. The storage and recall of memories in the hippocampo-cortical system. *Cell and Tissue Research* 373, 577–604 (2018).
- [30] Richards, B. A. & Frankland, P. W. The persistence and transience of memory. *Neuron* 94, 1071–1084 (2017).

- [31] Zeng, M., Chen, X., Guan, D., Xu, J., Wu, H., Tong, P. & Zhang, M. Reconstituted postsynaptic density as a molecular platform for understanding synapse formation and plasticity. *Cell* 174, 1172–1187 (2018).
- [32] Zhu, J., Shang, Y. & Zhang, M. Mechanistic basis of maguk-organized complexes in synaptic development and signalling. *Nature Reviews Neuroscience* 17, 209–223 (2016).
- [33] Zhang, Y., Cudmore, R. H., Lin, D.-T., Linden, D. J. & Huganir, R. L. Visualization of nmda receptor–dependent ampa receptor synaptic plasticity in vivo. *Nature neuroscience* 18, 402–407 (2015).
- [34] Yang, G., Pan, F. & Gan, W. Stably maintained dendritic spines are associated with lifelong memories. *Nature* 462, 920–924 (2009).
- [35] Xu, T., Yu, X. & Perlik, A. Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature* 462, 915–919 (2009).
- [36] Hayashi-Takagi, A., Yagishita, S. & Nakamura, M. Labelling and optical erasure of synaptic memory traces in the motor cortex. *Nature* 525, 333–338 (2015).
- [37] Roberts, T., Tsuchida, K. & Klein, M. Rapid spine stabilization and synaptic enhancement at the onset of behavioural learning. *Nature* 463, 948–952 (2009).
- [38] Sanders, J., Cowansage, K., Baumgartel, K. & Mayford, M. Elimination of dendritic spines with long-term memory is specific to active circuits. *Journal of Neuroscience* 32, 12570–12578 (2012).
- [39] Buibas, M. & Silva, G. A framework for simulating and estimating the state and functional topology of complex dynamic geometric networks. *Neural Computation* 23, 183–214 (2011).
- [40] Bi, G. & Poo, M. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of Neuroscience* 18, 10464–10472 (1998).
- [41] Clopath, C., Büsing, L., Vasilaki, E. & Gerstner, W. Connectivity reflects coding: a model of voltage-based stdp with homeostasis. *Nature Neuroscience* 13, 344–352 (2010).
- [42] van Rossum, M., Bi, G. & GG, T. Stable hebbian learning from spike timing-dependent plasticity. *Journal of Neuroscience* 20, 8812–8821 (2000).
- [43] Hennequin, G., Agnes, E. J. & Vogels, T. P. Inhibitory plasticity: Balance, control, and codependence. *Annual Review of Neuroscience* 40, 557–579 (2017).
- [44] Izhikevich, E. M., Gally, J. A. & Edelman, G. M. Spike-timing dynamics of neuronal groups. *Cerebral Cortex* 14, 933–944 (2004).

- [45] Dale, H. Pharmacology and nerve-endings. *Proceedings of the Royal Society of Medicine* 28, 319–332 (1935).
- [46] Xue, M., Atallah, B. V. & Scanziani, M. Equalizing excitation-inhibition ratios across visual cortical neurons. *Nature* 511, 596–600 (2014).
- [47] Carothers, C. D., Perumalla, K. S. & Fujimoto, R. M. Efficient optimistic parallel simulations using reverse computation. *ACM Transactions on Modeling and Computer Simulation* 9, 224–253 (1999).
- [48] Deng, L. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine* 29, 141–142 (2012).
- [49] Goldberger, J., Roweis, S., Hinton, G. & Salakhutdinov, R. Neighbourhood components analysis. *Advances in Neural Information Processing Systems* 17, 513–520 (2005).
- [50] McInnes, L. & Healy, J. Umap: Uniform manifold approximation and projection for dimension reduction. *ArXiv* (2018).
- [51] Vereczki, V. K., Müller, K., Krizsán, É., Máté, Z., Fekete, Z., Rovira-Esteban, L., Veres, J. M., Erdélyi, F. & Hájos, N. Total number and ratio of gabaergic neuron types in the mouse lateral and basal amygdala. *Journal of Neuroscience* 41, 4575–4595 (2021).
- [52] Tremblay, R., Lee, S. & Rudy, B. Gabaergic interneurons in the neocortex: From cellular properties to circuits. *Neuron* 91, 260–292 (2016).
- [53] Pelkey, K. A., Chittajallu, R., Craig, M. T., Tricoire, L., Wester, J. C. & McBain, C. J. Hippocampal gabaergic inhibitory interneurons. *Physiological Reviews* 97, 1619–1747 (2017).
- [54] Buzsáki, G. Neural syntax: Cell assemblies, synapse ensembles, and readers. *Neuron* 68, 362–385 (2010).
- [55] Constantinidis, C. & Klingberg, T. The neuroscience of working memory capacity and training. *Nature Reviews Neuroscience* 17, 438–449 (2016).
- [56] Lee, S., Kim, S., Kwon, O., Lee, J. & Kim, J. Inhibitory networks of the amygdala for emotional memory. *Frontiers in Neural Circuits* 7, 129 (2013).
- [57] Pfeffer, C. K., Xue, M., He, M., Huang, Z. J. & Scanziani, M. Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nature Neuroscience* 16, 1068–1076 (2013).
- [58] Booker, S. A. & Vida, I. Morphological diversity and connectivity of hippocampal interneurons. *Cell Tissue Research* 373, 619–641 (2018).

- [59] Zeng, H., Chattarji, S., Barbarosie, M., Bondi-Reig, L., Philpot, B. D., Miyakawa, T., Bear, M. F. & Tonegawa, S. Forebrain-specific calcineurin knockout selectivity impairs bidirectional synaptic plasticity and working/episodic-like memory. *Cell* 107, 617–629 (2001).
- [60] Froemke, R., Poo, M. & Dan, Y. Spike-timing-dependent synaptic plasticity depends on dendritic location. *Nature* 434, 221–225 (2005).
- [61] Chaudhuri, R. & Fiete, I. Computational principles of memory. *Nature Neuroscience* 19, 394–403 (2016).
- [62] Wills, T. J., Lever, C., Cacucci, F., Burgess, N. & O’Keefe, J. Attractor dynamics in the hippocampal representation of the local environment. *Science* 308, 873–876 (2005).
- [63] Fields, R. D., Araque, A., Johansen-Berg, H., Lim, S., Lynch, G., Nave, K., Nedergaard, M., Perez, R., Sejnowski, T. & Wake, H. Glial biology in learning and cognition. *The Neuroscientist* 20, 426–431 (2014).
- [64] Santello, M., Toni, N. & Volterra, A. Astrocyte function from information processing to cognition and cognitive impairment. *Nature Neuroscience* 22, 154–166 (2019).
- [65] Araque, A. & Navarrete, M. Glial cells in neuronal network function. *Philosophical Transactions of the Royal Society B: Biological Sciences* 365, 2375–2381 (2010).
- [66] Khakh, B. S. & Sofroniew, M. V. Diversity of astrocyte functions and phenotypes in neural circuits. *Nature Neuroscience* 18, 942–952 (2015).
- [67] Perea, G., Navarrete, M. & Araque, A. Tripartite synapses: astrocytes process and control synaptic information. *Trends in Neurosciences* 32, 421–431 (2009).
- [68] Araque, A., Carmignoto, G., Haydon, P. G., Oliet, S. H. R., Robitaille, R. & Volterra, A. Gliotransmitters travel in time and space. *Neuron* 81, 728–739 (2014).
- [69] Giaume, C., Koulakoff, A., Roux, L., Holcman, D. & Rouach, N. Astroglial networks: a step further in neuroglial and gliovascular interactions. *Nature Reviews Neuroscience* 11, 87–99 (2010).
- [70] Verkhratsky, A. & Toescu, E. Neuronal-glia networks as substrate for CNS integration. *Journal of Cellular and Molecular Medicine* 10, 869–879 (2006).
- [71] Sasaki, T., Ishikawa, T., Abe, R., Nakayama, R., Asada, A., Matsuki, N. & Ikegaya, Y. Astrocyte calcium signalling orchestrates neuronal synchronization in organotypic hippocampal slices. *The Journal of Physiology* 592, 2771–2783 (2014).
- [72] Butler, C. R., Graham, K. S., Hodges, J. R., Kapur, N., Wardlaw, J. M. & Zeman, A. Z. J. The syndrome of transient epileptic amnesia. *Annals of Neurology* 61, 587–598 (2007).

- [73] Butler, C. R. & Zeman, A. Z. Recent insights into the impairment of memory in epilepsy: transient epileptic amnesia, accelerated long-term forgetting and remote memory impairment. *Brain* 131, 2243–2263 (2008).
- [74] Holmes, G. L. Cognitive impairment in epilepsy: the role of network abnormalities. *Epileptic Disorders* 17, 101–116 (2015).
- [75] Seifert, G., Schilling, K. & Steinhäuser, C. Astrocyte dysfunction in neurological disorders: a molecular perspective. *Nature Reviews Neuroscience* 7, 194–206 (2006).
- [76] Cui, Y., Prokin, I., Mendes, A., Berry, H. & Venance, L. Robustness of stdp to spike timing jitter. *Nature Scientific Reports* 8, 1–15 (2018).
- [77] Faisal, A. A., Selen, L. P. & Wolpert, D. M. Noise in the nervous system. *Nature Reviews Neuroscience* 9, 292–303 (2008).
- [78] van Strein, N., Cappaert, N. & Witter, M. The anatomy of memory: an interactive overview of the parahippocampal-hippocampal network. *Nature Reviews Neuroscience* 10, 272–282 (2009).
- [79] Burwell, R. The parahippocampal region: Corticocortical connectivity. *Annals of the New York Academy of Sciences* 911, 25–42 (2006).
- [80] Chindemi, G., Abdellah, M., Amsalem, O., Benavides-Piccione, R., Delattre, V., Doron, M., Ecker, A., Jaquier, A. T., King, J., Kumbhar, P., Monney, C., Perin, R., Rössert, C., Tuncel, A. M., Geit, W. V., DeFelipe, J., Graupner, M., Segev, I., Markram, H. & Muller, E. B. A calcium-based plasticity model for predicting long-term potentiation and depression in the neocortex. *Nature Communications* 13, 1–19 (2022).
- [81] Bartol Jr, T. M., Bromer, C., Kinney, J., Chirillo, M. A., Bourne, J. N., Harris, K. M. & Sejnowski, T. J. Nanoconnectomic upper bound on the variability of synaptic plasticity. *Elife* 4, e10778 (2015).
- [82] Buzsáki, G. & Draguhn, A. Neuronal oscillations in cortical networks. *Science* 304, 1926–1929 (2004).
- [83] Marder, E. Neuromodulation of neuronal circuits: back to the future. *Neuron* 76, 1–11 (2012).
- [84] Bargmann, C. I. Beyond the connectome: how neuromodulators shape neural circuits. *Bioessays* 34, 458–465 (2012).
- [85] George, V. K., Morar, V., Yang, W., Larson, J., Tower, B., Mahajan, S., Gupta, A., White, C. & Silva, G. A. Learning without gradient descent encoded by the dynamics of a neurobiological model (2021). URL <https://arxiv.org/abs/2103.08878>.

- [86] Silva, G. A. Neuroscience and artificial intelligence can help improve each other. *The Conversation* (2019). URL <https://theconversation.com/neuroscience-and-artificial-intelligence-can-help-improve-each-other-110869>.
- [87] Tavanei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T. & Maida, A. Deep learning in spiking neural networks. *Neural Networks* 111, 47–63 (2019).
- [88] Tang, J., Yuan, F., Shen, X., Wang, Z., Rao, M., He, Y., Sun, Y., Li, X., Zhang, W., Li, Y., Gao, B., Qian, H., Bi, G., Song, S., Yang, J. J. & Wu, H. Bridging biological and artificial neural networks with emerging neuromorphic devices: Fundamentals, progress, and challenges. *Advanced Materials* 31, 1–33 (2019).
- [89] Taherkhani, A., Belatreche, A., Li, Y., Cosma, G., Maguire, L. P. & McGinnity, T. M. A review of learning in biologically plausible spiking neural networks. *Neural Networks* 122, 253–272 (2020).
- [90] Kriegeskorte, N. & Golan, T. Neural network models and deep learning. *Current Biology* 29, R225–R240 (2019).

Chapter 3 Simplifying Biology Vocabulary via Morphology

3.1 Abstract

In all fields of biology, understanding technical terminology is a challenge for students. In many cases, this may distract them from focusing on fundamental processes and concepts. Across the biology subfields, much of the vernacular shares similar etymology and morphology. However, students lack the exposure necessary to identify these key features, which often explain the meaning of terms without requiring any context at all. Therefore, instead of encouraging students to memorize many terms independently, it could be more beneficial to show them how words are constructed. Here, I propose an activity designed to help students recognize terms that may be connected, understand how vocabulary is often constructed to reflect its idea, and develop comfortability using these terms themselves in discussions. Through a guided group activity, students will have a chance to break down terms they have previously encountered and to draw connections between novel words. If students are capable of relating words to each other before even knowing what they mean, they may learn more effectively. Without being intimidated by enigmatic vocabulary, they can focus on broader concepts. In addition, when students understand how biological terminology is constructed, they may even dissect new words without needing the context surrounding them. This activity is applicable to courses in any specialty of biology, as various molecules, tissues, and processes follow general naming principles.

3.2 Introduction

Novice learners often struggle to discern which concepts are most important when presented with novel material [4]. As such, it is easy for students to get stuck on small details when they should be focusing on the bigger picture. In biology, there are myriad names for

molecules, tissues, and processes that may confuse students [16]. Without a full understanding of these terms, students are left with an incongruence between what a word looks like and what it refers to. This divide forces them to memorize the foreign word while also trying to grasp the biology behind why it is important [23, 6]. Introductory and advanced biology courses introduce as many as 1,000 technical terms [19, 7, 26]. Of those, many are heavily rooted in Greek or Latin. If students were primed beforehand to identify common biological morphemes, they could spend less time memorizing words and instead focus on the concepts themselves [7, 8, 2, 1]. Comparing the morphologies of words may even result in novel connections for the students [26]. Educational schema theory suggests that students can implement this knowledge of terminology to improve comprehension and retention of deeper ideas [25, 24, 12].

For students whose first language is not English, this approach may simplify scientific terms for them [10]. Students with first languages based in Latin or Greek more easily access STEM terminology when they can connect the etymologies to their home language [21, 13, 9]. Specifically, English Language Learner students (ELLs) in the U.S. received higher scores on science units when extra attention was placed on the vocabulary itself [11, 5]. Notably, verbal usage of academic language was critical to making ELLs comfortable with the terminology [15]. Further, ELLs with reading disabilities in the U.S. significantly improved their understanding of biological vocabulary following instruction on different morphemes [14, 13].

More broadly, if the countless terms in biology are thought of as a new language [19, 17, 23], their instruction may benefit from approaches designed to teach students languages. In particular, the Model of Domain Learning supports the idea that memorizing a vocabulary list is not as effective for learning a new language as identifying context clues necessary to discern the meaning of an unfamiliar word [20].

To this point, significant work has been done to effectively introduce students to biological terminology. In one study, an anatomy course in Minnesota implemented a five-day mini-unit focused on word roots and observed significant improvement on a vocabulary exam [18]. In addition to quantitative data, students provided positive feedback about how the mini-unit helped them. Similarly, a high school class was split into groups studying Greek and Latin roots either via rote memorization or an active mnemonic activity [3]. After nine weeks of instruction, there was a significant improvement in medical terminology recall in the groups that participated in the mnemonic activity, demonstrating the benefit of familiarity with these roots. Another study involved starting a biology course with a worksheet and a test on 40 common root words [17]. While there was no student feedback or performance data, priming students for terminology could only help them prepare [25]. An instructor in Texas even tried to incorporate teaching biological etymology into a dissection activity to improve students' understanding of what they were doing [8]. While time constraints impacted the performance of that activity, the idea that root words are useful persists ubiquitously.

Beyond etymology, active discussion is an important part of my proposal as well. One group prompted student pairs to hold a dialogue incorporating scientific word lists [19]. Conversing using technical terminology proved to be more fruitful for students than having a simple reading assignment. In a liberal arts college in the Midwest, one instructor designed a card game to teach students biological vocabulary and received feedback that it “helped with terminology and figuring out and connecting words and ideas together [6].” Further, one instructor put together a card sorting assignment to make students organize chemical substances [22]. By learning about the substances as they sorted them, students quickly identified what distinguished one from the others and what details were most important. Actively engaging

students to deliberate and organize scientific terms demonstrably benefits their understanding of broader concepts.

To teach students how to break down biological vocabulary, I propose a guided group discussion in which students identify key similarities between words and interrogate why said words are constructed as they are. I recommend this activity early in the school year for students between late high school and early graduate school, so they will have encountered some technical terms beforehand. In its current form, this activity can be applied to classes of any size and topic, but for large classes, instructional assistants may be necessary to aid with monitoring the groups. By allowing students to work together to hypothesize about connections, they will approach biological ideas from a new perspective. They will acquire an awareness of what goes into these words, which is crucial to reverse engineer what they mean.

3.3 Activity Preparation

To properly scaffold conversations between students, instructors must prepare a small introduction for the activity. This should include a general definition of morphology, noting how many technical words are constructed from Latin or Greek roots. One example I recommend is adrenaline-epinephrine. Both words refer to the same hormone and their similarities permeate to their etymologies. For adrenaline, the roots are “ad-” (towards), “-renal” (kidneys), and “-ine” (made of). Similarly, the roots of epinephrine are “epi-” (on top of), “nephro-” (kidney), and “-ine” (made of). These words both highlight where the hormone is produced.

However, instead of diving right into this etymological homology, I recommend introducing these roots individually first. For example, an instructor can guide students to break down “adapt” to understand the root “ad-”. Similarly, “epilogue” demonstrates the flexibility of

“epi-”. By starting with familiar words, it will be easier for students to grasp the idea of defining words via roots.

Other than this introduction, a word bank must be compiled. This word bank may focus on the specific class or may generalize to various subjects in biology. Words selected should have straightforward morphemes and among them, there should be overlaps in root words for students to identify. To increase student participation, students can be tasked with finding five to ten technical terms themselves. To guide their search, instructors can recommend aggregating terms from other classes or that students may have encountered in recent news.

Along with a bank of words, a bank of commonly used roots must be compiled [17]. Special care should be placed on the creation of this bank, as students may take it with them to future classes as well. The more comprehensive the list, the more useful it could be for them in other settings.

3.4 Activity Implementation

Overall, this group discussion should take around one hour in total regardless of class size. To begin, the instructor shall introduce morphemes and how common they are in biology. Once this brief lecture is complete, students should divide into groups of three or four. Per the instructor’s discretion, a different group size could be chosen, but too large would risk some students not having the time or space to contribute and too small may lead to discussions with too few observations.

With the word bank on screen or accessible through their devices, students should note down as many similarities between words as they can. By connecting words superficially, they will create clusters to work with when trying to understand roots. Students should have 5 minutes to group as many words as they can. After this, students should explicitly identify possible roots

and hypothesize about what they might mean in the context of these words. One student per group should be chosen via the instructor's preferred method to write down the groups of words and hypothetical roots as the discussion proceeds. If students are particularly stuck on a cluster, they may need to look up what certain words mean. This experimental period for students should last around ten minutes, but may be extended if the students are very active. When assessing how students are doing, a good discussion may include references to particular organs or a specific chemical process. References to the anatomical orientation described by roots would also indicate an ideal conversation. Poor discussions may involve students getting caught up on the location of roots within a word or spending too much time trying to define a particular root that they've identified.

Following this discussion, groups should come back together as a class and each one should highlight one root of interest. This should include all the words they identified with the root and possible definitions for them. After a group shares, other groups should have the opportunity to comment if they noticed other candidate words or if they think certain words actually do not use that particular root.

Once the class has finished, the instructor should provide the bank of roots to students on screen and on their devices. Students should return to their groups and reflect on how accurate their word clusters were. After assessing their previous hypotheses, students should tackle defining unknown words in the word bank once again, but this time using the bank of roots provided. In addition to the clusters they formed, this will allow them to directly dissect technical terms, possibly for the first time ever. Students should be given around ten minutes for this as well, so they can fully apply the roots together.

Following this discussion period, the class should come together and share one final time. Groups should share one word that they found interesting, either because of its etymological basis or because the roots appropriately contextualize what is important about it. For instance, the adrenaline example demonstrated an interesting etymological basis, but did not contextualize what its function is. Different students may find either case more appealing to them.

Once sharing concludes, an instructor may briefly lecture to reflect on the value of morphology. Other interesting terms may be highlighted here at the instructor's discretion. To conclude, a transition between this exercise and key terms of the next unit could serve the students well.

3.5 Conclusion

This exercise informs students about the importance of morphology in biology while encouraging them to have discussions about biology with each other. By introducing these word roots in a group setting rather than as a worksheet, students can work together to highlight interesting trends or troubleshoot inconsistencies. In addition, making connections with their peers may lead to the exercise being more memorable. However, without the proper tool set, the takeaways of this activity cannot be fully utilized. As such, the root bank is an important resource to provide so students can reference it when they encounter new fields of biology. With the root bank and the experience acquired through this exercise, students can tackle novel vocabulary without concern. In addition, when concepts are introduced, students can focus on them without being distracted by technical jargon that may steal their attention. While this exercise will highlight the utility of morphology to the students, many of who may only be learning this concept for the first time, repetition is always helpful for learning. This exercise in isolation will likely benefit the students greatly, but it may be repeated with new words or

incorporated into a larger morphology unit should the instructor feel there is time to emphasize the importance of this skill.

3.6 Acknowledgements

Chapter 3, in full, is a reprint of the material as it appears in *The American Biology Teacher*, Morar, Vikash, University of California Press, 2023. The dissertation author was the primary author of this paper.

3.7 References

- [1] August, D., Branum-Martin, L., Cardenas-Hagan, E. & Francis, D. J. The impact of an instructional intervention on the science and language learning of middle grade English language learners. *Journal of Research on Educational Effectiveness* 2.4 pp. 345–376 (2009).
- [2] Boyd, F. B., Sullivan, M. P., Popp, J. S. & Hughes, M. Vocabulary instruction in the disciplines. *Journal of Adolescent & Adult Literacy* 56.1, pp. 18–20 (2012).
- [3] Brahler, C. J. & Walker, D. Learning scientific and medical terminology with a mnemonic strategy using an illogical association technique. *Advances in physiology education* 32.3, pp. 219-224 (2008).
- [4] Bransford, J. D., Brown, A. L & Cocking, R. R. *How people learn*. Vol. 11. Washington, DC: National academy press (2000).
- [5] Bravo, M. A. & Cervetti, G. N. Attending to the language and literacy needs of English learners in science. *Equity & Excellence in Education* 47.2, pp. 230–245 (2014).
- [6] Burleson, K. M. & Olimpo, J. T. ClueConnect: a word array game to promote student comprehension of key terminology in an introductory anatomy and physiology course. *Advances in Physiology Education* 40, pp. 223–228 (2016).
- [7] Cardinale, L. A. Using explication to improve vocabulary acquisition. *The American Biology Teacher* 54.3, pp. 177–178 (1992).
- [8] Chapman, A. M., Ward, H. C., Thvari, A., Weimer, A., Duran, J. B., Guerra, F. & Sale, R. P. The role of language in anatomy and physiology instruction. *The American Biology Teacher* 79.3, pp. 184–190 (2017).
- [9] De Oliveira, L. C. & Wilcox, K. C. (Eds.). *Teaching science to English language learners: Preparing pre-service and in-service teachers*. Springer (2017).

- [10] Echevarria, J., Vogt, M. E. & Short, D. J. Making Content Comprehensible for English Learners: The SIOP Model. Pearson Education, ISBN: 9780134045320 (2016).
- [11] Garza, T., Huerta, M., Spies, T. G., Lara-Alecio, R., Irby, B. J., & Tong, F. Science classroom interactions and academic language use with English learners. *International Journal of Science and Mathematics Education* 16.8, pp. 1499–1519 (2018).
- [12] Hayden, E., Baird, M. E. & Singh, A. Learning the words AND knowing the concepts: an in-depth study of one expert teacher’s use of language as a cultural tool to support inquiry. *Literacy* 54.1, pp. 18–28 (2020).
- [13] Helman, A., Dennis, M. S. & Kern, L. Clues: Using Generative Strategies to Improve the Science Vocabulary of Secondary English Learners With Reading Disabilities. *Learning Disability Quarterly* 45.1, pp. 19–31 (2022).
- [14] Helman, A. L., Calhoun, M. B. & Kern, L. Improving science vocabulary of high school English language learners with reading disabilities. *Learning Disability Quarterly* 38.1, pp. 40–52 (2015).
- [15] Hoffman, L. & Zollman, A. What STEM teachers need to know and do for English language learners (ELLs): Using literacy to learn. *Journal of STEM Teacher Education* 51.1, p. 9 (2016).
- [16] Hoots, R. A. Biological Bywords. *The American Biology Teacher*, pp. 335–337 (1991).
- [17] Kessler, J. W. An alternative approach to teaching biological terminology. *The American Biology Teacher*, pp. 688–690 (1999).
- [18] King, E. The Effect of Increased and Varied Vocabulary Instruction on Science Content Acquisition. (2022).
- [19] May, S. R., Cook, D. L. & May, M. K. Biological dialogues: how to teach your students to learn fluency in biology. *The American Biology Teacher* 75.7, pp. 486–493 (2013).
- [20] Parkinson, M. M. & Dinsmore, D. L. Understanding the developmental trajectory of second language acquisition and foreign language teaching and learning using the Model of Domain Learning, *System* 86, p. 102125 (2019).
- [21] Suh, E. K., Hoffman, L. & Zollman, A. STEM inclusion research for English language learners (ELLs): Making STEM accessible to all. *Handbook of research on STEM education*. Routledge, pp. 311–322 (2020).
- [22] Wallon, R. C., Jasti, C. & Hug, B. A Card-Sorting Activity to Engage Students in the Academic Language of Biology. *The American Biology teacher* 79.3, pp. 233–237 (2017).

[23] Wandersee, J. H. The terminology problem in biology education: A reconnaissance. *The American Biology Teacher* 50.2, pp. 97–100 (1988).

[24] Wilhelmsson, N., Bolander-Laksov, K., Dahlgren, L.O., Hult, H., Nilsson, G., Ponzer, S., Smedman, L. & Josephson, A. Long-term understanding of basic science knowledge in senior medical students. *Int J Med Educ* 4, pp. 193–197 (2013).

[25] Wright, K.L., Franks, A.D., Kuo, L.J., McTigue, E.M. & Serrano, J. Both theory and practice: Science literacy instruction and theories of reading. *International Journal of Science and Mathematics Education* 14.7, pp. 1275–1292 (2016).

[26] Zoski, J. L., Nellenbach, K. M. & Erickson, K. A. Using morphological strategies to help adolescents decode, spell, and comprehend big words in science. *Communication Disorders Quarterly* 40.1, pp. 57–64 (2018).

CONCLUSION

Across all three chapters, one unifying thread was the importance of learning in the brain and in networks attempting to mimic the brain. From the myriad learning rules described in Chapter 1 to the methodology proposed in Chapter 3, this has been an endeavor to better understand how learning works from the molecular scale all the way to the behavioral scale in a classroom. While the two may seem disparate at first glance, I think there is interesting perspective to gain approaching learning from both ways. Machine learning models may require out-of-the-box methods to best comprehend an input, just as a troubled student might in a classroom. Similarly, students require constant repetition to properly learn difficult concepts, just like neural networks in their training stages. In between both, neuroscientific learning rules are employing tools with various mechanisms and timescales to achieve a variety of functions in the brain. While the technical details are unique to each learning challenge, the fundamental approaches can inform solutions at every level.