

Learning Spiking Neural Network Models of *Drosophila* Olfaction

John Carter*
Department of Computer Science
Drexel University
jmc683@drexel.edu

Jocelyn Rego*
Department of Computer Science
Drexel University
jr3548@drexel.edu

Daniel Schwartz*
Department of Computer Science
Drexel University
des338@drexel.edu

Vikas Bhandawat
Department of Biomedical Engineering
Drexel University
vikas.bhandawat@drexel.edu

Edward Kim
Department of Computer Science
Drexel University
ek826@drexel.edu

ABSTRACT

We present research in the modeling of neurons within *Drosophila* (fruit fly) olfaction. We describe the process from data collection, to model creation, and spike generation. Our approach utilizes computational elements such as spiking neural networks that employ leaky integrate-and-fire neurons with adaptive firing behavior that more closely mimick biological neurons. We describe the methods of several learning implementations in both software and hardware. Finally, we present both quantitative and qualitative results on learning spiking neural network models.

CCS CONCEPTS

• Computing methodologies → Machine learning algorithms; Modeling methodologies; • Hardware → Neural systems.

KEYWORDS

neuro-inspired artificial intelligence, machine learning, spiking neural networks, olfaction, neuromorphic computing

ACM Reference Format:

John Carter, Jocelyn Rego, Daniel Schwartz, Vikas Bhandawat, and Edward Kim. 2020. Learning Spiking Neural Network Models of *Drosophila* Olfaction. In *International Conference on Neuromorphic Systems 2020 (ICONS 2020)*, July 28–30, 2020, Oak Ridge, TN, USA. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3407197.3407214>

1 INTRODUCTION

In the past several decades, we have seen tremendous progress in understanding the connectivity between neurons in the brain, in developing methods to record from neurons in awake behaving animals, and in perturbing neural circuits to assess brain function. Despite this progress, we do not understand the ability of the

*Authors contributed equally to this research.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ICONS 2020, July 28–30, 2020, Oak Ridge, TN, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8851-1/20/07...\$15.00

<https://doi.org/10.1145/3407197.3407214>

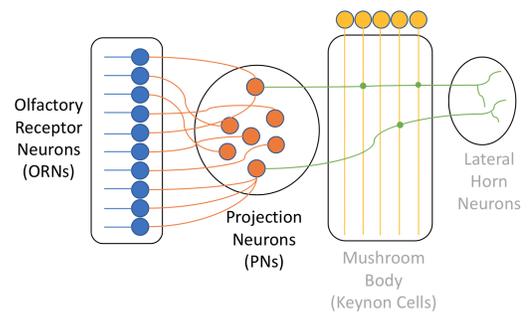


Figure 1: Illustration of the *Drosophila* Olfactory system. We focus on recorded activity from the first and second order neurons, i.e. Olfactory Receptor Neurons (ORNs) and the Projection Neurons (PNs), after neuronal stimulation. We then learn the synaptic weights of the connections between the first and second order neurons through various statistical and neuromorphic approaches.

nervous system to learn the structure of the world as the animal observes and interacts with it. Without this knowledge, we are not able to explain or replicate an animal’s ability to perform a large suite of behaviors, ability to learn new behaviors, and ability to adapt on the fly.

In an effort to advance knowledge in the mechanisms of low level perception, we investigate and model neurons of *Drosophila* (fruit fly) olfaction, see Figure 1. We observed the neural architecture of these circuits and aimed to replicate an in-silico architecture that resembles these circuits. We utilized biologically plausible computational elements such as spiking neural networks that employ leaky integrate-and-fire neurons with adaptive firing behavior that more closely mimicked the biological neurons. We describe several learning implementations of first and second order neurons to infer representations that support subsequent processing tasks. Since our data inputs and outputs are measured and accurately defined, our networks can be trained using high-fidelity data.

2 BACKGROUND

Our research employs odor-tracking in *Drosophila* (fruit flies) as a model to understand brain function. Odors are dispersed by the turbulent flow of air; the spatial and temporal distribution of odor is

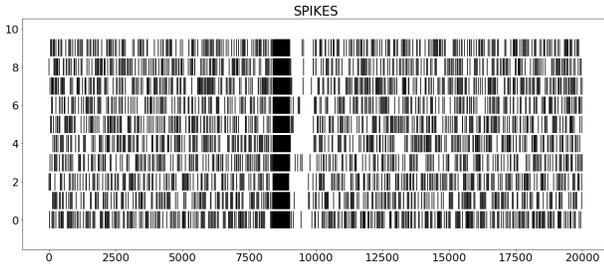


Figure 2: 10 spike trains collected from the ORN neurons. The recorded time is 20 seconds, with the onset of the stimulus at 8 seconds lasting for 500 ms.

highly dynamic and is therefore unpredictable. This task represents exactly the kind of task that animals excel at, and the exact process and mechanisms are not well understood. Furthermore, although modern day artificial intelligence and machine learning have neuro-inspired foundations, much of the functions of biological neurons have been oversimplified [8]. We investigate the anatomy and physiology of the olfactory system in *Drosophila* for several reasons - it is at a scale that can be modeled in computational models, and the mechanisms and architecture are observable at the neuronal level. Genetic tools are already in place to both assess neural function and to perturb circuits in the brain. Thus, the overall goal is to understand the basis of odor-tracking in the fly brain, and rigorously test this understanding by creating a realistic in-silico model.

Our research laboratory has built a novel apparatus in which flies are exploring a circular arena 0.5 meter in diameter. Odors were delivered with an olfactometer producing a continuous stream of air directed over the fly. Ultimately, we plan to perform in vivo imaging and physiology from three olfactory circuits in the fly brain – antennal lobe, mushroom body and lateral horn to assess the role of these circuits in naturalistic behavior. Because the flies are being tracked in real-time, we plan to use optogenetics in the future (a method whereby light is employed to activate a genetically defined population), whereby their olfactory neurons can be stimulated in arbitrary spatial and temporal patterns based on real-world measurement of odor dispersion, and we can assess how a fly adapts its behavior. For the purposes of this research experiment, we focus on modeling the first and second order neurons (ORN and PN) on neuromorphic software and hardware. We collect and show 10 example ORN recordings in Figure 2, and 10 PN responses.

3 METHODOLOGY

In this section, we briefly describe the methods on the data collection process from the *Drosophila* brain. Next, we will describe a series of neuronal models that we created, that gradually increase in function and complexity. We delineate the process and parameters used to create these models in neuromorphic software and hardware platforms. Finally, we describe the experimentation procedure, and report both quantitative and qualitative results on modeling the behavior of the first and second order neurons.

3.1 Data Collection

Olfactory Receptor Neurons (ORN) Recordings - Flies were immobilized in the trimmed end of a plastic pipette tip. A reference

electrode filled with *Drosophila* saline was inserted into the eye, and a sharp saline-filled glass capillary (tip diameter < 1µm) was inserted into a sensillum. Sensilla were visualized using an Olympus BX51WI microscope with a 50x air objective. Sensillum types were identified based on their morphology and their characteristic responses to a panel of odors [5, 6]. Voltage signals were acquired with an A-M Systems Model 2400 amplifier. Signals were low-pass filtered at 2 kHz and digitized at 10 kHz and the final ORN spikes were detected off-line using routines in IgorPro (Wavemetrics).

Projection Neurons (PN) Recordings - Whole-cell recordings from PNs were performed in vivo as previously described [3, 13, 14], using an internal patch-pipette solution. Voltage recordings were obtained with an A-M Systems Model 2400 amplifier in the current clamp mode and signals were low-pass filtered at 5 kHz and digitized at 10 kHz. An Olympus BX51WI microscope with a 40x water-immersion objective, IR-DIC optics, and a fluorescence attachment was used to obtain recordings under visual control. One neuron was recorded per brain, and the morphology of each cell was visualized post hoc with biocytin histochemistry.

Olfactory Stimulation - Odors were diluted in paraffin oil at a ratio of 1:100. Odors were delivered with a custom-built olfactometer. A continuous stream of charcoal-filtered air (2.2 l/min) was directed over the fly. Switching of a three-way solenoid redirected 200 ml/min of this air through an odor vial, which rejoined the air stream 12 cm from the end of the odor tube. Thus, all odors were diluted 10-fold in air just before reaching the fly. The odor tube was 8 mm in diameter and terminated 8 mm from the fly.

3.2 Rate Coded Models

It is known that learning the parameters of a neural network model using spiking data is a difficult endeavor. The main reasons for this are that spike activations are non-differentiable, and occur with sparse frequency. Thus, our first, more simplistic models are based upon learning on rate codes instead of on the spike train data directly. The data analyzed in our research consisted of spike times of ORNs and PNs in response to two different odors. Spikes were recorded from ORNs and PNs over ten 20-second trials. Odor was presented at the 8 second mark. First, we transformed this data, which was initially represented by timestamps of spikes, into bins of spike counts. We summed the number of spikes in bins of 50 milliseconds, with an overlap of 25 milliseconds. This resulted in each trial being represented as an array of 800 spike counts.

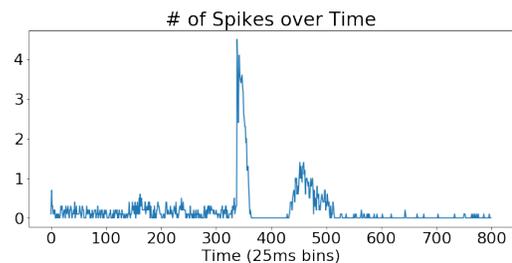


Figure 3: Histogram showing the average rate code (number of spikes) in projection neurons (PN) over 20 seconds, binned into 800 time segments.

Shown in Figure 3 is a visualization of the average PN activity, with the number of spikes indicated by the y-axis, and the time bins indicated by the x-axis. The largest number of spikes appear soon after odor onset, at about the 8-second mark, as shown around the 350th bin.

Gaussian Generalized Linear Model - A Generalized Linear Model (GLM) is a class of linear regression models that contain response variables that can have different error distribution models. A GLM is composed of a linear predictor, a link function that describes the relationship between the mean and the linear predictor, and a distribution that fits the dataset. The Generalized Linear Model has been shown to be an effective tool for modeling statistical relationships in data [12]. GLMs have been particularly successful in modeling the activity of a neuron’s reaction to a sensory stimulus. Furthermore, a GLM can represent a variety of factors of a spike train to probabilistically model spiking activity [7].

Using the rate coded data, we modeled our data using a Gaussian GLM (this is identical to the ordinary least squares linear regression model). This linear regression model mapped ten ORNs to each of our ten PNs. The coefficients from these ten models were then averaged together, which resulted in ten coefficients representing the average impact of each ORN on the activity of a PN.

Poisson Generalized Linear Model - Poisson GLMs resolve the flaws of general linear models where the range of the model is restricted and the variance depends on the mean. The Poisson distribution is a commonly used statistical distribution and is of the exponential family. A Poisson distribution can be used within a GLM and is often referred to as a Poisson regression model. The linear predictor of a Poisson regression model is simply just the linear combination of an explanatory variable and a bias term. The link function within this type of model is the canonical logit link function. Lastly, the probability distribution that generates the observed variable in this type of model is appropriately the Poisson distribution, which has a single parameter to represent both the mean and standard deviation of the distribution. In a study analyzing how olfactory cells affect coding of stimuli, GLMs modeled olfactory cells activity better than a simpler model that did not include spike history effects [11]. Additionally, the work presented that for all neurons the GLM fits were better than the linear-nonlinear-Poisson (LNP) models [11].

For this model, we used Poisson regression with a log link function to model the effect of the binned spike counts of ten ORNs on the spike counts of each PN. We constructed a GLM for each of the 10 PN trials, mapping the ten ORNs to each PN. We then averaged the resulting coefficients of the models to predict the spike counts of a PN. The Poisson GLM was mostly successful in predicting the actual values for the PNs, especially in anticipating the large increase in spikes around bin 350.

3.3 Spiking Neural Network Models

In this section, we move away from the rate coded models and learn directly on the spike input. We describe two neuromorphic implementations that utilize local learning rules to inform the synapse connections between the ORN and PN neurons.

Neural Engineering Framework - In the field of neuroscience, the Neural Engineering Framework (NEF) is a set of principles that

can be used to build large-scale networks of single-neuron models which enforces biological constraints, which are detailed in [1].

Nengo is a software module that is used to simulate models based on the NEF principles. We used a Nengo model to learn according to a biologically plausible supervised learning rule frequently used with the Neural Engineering Framework (NEF) called Prescribed Error Sensitivity (PES). The PES rule adjusts connection weights between populations of neurons to minimize an external error signal over time. The PES learning rule on decoders is represented in equation 1 where \mathbf{e}_j are the encoders of the “post” population, κ is a scalar learning rate, n is the number of neurons in the “pre” population, and a_j is the gain of the “post” population.

Nengo provides functionality to developers to configure each neuron in their networks to be sensitive to a direction in the n -dimensional space and can be parameterized by a floating-point number. In addition, the encoder aids in the PES learning rule analogous to backpropagation such that an individual synaptic connection weight is adjusted by mapping a global error signal to its local error signal. This mapping is computed by imposing the portion of the error vector space via its encoder [2].

$$\Delta\omega_{ij} = -\frac{\kappa}{n}\alpha_j\mathbf{e}_j \cdot \mathbf{E}a_i \quad (1)$$

We initialize ten Nengo Nodes and present the spike trains as input. Next, the ten Nodes representing the ORNs are all connected to a single Adaptive LIF neuron, which represents the PN. This model attempts to learn the connection weights of the ten ORNs to the single PN via PES and its error is measured by the difference of the simulated PN and the observed PN. All neurons are configured such that their encoders are set to 1 and their intercepts or the point along each neuron where its activity is zero is set to 0.5. Additionally, we used Nengo’s Adaptive spiking version of the Leaky Integrate-and-Fire (LIF) neuron model that works similar to the LIF model, but has an adaptation time constant of $\tau_{adp} = 1$ and an adaptation increment of $inc_n = 0.5$, which is how much the adaptation state is increased after each spike. Lastly, we define the presentation time of the input to be 0.001 seconds and the synapse used to filter the connections with a 0.1 second delay.

Loihi Hardware Model - Intel’s Loihi is the latest chip released for research in 2018 [4]. Loihi is a manycore spiking neural network, where a single chip contains 128 neuromorphic cores, 128k neurons, 128M synapses. The chip is fully asynchronous, fully digital, and deterministic. At any given time, the neurons may send out a spike to its neighbors by use of virtual synapses. Neurons have two internal state variables - a synaptic response current, a weighted sum of the input spikes and a constant bias, and membrane potential that leaks over time and will send out a spike when the potential crosses the firing threshold. Once enough spikes accumulate and exceed some threshold, a spike message is generated and sent to other connected neurons. This chip is programmable using the neuromorphic software development kit (Nx SDK) for use in hardware accelerated, low power neuronal applications [9].

Our network is built using three primary objects: compartments, spike generators, and learning rules. Compartments in Loihi are the basic building blocks of the spiking neural networks. The compartment voltage integrates the bias current and when it exceeds a specified membrane potential, it produces a spike. The voltage

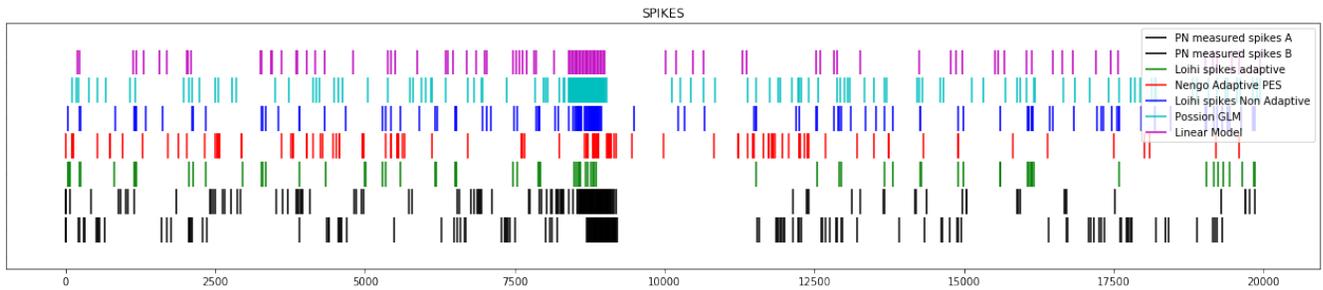


Figure 4: Spike trains of a simulated PN generated from all different models using 10 observed ORN as inputs to the system. The spikes shown from each of the methods are randomly generated from the parameters of the model.

then resets and the integration process begins again. We can also specify that the voltage decay over some time constant to simulate the behavior of a leaky integrate and fire neuron (LIF). For the time being, we only create excitatory connections within the model.

Input spikes are injected into the system using built-in spike generators. The spike generator is an object that can generate spikes by some algorithm, through a sensor, or in our case, spike trains extracted from the *Drosophila* ORNs. We create a more adaptive version of the leaky integrate and fire neuron by modulating the activity range homeostasis. Thus, the membrane threshold of the compartment becomes a function of its firing activity, see Figure 5. As the membrane activity increases over a specified max activity level, the membrane threshold will adapt to prevent the compartment from firing.

The synapse weights connecting the ORN to PN neurons are learned using a pairwise STDP rule of the following form.

$$W(t + 1) = W(t) - A_-x_0(t)y_1(t) + A_+x_1(t)y_0(t) \quad (2)$$

Where the weight, W , at time= $t+1$ is updated by the x and y , pre and post synaptic spikes, respectively. This model is able to reproduce key features of long term potentiation and depression.

4 EXPERIMENTS AND RESULTS

For our experiments, we modeled the connection from 10 ORN neurons to a single PN neuron. We learned the weights of these connections using the different models as described in our methodology. Given the learned weights, we then simulated spike traces from the models and compared these spike traces to actual PN spike trains collected from the *Drosophila* cells.

In Figure 4, we show an example spike train generated from all of the different methods. We also show two trials of PN data collected from the fruit fly. These two ground truth spike trains illustrate the variability that exists within real world recordings of the same stimulus. In order to quantify these results, we use a measure of spike synchrony called the inter-spike interval metric [10]. The ISI distance quantifies information from the interspike intervals by evaluating the ratio of the instantaneous firing rates. We normalize the value between 0-1, where 0 means perfectly synchronous. We show the ISI measure between spike trains in Table 1. As a baseline, we show the average ISI between all PNs in the first row, and of two ground truth PNs in the second row.

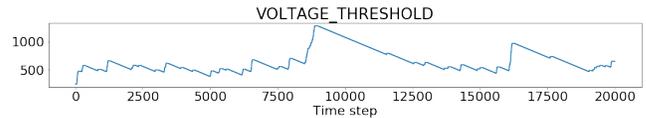


Figure 5: Adaptive voltage threshold of the Loihi model. The rapid spiking behavior of a neuron increases the voltage threshold, inhibiting the neuron from firing for a period of time.

Method	0-8s	8-9s	9-11s	11-20s
Average of All PNs	0.656	0.464	0.122	0.123
Observed PN A to B	0.675	0.521	0.104	0.110
Linear Regression	0.678	0.663	0.640	0.767
Poisson GLM	0.703	0.672	0.633	0.853
Nengo Adaptive LIF	0.612	0.670	0.704	0.776
Loihi Non-Adaptive	0.700	0.686	0.703	0.879
Loihi Adaptive	0.623	0.642	0.352	0.796

Table 1: Spike ISI similarity between methods at different times during the trials. Lower values indicate better synchrony.

5 CONCLUSION

In conclusion, we investigated several learning implementations of gradually increasing biological plausibility for modeling *Drosophila* (fruit fly) olfaction ORN and PN neurons. We present both quantitative and qualitative results on learning rate coded and spiking neural network models. Our data shows that most methods model the spike train well during the pre-stimulus (0-8 seconds) time, and during the presentation of the stimulus (8-9 seconds). The adaptive Loihi implementation greatly outperforms the other methods during the time of adaptation (9-11 seconds) due to activity range homeostasis.

ACKNOWLEDGMENTS

This material is based upon work partially supported by the Intel Corporation and the National Science Foundation under Grant No. 1954364.

REFERENCES

- [1] Trevor Bekolay. 2011. Learning in large-scale spiking neural networks.
- [2] Trevor Bekolay, Carter Kolbeck, and Chris Eliasmith. 2013. Simultaneous unsupervised and supervised learning of cognitive functions in biologically plausible spiking neural networks. In *CogSci*.
- [3] Vikas Bhandawat, Shawn R Olsen, Nathan W Gouwens, Michelle L Schlief, and Rachel I Wilson. 2007. Sensory processing in the Drosophila antennal lobe increases reliability and separability of ensemble odor representations. *Nature neuroscience* 10, 11 (2007), 1474–1482.
- [4] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, et al. 2018. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38, 1 (2018), 82–99.
- [5] Marien De Bruyne, Peter J Clyne, and John R Carlson. 1999. Odor coding in a model olfactory organ: The Drosophila maxillary palp. *Journal of Neuroscience* 19, 11 (1999), 4520–4532.
- [6] Marien De Bruyne, Kara Foster, and John R Carlson. 2001. Odor coding in the Drosophila antenna. *Neuron* 30, 2 (2001), 537–552.
- [7] Robert E Kass and Valérie Ventura. 2001. A spike-train probability model. *Neural computation* 13, 8 (2001), 1713–1720.
- [8] Edward Kim, Darryl Hannan, and Garrett Kenyon. 2018. Deep sparse coding for invariant multimodal halle berry neurons. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 1111–1120.
- [9] Edward Kim, Jessica Yarnall, Priya Shah, and Garrett T Kenyon. 2019. A Neuromorphic Sparse Coding Defense to Adversarial Images. In *Proceedings of the International Conference on Neuromorphic Systems*. 1–8.
- [10] Thomas Kreuz, Julie S Haas, Alice Morelli, Henry DI Abarbanel, and Antonio Politi. 2007. Measuring spike train synchrony. *Journal of neuroscience methods* 165, 1 (2007), 151–161.
- [11] Shreejoy J Tripathy, Krishnan Padmanabhan, Richard C Gerkin, and Nathaniel N Urban. 2013. Intermediate intrinsic diversity enhances neural population coding. *Proceedings of the National Academy of Sciences* 110, 20 (2013), 8248–8253.
- [12] Alison I Weber and Jonathan W Pillow. 2017. Capturing the dynamical repertoire of single neurons with generalized linear models. *Neural computation* 29, 12 (2017), 3260–3289.
- [13] Rachel I Wilson and Gilles Laurent. 2005. Role of GABAergic inhibition in shaping odor-evoked spatiotemporal patterns in the Drosophila antennal lobe. *Journal of Neuroscience* 25, 40 (2005), 9069–9079.
- [14] Rachel I Wilson, Glenn C Turner, and Gilles Laurent. 2004. Transformation of olfactory representations in the Drosophila antennal lobe. *Science* 303, 5656 (2004), 366–370.