SANDIA REPORT
SAND2021-11680
Printed September 2021

Sandia National Laboratories

Short Term Plasticity for Artificial Neural Networks

Corinne M. Teeter

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico
87185 and Livermore,
California 94550

Issued by Sandia National Laboratories, operated for the United States Department of Energy by National Technology & Engineering Solutions of Sandia, LLC.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@osti.gov
Online ordering: http://www.osti.gov/scitech

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5301 Shawnee Rd
Alexandria, VA 22312

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.gov
Online order: https://classic.ntis.gov/help/order-methods/

ABSTRACT

Achieving efficient learning for AI systems was identified as a major challenge in the DOE's recently released, *AI for Science*, report**.** The human brain is capable of efficient and low-powered learning. It is likely that implementing brain-like principles will lead to more efficient AI systems. In this LDRD, I aim to contribute to this goal by creating a foundation for implementing and studying a brain phenomenon termed short term plasticity (STP) in spiking artificial neural networks within Sandia. First, data collected by the Allen Institute for Brain Science (AIBS) was analyzed to see if STP could be classified into types using the data collected. Although the data was inadequate at the time, AIBS has updated their database and created models that could be utilized in the future. Second, I began creating a software package to assess the ability of a Boltzmann machine utilizing STP to sample from national security data.

ACKNOWLEDGEMENTS

I thank the LDRD S&T CIS IAT (Laboratory Directed Research and Development, Science and Technology, Computer and Information Sciences) for this fantastic training opportunity.

Table of Contents

LIST OF FIGURES

This page left blank

ACRONYMS AND DEFINITIONS

Abbreviation	Definition
STP	Short term plasticity: Biological phenomena where the amplitude of the post-synaptic response depends on the spiking history of the pre-synaptic neuron.
PSP	Post-synaptic potential: Voltage deflection elicited in a post-synaptic neuron due to neurotransmitter released from a firing pre-synaptic neuron.
STDP	Spike time dependent plasticity: Mechanism associated with learning that is different from STP.
BI	Biological intelligence: computational abilities that have evolved through evolution in animals.
AI	Artificial intelligence: Algorithms developed by humans to perform computation. In contrast to biological intelligence (BI).
RBM	Restricted Boltzmann machine
ANN	Artificial neural network
DNN	Deep neural network
Synaptic Facilitation	Amplitude of PSP increase as pre-synaptic neuron fires more
Synaptic Depression	Amplitude of PSP decrease as pre-synaptic neuron fires more
AIBS	Allen Institute for Brain Science
MNIST	Database of handwritten digits
aisynphys	Python package provided by AIBS for accessing their synaptic physiology data.

1. INTRODUCTION

1.1. Funding Background and Intention

This report is a summary of the work conducted for CIS LDRD project *Short Term Plasticity for Artificial Neural Networks*. This funding was provided as part of a hiring package to enable me as a strategic hire to assimilate into Sandia. As such, the subject material of the project was high risk but provided me valuable training on the internal LDRD process, allowed me to familiarize myself with the work being performed at Sandia, and enabled me to create a foundation for a body of work that has not previously been utilized at Sandia. This work will be important for algorithms in artificial intelligence (AI) and neuromorphic hardware utilizing spiking nodes.

1.2. Purpose and Overview of Work

There are many fundamental phenomena found in the biological brain that are not yet integrated in modern ANNs; it is likely that their implementation will yield more capable and low-powered AI useful for national security applications. In this report, I focus on one brain phenomenon, short term plasticity (STP), and aim to help create a foundation within Sandia for its implementation in future ANNs and neuromorphic hardware.

There were two stages in this project, the first was to utilize a recent, large database created by the Allen Institute for Brain Science (AIBS) to characterize potential STP types to constrain networks in future research. In the end, the AIBS database did not have enough values to perform a meaningful analysis. As a result, I moved on to the backup plan: to explore if preliminary work conducted outside the laboratory could be useful for national security applications. This previous work illustrated that STP could be used to perform superior probability distribution sampling enabling Boltzmann Machines (BM) to adjust their weights to both classify and generate images (Leng et al., 2018). To this point, I have started creating the code to replicate the results of Leng and colleagues (Leng et al., 2018). This project will likely be continued in future ASC work as the STP phenomenon is a promising algorithm for spiking networks and thus could be useful for implementation on low-power, neuromorphic hardware

Below, I describe STP, the process for accessing the AIBS synaptic physiology database and how STP could be used to create better ANNs.

2. BACKGROUND

2.1. Why Use the Brain for Inspiration?

The concept and creation of artificial neural networks (ANNs) was inspired by the brain. Now, deep neural networks (DNNs) have enabled non-linear function approximation and are surpassing human performance on specific tasks such as image processing (Russakovsky et al., 2015), natural language processing (LeCun, Bengio, & Hinton, 2015), and playing games (Mnih et al., 2015; Silver et al., 2016). DNNs trained via supervised learning yield representations like those found in the brain (Mante, Sussillo, Shenoy, & Newsome, 2013; Yamins & DiCarlo, 2016). Modern deep reinforcement learning algorithms have provided testable predictions leading to insights about how the brain functions(Dabney et al., 2020; J. X. Wang et al., 2018).

Despite the successes of modern day ANNs, there are severe drawbacks. They are power hungry, and a great deal of labeled data or simulation time is required for training, making real world problems computationally expensive or infeasible. Furthermore, DNNs implementing supervised learning need to be retrained when the statistics of the input change. For a more detailed description of the deficiencies of modern artificial intelligence (AI), please see the "Achieve efficient learning for AI systems" section on page 101 of the DOE, *AI for Science report*.

The human brain provides proof that a network can learn using limited amounts of data, transfer knowledge between different tasks, and easily adapt to changing environments: all while consuming extremely low power. Although we do not yet know the secrets of how the brain achieves it impressive computational feats, especially while consuming so little power, the brain implements many phenomena that the AI field does not yet widely utilize.

Importantly, the brain uses spiking nodes as opposed to the continuous nodes utilized in modern ANNs. It is likely that spiking networks will be both low power and add computational capabilities due to the ability to perform event-based programing. Algorithms such as STP evolved from a biological spiking networks and therefore will likely be important algorithms to enable artificial spiking networks to perform computation.

2.2. The Complexity of the Brain and Allen Institute for Brain Science (AIBS) Database

Although utilizing biological brain principles in ANNs might seem like a "no-brainer", in practice it is quite difficult given the complexity of the brain and how difficult it is to study. In standard ANNs there are only several type of node characterized by their activation functions and the connection weights between the nodes are trained via extensions of the traditional backpropagation algorithm (Rumelhart, Hinton, & Williams, 1986). In the brain, the diversity of neurons and connections between them can seem overwhelming. There are many different classes of neurons, with different transcriptomic profiles, different morphologies, and different firing behaviors. In addition, different synapses have different dynamic behaviors. Understanding the components, and connections between them, are fundamental to understanding any network; thus, the characterization of neurons and synapses has been an active area of research for decades.

To help characterize the complexity of the brain the AIBS created the Cell Types program to collect transcriptomic, morphological, and electrical data from neurons at a large scale (this effort has now been taking place for more than 10 years). After AIBS data has been collected and has passed quality control, it is released to the public via their website (https://portal.brain-map.org). The hope is that the community will utilize the data and potentially discover insights to how the brain operates.

In the first stage of my work, I will analyze data in the synaptic physiology database to characterize STP parameters for future implementation AI.

2.3. Short Term Plasticity (STP)

STP describes how communication between neurons evolves based on pre-synaptic activity over a short time scale (milliseconds to seconds). STP has the capability to enable many computations including filtering (Figure 1), enhancement of transients, decorrelation, burst detection, stimuli presentation duration, sound location over different intensity ranges, dynamic input compression, working memory, and spatiotemporal processing (Abbott & Regehr, 2004; Buonomano & Maass, 2009; Mongillo, Barak, & Tsodyks, 2008). Note, that STP should not be confused with spike time dependent plasticity (STDP) also known as Hebbian learning which takes place on the order of minutes to hours. These are two different synaptic plasticity mechanisms. Generally, STDP takes place in learning where STP enables the computations mentioned above.

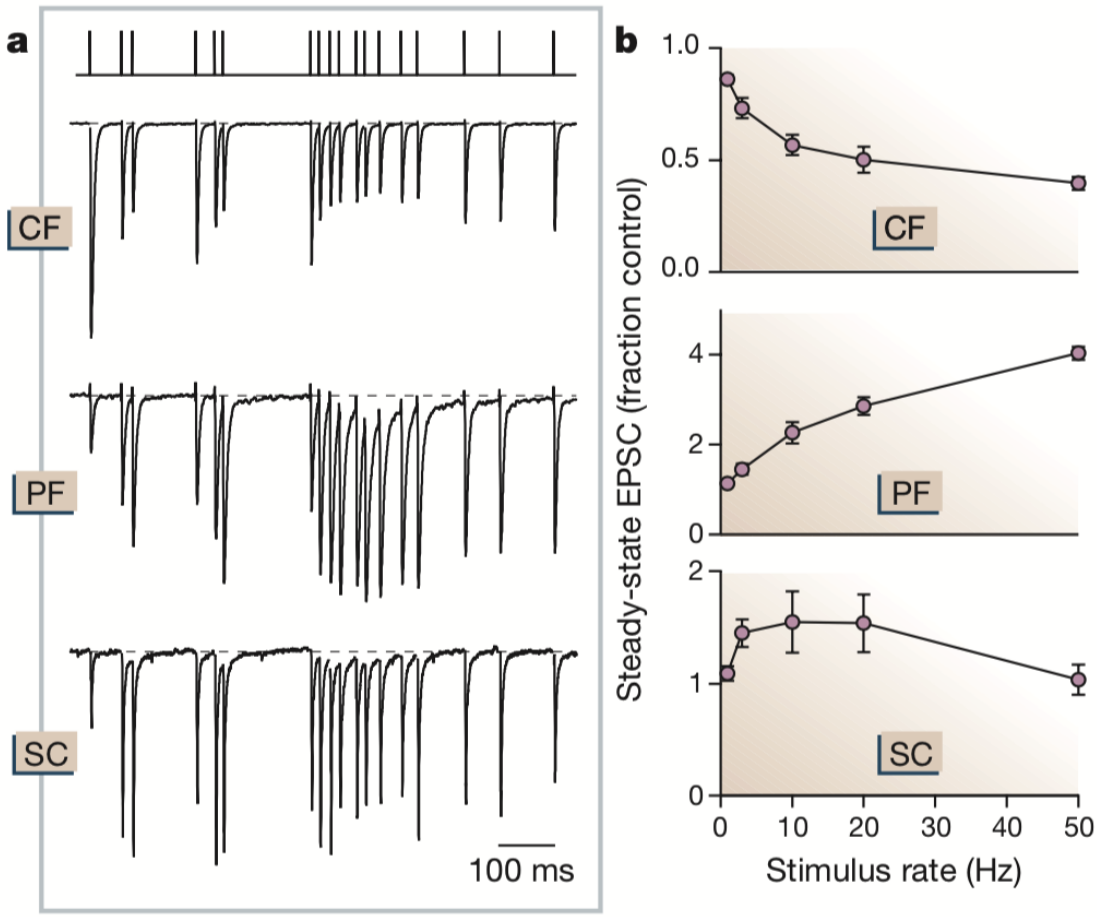

Figure 1. Example of filtering. Figure taken from Abbott & Regehr, 2004. a) Illustrates the diversity of STP. Top line is the firing of the presynaptic neuron. CF, PF, and SC show the resulting PSPs in post-synaptic neurons from different regions of the brain (An explanation of specific regions is outside the scope of this document). Notice that the amplitude of PSPs change in different ways for the different neurons. b) Illustrates the corresponding filtering properties. Y-axis represents the normalized relative height of PSPs given steady input frequencies. Thus CF is a low pass filter (amplitude gets smaller when there is a high frequency of presynaptic spiking), PF is high pass

filter (amplitude gets larger when there is a high frequency of presynaptic spiking), and SC acts as a band pass filter.

2.3.1. Mechanism

When a neuron receives enough positive stimulation from upstream neurons to reach its voltage threshold, a large and quick voltage fluctuation, referred to as a spike (the change in voltage during a spike is usually between 100 to 200 mV), is evoked. The spike causes neurotransmitter to be released from the firing (pre-synaptic) neuron in small packets called vesicles. This neurotransmitter travels to the post-synaptic neuron and elicits a small (usually between 0 and 5 mV) voltage fluctuation referred to as a post synaptic potential (PSP). The size and the shape of these PSPs can change as the result of the firing patterns of the pre-synaptic neurons**.** For example, if the pre-synaptic neuron fires several sequential spikes, the amplitude of the PSPs can sequentially increase (referred to as a facilitating synapse), decrease (referred to as a depressing synapse), or have some other more complicated behavior (Abbott & Regehr, 2004).

2.3.2. Difficulties in Characterizing STP

How similar is the STP evoked from different pre-synaptic cell types (or post-synaptic types)? Are there different behaviors between specific pre and post-synaptic type pairs? What is the most important characteristic for characterizing STP? Are there parameters or combinations of parameters that seem to define synaptic types? Addressing these questions is non-trivial due to the inherent caveats of synaptic physiology data. PSPs can be very small, and hard to discriminate due to noise in the recording. The typical solution of taking the signal mean to average out the noise can be misleading because synapses are stochastic. Different synapses have different levels of stochasticity which is an important characteristic of synapses. For example, say you want to characterize the amplitude of the PSPs from a particular synapse. When the synapse is activated, it typically releases one vesicle of neurotransmitter, but it only activated half of the time a spike occurs. One vesicle elicits a PSP of 1 mV and this signal is difficult to observe in the noise. Naively, the mean of the signals is taken, and it is concluded that the amplitude of the PSP is 0.5 mV, when in actuality, the synapse has an amplitude of 1 mV with a release probability of 50 percent. This is just one of many difficulties in analyzing synaptic physiology data, the full extent of which, is outside the scope of this document. However, there is one situation which requires mentioning as it directly affects this analysis. Occasionally, the presynaptic neuron will fail to spike even though current was injected. In the case where structured data analysis is being utilized (further discussed below), the entire train of data needs be discarded because the synaptic dynamics are subject to different patterns of pre-synaptic firing.

In an attempt overcome the challenges associated with synaptic physiology data and address the aforementioned questions, two different methods (modeling and data analysis) can be employed. These techniques have different strengths and weaknesses. The model-based approach was currently being developed at AIBS. My aim was to perform dimensionality reduction and clustering.
In the model-based approach, models will be fit to the data and the resulting parameters can be characterized. The strength of this approach is that it does not require structured data; an entire spike train does not need to be discarded if the pre-synaptic neuron fails to spike. Instead, the relevant variable is the time that has passed since the previous spike; the spike failure can be ignored, and the progression in time can continue to evolve from the last successful spike. In fact, as long as the current injection stimulus elicits spikes that cover the dynamic range of the synapses, structured frequencies are not necessary for this method. The drawback of the model-based approach is that

current models only utilize the amplitude of the PSP, thus the other measurements that may be informative for characterization are ignored.

2.4. Short Term Plasticity in Restricted Boltzmann Machines

Leng et al., 2018 illustrated the potential for STP algorithms to sample from probability distributions. In their study, they demonstrate that restricted Boltzmann machines (RBMs) with connections implementing STP could classify images. In addition, they can generate a larger variety of realistic images than algorithms that implement Gibbs sampling or adaptive simulated tempering (AST) (Figure 2). Although Boltzmann Machines are not currently a mainstream ANN algorithm (due to the amount of time they take to train), they have well understood mathematical properties that make them an ideal to demonstrate the sampling capabilities of STP.

Figure 2: Section of Figure 4 from Leng et al., 2018. d) shows input given to trained Boltzmann Machine. e) and f) show how "spiking" sampling produces both 3's and 5's when the input, in d) could realistically result in either. AST and Gibbs sampling only generates 3's.

2.4.1. Restricted Boltzmann Machines (RBMs)

RBMs are a version of a Boltzmann Machine (Ackley, Hinton, & Sejnowski, 1985; Hinton, 2007, 2012; Hinton, Sejnowski, & Ackley, 1984) where the nodes are organized within two layers. The layers are fully connected in between layers but there are no connections within layers (Figure 3).

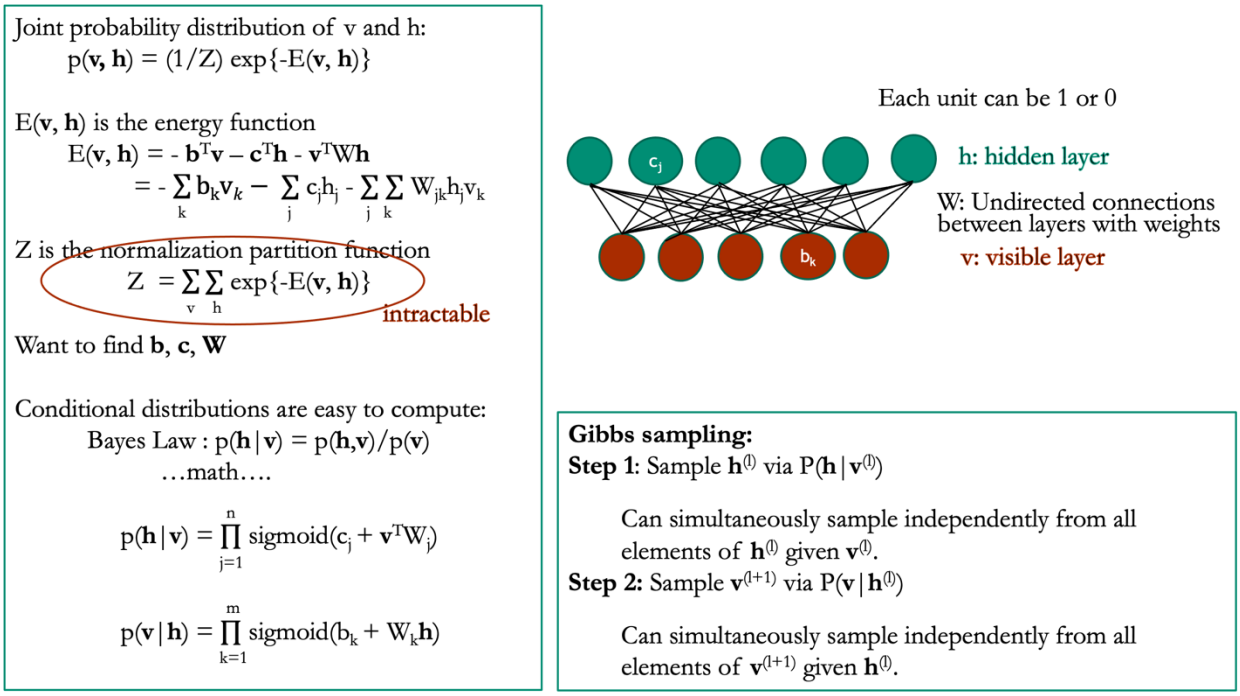

Figure 3. Illustration of Restricted Boltzmann Machine (RBM)

3. PROJECT WORK AND OUTCOME

3.1. Allen Institute Synaptic Physiology Database

To help characterize STP, the Synaptic Physiology team within the Cell Types program at AIBS has recorded the electrical activity evoked in post-synaptic neurons as a result of the spiking of pre-synaptic neurons. Synaptic physiology experiments are notoriously difficult and slow; the team has spent many years recording the activity of thousands of pairs on neurons. The experiments are performed by inserting separate electrodes into the pre and post-synaptic neurons in a slice of brain tissue viewed though a microscope. Current is injected via the pre-synaptic electrode in order to elicit a spike in the pre-synaptic neuron. The PSP in the post-synaptic neuron resulting from the pre-synaptic spike, is recorded via the post-synaptic electrode. An experimenter does not know which neurons are connected a priori, so connected neurons must be found via trial and error. To characterize how the size and shape of the PSPs change based on the specific behavior of the pre-synaptic spiking, many pulses are injected into the pre-synaptic neuron at different frequencies (typically between 20 and 100 Hz) to elicit a train of spikes at the corresponding frequencies in the pre-synaptic neuron. The resulting PSPs in the post-synaptic neuron are recorded for analysis. In general, the shape of a PSP is an asymmetric bump which can be modeled as a double exponential (Campagnola et al., 2021; Seeman et al., 2018). From each PSP, four parameters are extracted: latency from pre-synaptic spike to the PSP initiation, the rise-time from the initiation to peak, the amplitude, and the time constant of the decay. These measurements are made for PSP elicited at different holding potentials and in both voltage and current clamp mode.

3.1.1. Accessing the Database and Analysis

The code I wrote to perform the analysis is in the cmt_analysis folder of the cmt_analysis branch at https://gitlab.sandia.gov/cmteete/aiephys/-/tree/cmt_analysis/cmt_analysis. This code utilizes the AIBS aisynphys package available at https://github.com/AllenInstitute/aisynphys.git. I used Python 3.8

3.1.1.1. Extracting the Data

The first step in the analysis is to be able to understand and query the available quality controlled data. This is done in the code at https://gitlab.sandia.gov/cmteete/aiephys/-/blob/cmt_analysis/cmt_analysis/get_data.py using helper modules in lib.py.

3.1.1.2. Analysis

The available parameters have various levels of correlation Figure 4. This plot is made via https://gitlab.sandia.gov/cmteete/aiephys/-/blob/cmt_analysis/cmt_analysis/look_for_correlations.py.

It is possible the parameters could be used to classify different types of STP using either supervised or unsupervised learning. I started down this path in feature_classification.py and wrote supervised classification algorithms for the data. However, I discovered that only 75 out of the 2335 synapses were not missing data. Given the correlation plot in Figure 4, I considered attempting to impute values based on correlations. However, I spoke with the synaptic team at AIBS and they said they experienced difficulties analyzing the data in this way too, and that it probably was not worth my time. In addition, they were having success with the modeling-based approach as discussed in section 2.3.2. Difficulties in Characterizing STP. At this point, I decided if we needed to constrain model parameters in future networks, we could use the outputs of their models (as

discussed in the next section, 3.1.2 Notes On Recent AIBS Database Updates). The code of accessing the database and the preliminary code for testing supervised algorithms is already in place.

At this point I decided it would be a better investment to have me move on to my back-up plan (as agreed upon in the project proposal, see) of extending the work of Leng and colleagues (Leng et al., 2018). This would train me to utilize these principles in ANNs which will be important in future work (see section 4. Future Work)

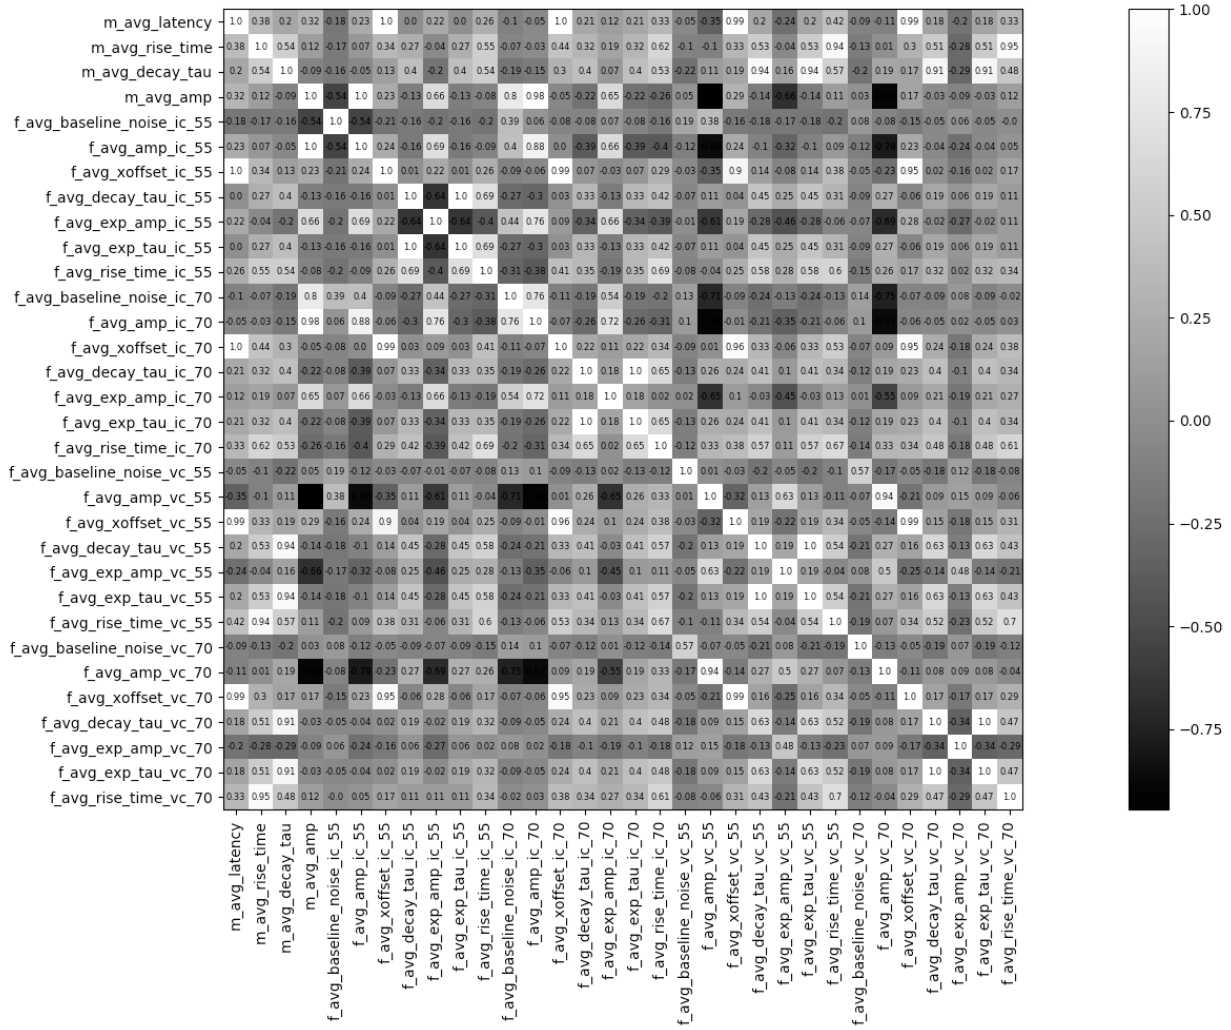

Figure 4: Correlation matrix of data within synaptic physiology database. Legend and values in each square report Spearman's correlation coefficient.

3.1.2. Notes On Recent AIBS Database Updates

Upon noticing that the parameters sets were not complete enough to perform supervised and unsupervised clustering techniques that would yield meaningful results. I contacted the AIBS Synaptic Physiology team who agreed and said they were moving forward using models for possible classification. Models have sense been released in the database and a description of the characterization of the parameters can be found at (Campagnola et al., 2021). In addition a tutorial on the most recent database can be found at

https://colab.research.google.com/github/AllenInstitute/CNS_2021/blob/main/1_Synaptic_Physiology/SynPhys_Tutorial.ipynb. If we need to constrain our networks we can use the code I have already written to extract the model parameters.

3.2. Short Term Plasticity in Boltzmann Machines for National Security Data

3.2.1. Limitations of Previous Studies

Although the work of Leng and colleagues (Leng et al., 2018), (described in the Background) is an intriguing example showing the potential of STP as a sampling algorithm, they used very limited data, the most complicated being the MNIST database (Figure 5). Thus, it is unclear if this method will be effective on more sophisticated data sets including the sort of data sets utilized in national security applications. In addition, the code for the Leng et al, 2018 study is out of date (written in python 2.7) https://github.com/electronicvisions/spike-based-sampling, and utilizes the PyNN software package (https://menloservice.sandia.gov/http://neuralensemble.org/PyNN/) to interface with neuroscience based simulators. Although PyNN is a useful language for those who need to interface with different neuron simulators to model brain function, there are various issues with memory consumption. Thus, this code is complicated to use and not suited to more general machine learning applications. Here, I aim to recreate a codebase that can implement the STP based sampling method in current software aimed at machine learning as opposed to simulating the brain.

3.2.2. Data

The (Leng et al., 2018) study uses only a limited amount of data such as MNIST (http://yann.lecun.com/exdb/mnist/). Here aim to explore to explore more sophisticated datasets such as FashionMNIST and CFAR100. Importantly, I will utilize data sets that may be used in national security type situations such as data SAR data of geophysical phenomena (C. Wang et al., 2019; Wang Chen, 2018)

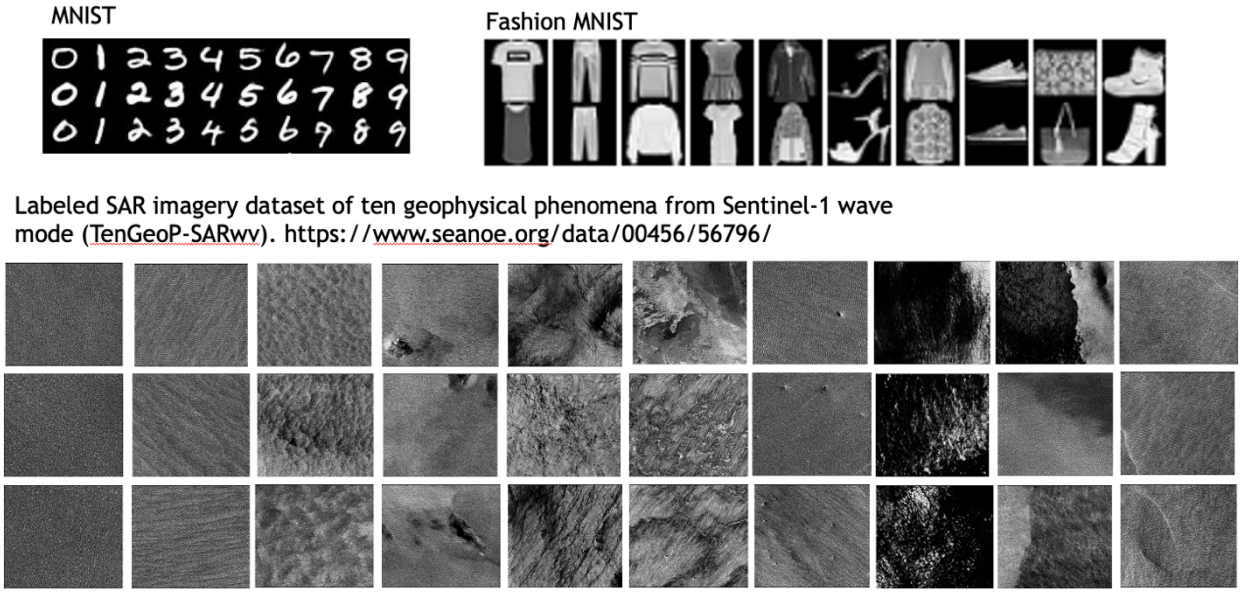

Figure 5. Example data sets for testing classification and generative capabilities of networks

3.2.3. Code and Status

I plan to use Python 3 along with the PyTorch neural network Python package to implement the Boltzmann machine (and potentially other networks). In order to implement spiking neurons, I plan to utilize BindsNET (https://github.com/BindsNET/bindsnet) (Hazan et al., 2018), which is a python spiking neuron package oriented toward machine learning (as opposed to understanding the brain) with an active user base. The development of this code is not yet complete. The code repository is located at https://cee-gitlab.sandia.gov/cmteete/stpnetwork and will be made available as it develops.

4. FUTURE WORK

Although the code to test STP algorithms is not yet complete, it is likely I will continue this work as part of the ASC Beyond Moore's Law project aimed to develop algorithms for neuromorphic hardware. STP is a fundamental phenomenon within the brain and will likely be an instrumental algorithm for complex capabilities in spiking neural networks

REFERENCES

Abbott, L. F., & Regehr, W. G. (2004). Synaptic computation. *Nature, 431*(7010), 796-803. doi:10.1038/nature03010

Ackley, D. H., Hinton, G. E., & Sejnowski, T. J. (1985). A learning algorithm for Boltzmann machines. *Cognitive science, 9*(1), 147-169.

Buonomano, D. V., & Maass, W. (2009). State-dependent computations: spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience, 10*(2), 113-125. doi:10.1038/nrn2558

Campagnola, L., Seeman, S. C., Chartrand, T., Kim, L., Hoggarth, A., Gamlin, C., . . . Jarsky, T. (2021). Local Connectivity and Synaptic Dynamics in Mouse and Human Neocortex. *bioRxiv*, 2021.2003.2031.437553. doi:10.1101/2021.03.31.437553

Dabney, W., Kurth-Nelson, Z., Uchida, N., Starkweather, C. K., Hassabis, D., Munos, R., & Botvinick, M. (2020). A distributional code for value in dopamine-based reinforcement learning. *Nature, 577*(7792), 671-675. doi:10.1038/s41586-019-1924-6

Hazan, H., Saunders, D. J., Khan, H., Patel, D., Sanghavi, D. T., Siegelmann, H. T., & Kozma, R. (2018). BindsNET: A Machine Learning-Oriented Spiking Neural Networks Library in Python. *Frontiers in Neuroinformatics, 12*(89). doi:10.3389/fninf.2018.00089

Hinton, G. E. (2007). Boltzmann machine. *Scholarpedia, 2*(5), 1668.

Hinton, G. E. (2012). A practical guide to training restricted Boltzmann machines. In *Neural networks: Tricks of the trade* (pp. 599-619): Springer.

Hinton, G. E., Sejnowski, T. J., & Ackley, D. H. (1984). *Boltzmann machines: Constraint satisfaction networks that learn*: Carnegie-Mellon University, Department of Computer Science Pittsburgh, PA.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436-444. doi:10.1038/nature14539

Leng, L., Martel, R., Breitwieser, O., Bytschok, I., Senn, W., Schemmel, J., . . . Petrovici, M. A. (2018). Spiking neurons with short-term synaptic plasticity form superior generative networks. *Scientific Reports, 8*(1), 10651. doi:10.1038/s41598-018-28999-2

Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature, 503*(7474), 78-84. doi:10.1038/nature12742

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., . . . Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature, 518*(7540), 529-533. doi:10.1038/nature14236

Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic Theory of Working Memory. *Science, 319*(5869), 1543. doi:10.1126/science.1150769

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature, 323*(6088), 533-536. doi:10.1038/323533a0

Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., . . . Fei-Fei, L. (2015). ImageNet Large Scale Visual Recognition Challenge. *arXiv:1409.0575v3*

Seeman, S. C., Campagnola, L., Davoudian, P. A., Hoggarth, A., Hage, T. A., Bosma-Moody, A., . . . Jarsky, T. (2018). Sparse recurrent excitatory connectivity in the microcircuit of the adult mouse and human cortex. *eLife, 7*, e37349. doi:10.7554/eLife.37349

Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., . . . Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature, 529*(7587), 484-489. doi:10.1038/nature16961

Wang, C., Mouche, A., Tandeo, P., Stopa, J. E., Longépé, N., Erhard, G., . . . Chapron, B. (2019). A labelled ocean SAR imagery dataset of ten geophysical phenomena from Sentinel-1 wave mode. *Geoscience Data Journal, 6*(2), 105-115. doi:https://doi.org/10.1002/gdj3.73

Wang Chen, M. A., Tandeo Pierre, Stopa Justin, Longépé Nicolas, Erhard Guillaume, Foster Ralph, Vandemark Douglas, Chapron Bertrand (2018). Labeled SAR imagery dataset of ten geophysical phenomena from Sentinel-1 wave mode (TenGeoP-SARwv). *SEANOE*. doi:https://doi.org/10.17882/56796

Wang, J. X., Kurth-Nelson, Z., Kumaran, D., Tirumala, D., Soyer, H., Leibo, J. Z., . . . Botvinick, M. (2018). Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience, 21*(6), 860-868. doi:10.1038/s41593-018-0147-8

Yamins, D. L., & DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nat Neurosci, 19*(3), 356-365. doi:10.1038/nn.4244

APPENDIX A. PROJECT SUMMARIES

A.1. New Proposal

20-1019 Proposal Summary:

Renewal Status: New
Title: Short Term Plasticity for Artificial Neural Networks Investment Area: Computing and Information Sciences Principal Investigator: TEETER,CORINNE MICHELLE ,01421
Project Manager: WAGNER,JOHN S. ,01421
Derivative Classifier: WAGNER,JOHN S. ,01421
Project Intent: Discover
Team Members: TEETER,CORINNE MICHELLE, 01421

Unclassified Unlimited Release (UUR) Congressional Summary:

Current artificial intelligence (AI) is power hungry and inflexible. Utilizing brain inspired algorithms will likely enable AI to adapt and run on low-powered systems. This project focuses on a phenomenon found in the brain termed short-term plasticity which can be implemented on low-powered neuromorphic platforms. Preliminarily implementation has demonstrated that this plasticity can enable AI to simultaneously classify difficult data and adapt to changing environments. This research will characterize short-term plasticity in order to realize adaptive and low powered AI. It will impact machine learning and AI capabilities in future Exascale systems and embedded national security systems.

Summary of Research Approach and Project Plan:

Despite the successes of modern day artificial neural networks (ANNs), there are severe drawbacks. They are power hungry, and a great deal of labeled data or simulation time is required for training, making real world problems computationally expensive or infeasible. Furthermore, deep neural networks (DNNs) implementing supervised learning need to be retrained when the statistics of the input change. For a more detailed description of the deficiencies of modern artificial intelligence (AI), please see the “Achieve efficient learning for AI systems” section on page 101 of the DOE, AI for Science report. The human brain provides proof that a network can learn using limited amounts of data, transfer knowledge between different tasks, and easily adapt to changing environments: all while consuming extremely low power.

There are many principles ubiquitously found across the brain that are not yet implemented in modern ANNs; it is likely that their implementation will yield more capable and low-powered AI useful for national security applications. However, as more brain-like complexity is introduced, more parameters will be necessary in the implementation. Finding the right combination of parameters that yield high performance will be challenging, if not impossible. Here, again we can turn to the brain for guidance. Characterizing how the brain implements phenomena will help limit the parameter space.

This project focuses on one brain phenomena, short term plasticity (STP), and aim to help create a foundation for implementation in ANNs. STP has the capability to enable many computations (Abbott and Regehr, 2004; Mongillo et al., 2008; Buonomano and Maass, 2009). Notably, recent work by Leng and colleagues (2018), have shown that spiking networks with STP can simultaneously perform well at classification and become good generative models. I aim to characterize STP using a large data set

collected by the Allen Institute for Brain Science (AIBS; https://alleninstitute.org /what-we-do/brain-science/). This characterization of STP will help provide a foundation for realistic implementations in ANNs. In addition, it is likely that this characterization will contribute to a pipeline manuscript currently being assembled at AIBS describing their methods and findings. In the event STP characterization and/or contribution to the manuscript is not possible, I will use current knowledge of STP to extend the work of Leng and colleagues (2018).

Challenge and Innovation:

This work will characterize and/or define types of short-term plasticity (STP) between different types of neurons. Linear and non-linear data analysis and dimensionality reduction techniques will be utilized to characterize STP in a new dataset released by the Allen Institute for Brain Science (https://portal.brain-map.org/explore/connectivity/synaptic- physiology).

Summary of FY20 Project Plan:

Month 1: Acquire and clean data.

Month 2 and 3: Perform basic PCA and other standard dimensionality techniques. Access how different standard characteristics are distributed across the space

End of month 3: Determine if interesting results are emerging and if there is a publication avenue with the Allen Institute for Brain Science or another potential publication avenue.

If yes: month 3-6: Refine analysis, create figures and write publication along with SAND report.

In no: month 3-6: Pursue advancing current minimal STP implementation of Leng et al., 2018 and complete SAND report.

Tie to Mission:

Recent advances in Artificial intelligence (AI) have transformed the world and national security. Current AI has power and training limitations; however, technology is quickly evolving. Biologically inspired algorithms will likely be the next breakthrough that enables low-power, next generation scientific HPC platforms for national security programs. This LDRD will impact DOE, NNSA, DHS, ASC, CSSE, and ATDM programs by contributing to new capabilities in machine learning and emerging/exascale computing programs. It will help provide new low-power computing and communication capabilities for embedded sensors and autonomous machines and help enable a new generation of ‘smart’ transportation, energy production and grid resiliency.

A.2. Proposal Continuation

21-1058 Proposal Summary

Renewal Status: Continuation
Title: Short Term Plasticity for Artificial Neural Networks Investment Area: Computing and Information Sciences Principal Investigator: TEETER,CORINNE MICHELLE ,01421 Project Manager: WAGNER,JOHN S. ,01421
Derivative Classifier: WAGNER,JOHN S. ,01421
Project Intent: Discover

Team Members: TEETER,CORINNE MICHELLE, 01421
Unclassified Unlimited Release (UUR) Congressional Summary:

Current artificial intelligence (AI) is power hungry and inflexible. Utilizing brain inspired algorithms will likely enable AI to adapt and run on low-powered systems. This project focuses on a phenomenon found in the brain termed short-term plasticity which can be implemented on low-powered neuromorphic
platforms. Preliminarily implementation has demonstrated that this plasticity can enable AI to simultaneously classify difficult data and adapt to changing environments. This research will characterize short-term plasticity in order to realize adaptive and low powered AI. It will impact machine learning and AI capabilities in future Exascale systems and embedded national security systems.

Project Description:

Despite the successes of modern day artificial neural networks (ANNs), there are severe drawbacks. They are power hungry, and a great deal of labeled data or simulation time is required for training, making real world problems computationally expensive or infeasible. Furthermore, deep neural networks (DNNs) implementing supervised learning need to be retrained when the statistics of the input change. For a more detailed description of the deficiencies of modern artificial intelligence (AI), please see the "Achieve efficient learning for AI systems" section on page 101 of the DOE, *AI for Science report*. The human brain provides proof that a network can learn using limited amounts of data, transfer knowledge between different tasks, and easily adapt to changing environments: all while consuming extremely low power.

There are many principles ubiquitously found across the brain that are not yet implemented in modern ANNs; it is likely that their implementation will yield more capable and low-powered AI useful for national security applications. However, as more brain-like complexity is introduced, more parameters will be necessary in the implementation. Finding the right combination of parameters that yield high performance will be challenging, if not impossible. Here, again we can turn to the brain for guidance. Characterizing how the brain implements phenomena will help limit the parameter space.

This project focuses on one brain phenomena, short term plasticity (STP), and aim to help create a foundation for implementation in ANNs. STP has the capability to enable many computations (Abbott and Regehr, 2004; Mongillo et al., 2008; Buonomano and Maass, 2009). Notably, recent work by Leng and colleagues (2018), have shown that spiking networks with STP can simultaneously perform well at classification and become good generative models. I aim to characterize STP using a large data set collected by the Allen Institute for Brain Science (AIBS; https://alleninstitute.org/what-we-do/brain-science/). This characterization of STP will help provide a foundation for realistic implementations in ANNs. In addition, it is likely that this characterization will contribute to a pipeline manuscript currently being assembled at AIBS describing their methods and findings. In the event STP characterization and/or contribution to the manuscript is not possible, I will use current knowledge of STP to extend the work of Leng and colleagues (2018).

Previous Accomplishments:

Recent implementations of short term plasticity (STP) mechanisms in artificial neural networks (ANNs) are starting to reveal their benefits in engineered systems. Here, I aim to use biological

data provided by the Allen Institute for Brain Science (AIBS) to explore possible characterizations of STP so they can be used to guide and constrain STP implementations in ANNs.

Accomplishments:

Synaptic physiology database and software package to access data have been acquired from the Allen Institute for Brain Science (AIBS).
Tools and data are up and running on local computer.
At the end of this fiscal year I will have demonstrated:

If there is low dimensional embedding provided by methods such as principle component analysis (PCA), autoencoders, t-distributed stochastic embedding (t-SNE), etc., that adequately represents the high dimensional data.
How well supervised learning methods such as random forests, support vector machines can differentiate preconceived biological characterizations.
If unsupervised learning methods such as k-means, gaussian mixtures, affinity propagation, and hierarchical clustering will reveal preconceived biological characterizations or expose new classification paradigms.
It is possible that the space of features is continuous which is also important knowledge for constraining network parameters.

Unfortunately, I believe that parameters for an STP model fit to raw data will be most useful for classification and ANN implementations. Someone at the AIBS is actively working on this, but they have not yet achieved this feat. If it is achieved, the code above will be ready to run on resulting parameters.

Summary of FY21 Project Plan: FY21 Project Plan:

Discuss output of analysis with the Allen Institute for Brain Science and determine if it will go into their pipeline publication.

If it looks promising, refine analysis and figures.

If no, start on ANN network implementation below. SAND report

In the event plus-up money becomes available:

In order to utilize STP in ANNs for national security applications, it would be ideal to have an in-house implementation. I plan to extend the work of Leng Et Al, 2018 by implementing biologically constrained STP.
Implementation and benchmarking of above ANN on new Loihi neuromorphic chip.

Tie to Mission:

Recent advances in Artificial intelligence (AI) have transformed the world and national security. Current AI has power and training limitations; however, technology is quickly evolving. Biologically inspired algorithms will likely be the next breakthrough that enables low-power, next generation scientific HPC platforms for national security programs. This LDRD will impact DOE, NNSA, DHS, ASC, CSSE, and ATDM programs by contributing to new capabilities in machine learning and emerging/exascale computing programs. It will help provide new low-power computing and communication capabilities for embedded sensors and autonomous

machines and help enable a new generation of ‘smart’ transportation, energy production and grid resiliency.

DISTRIBUTION

Email—Internal

Name	Org.	Sandia Email Address
Technical Library	01977	sanddocs@sandia.gov

This page left blank

This page left blank

Sandia National Laboratories

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia LLC, a wholly owned subsidiary of Honeywell International Inc. for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.