Inferring Network Structure from Neural Activity in a Biologically-Constrained Model of the Insect Olfactory System

Shruti Joshi (s4joshi@ucsd.edu)
Department of Electrical and Computer Engineering,
University of California San Diego,
La Jolla, CA

M. Gabriela Navas-Zuloaga (mnavaszuloaga@ucsd.edu)
Department of Medicine,
University of California San Diego,
La Jolla, CA

Autumn McLane-Svoboda (mclanesv@msu.edu)
Department of Biomedical Engineering,
Institute for Quantitative Health Science and Engineering, Michigan State University,
East Lansing, MI

Simon Sanchez (sanch364@msu.edu)
Department of Biomedical Engineering,
Institute for Quantitative Health Science and Engineering, Michigan State University,
East Lansing, MI

Debajit Saha (sahadeb3@msu.edu)
Department of Biomedical Engineering,
Institute for Quantitative Health Science and Engineering, Michigan State University,
East Lansing, MI

Maxim Bazhenov (mbazhenov@ucsd.edu)
Department of Medicine,
University of California San Diego,
La Jolla, CA

Abstract

Understanding olfactory processing in insects requires characterizing the complex dynamics and connectivity of the first olfactory relay - antennal lobe (AL). We leverage *in vivo* electrophysiology to train recurrent neural network (RNN) model of the locust AL, inferring the underlying connectivity and temporal dynamics. The RNN comprises 830 projection neurons (PNs) and 300 local neurons (LNs), replicating the locust AL anatomy. The trained network reveals sparse connectivity, with different connection densities between LNs and PNs and no PN-PN connections, consistent with *in vivo* data. The learned time constants predict slower LN dynamics and diverse PN response patterns, with low and high time constants correlating with early and late odor-evoked activity, as reported in vivo. Our approach demonstrates the utility of biologically-constrained RNNs in inferring circuit properties from empirical data, providing insights into mechanisms of odor coding in the AL.

Keywords: recurrent neural network; olfactory system; synaptic connectivity

Introduction

Olfactory processing in insects relies on the intricate circuitry of the antennal lobe (AL) - the first relay of the olfactory system, where excitatory PNs and inhibitory LNs interact to shape odor representations. The AL exhibits complex dynamics, with PNs and LNs displaying diverse and odor-specific temporal response patterns (Laurent, 2002). Understanding these dynamics is crucial for unraveling the mechanisms underlying odor coding in insects.

Many existing computational models of the AL rely on hand-tuned parameters to recreate patterns of activity observed *in vivo* (Patel et al., 2009; Bazhenov et al., 2013, 2001; Chen et al., 2015). While these models provided many valuable insights, they do not fully leverage the wealth of available *in vivo* data to construct the models. Machine learning approaches, such as recurrent neural networks (RNNs), offer a powerful alternative to investigate the AL's dynamics and generate testable hypotheses. RNNs can learn from large sets of empiri-

cal data and capture complex temporal dependencies, making them well-suited for modeling the intricate dynamics of the AL.

In this study, we construct a continuous rate RNN model of the locust AL to capture the temporal patterns and diverse responses of the AL PNs and LNs. Our approach allows to train the model parameters based on empirical data, enabling us to uncover the key mechanisms that shape odor representations in the AL. Through this data-driven modeling framework, we seek to advance our understanding of the computational principles governing olfactory processing in insects and generate testable predictions for future experimental studies.

Methods

Experimental data Electrophysiology recordings were obtained from the AL of 12 locusts exposed to 5 odors. The data were processed to extract spike times of 69 PNs. The instantaneous firing rates of these PNs were calculated by binning the spiking data into 100ms intervals and convolving with Gaussian kernel.

Recurrent Neural Network Model We constructed a continuous rate RNN model of the AL based on (Kim et al., 2019). The model (Fig. 1) comprises 830 PNs and 300 LNs, same as the numbers found in the locust AL (Laurent, 2002). The temporal dynamics of the synaptic current for each RNN unit are governed by the following equation:

$$\tau^d \frac{dx}{dt} = -x + W_{rec} r_{rate} + W_{in} u \;\; ; \;\; r_{rate} = \phi(x) \quad (1)$$

where τ^d represents the membrane time constant, x denotes the synaptic currents, W_{rec} signifies the recurrent connectivity, W_{in} indicates the input weight between olfactory receptor neurons (ORNs) and AL neurons, u corresponds to the input from the ORNs, and ϕ is a non-linear transfer function that determines the firing rate.

Among the 820 PNs, 69 were trained to replicate the firing rates of the 69 recorded *in vivo* PNs. The model was trained using backpropagation through time (BPTT) (Mozer, 1995) to minimize the root mean squared error (rMSE) loss between the firing rates of the RNN units and their corresponding *in vivo* PNs. During training, the input weights W_{in}, recurrent weights W_{rec}, and time constants τ^d were optimized. To maintain biological plausibility, W_{in} was constrained to be positive, W_{rec} was masked to have separate excitatory and inhibitory connections (Song et al., 2016) according to Dale's principle (Eccles et al., 1954), and the values for τ^d were restricted between 10ms and 1000ms.

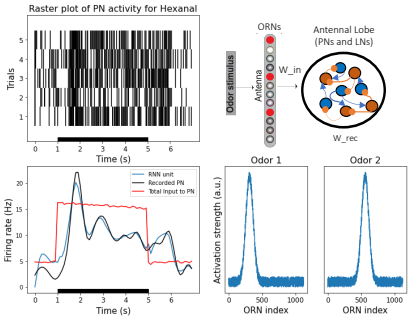

Figure 1: RNN model of the locust AL. (Left) Example *in vivo* PN spike train over 5 trials (top) and corresponding firing rates of the biological PN (black) and trained model unit (blue). Black bar indicates odor stimulation. (Right) Schematic of the RNN model structure and input patterns for two odors

Results

Neural activity The trained RNN units closely match the recorded activity of their corresponding PNs (Fig. 1). Moreover, the remaining RNN units exhibit distinct temporal dynamics, with some responding at the onset of odors, some at the offset, and a few at both onset and offset (Fig. 3), matching experimentally observed patterns (Saha et al., 2017).

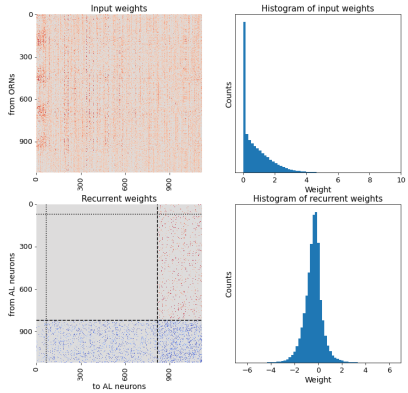

Figure 2: Synaptic weights. (Top) Trained RNN input weight adjacency matrix. (Bottom) Trained RNN recurrent weights adjacency matrix. Neurons before the dashed line are the PNs and after are LNs. Weights in blue are negative and red are positive.

Trained synaptic connectivity The initial connection probabilities were set to 0.3 for PN-LN, LN-LN, and LN-PN connections, with PN-PN connections constrained to 0 throughout training. After training, the connection probabilities were found to be 0.1 for PN-LN, 0.23 for LN-PN, and 0.28 for PN-LN (Fig. 2). Even though there is no direct estimate of connection probabilities from experiments, this is consistent with values predicted by

modeling studies (Patel et al., 2009). Importantly, when non-zero PN-PN connections were allowed, the training eliminated the vast majority of them, took longer to converge, and had a worse fit to data. This matches the experimental observation of absence of PN-PN connections (Wilson, 2011).

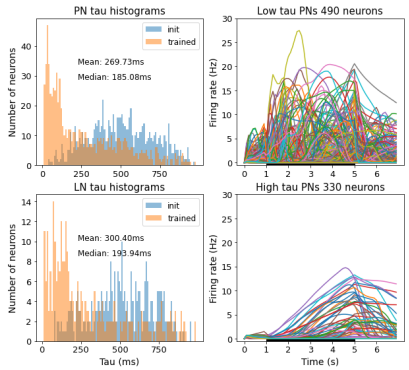

Figure 3: Time constants. (Left) Trained time constants for LNs and PNs. (Right) PN activities split by their time constants. Top shows activities of PNs with low time constants and bottom shows activities of PNs with high time constants.

Trained time constants After random initialization and training, the time constants for both PNs and LNs exhibited a distribution with mostly low values and a small number of units with high values (Fig. 3). The median time constant for LNs was approximately 193ms (mode=70ms), while for PNs, it was about 185ms (mode=30ms). The higher time constants of LNs are consistent with experimental observations, partially attributed to LNs generating slower Ca^{2+} spikelets instead of Na^{+} spikes in locusts (Laurent et al., 1993).

Splitting PNs according to their time constants revealed markedly different activity patterns. Most PNs responding at an odor onset belonged to the low time constant group, while those responding with a delay or at odor offset primarily had high time constants (Fig. 3). The locust AL includes fast and slow inhibitory synapses between LNs and PNs (MacLeod and Laurent, 1996). As our model does not explicitly include synaptic time constants, the unit time constants integrate these synaptic time scales, potentially predicting the presence of units with high time constants (400-500ms) that could account for the slow inhibitory synapses.

Conclusion

We developed a continuous rate RNN model of the locust AL to study olfactory processing dynamics in insects. The model captured the temporal patterns and diverse responses of PNs and LNs. Analysis of the trained weights and time constants provided insights into the functional roles of neuron types and synaptic connectivity in the AL, demonstrating the effectiveness of using RNNs to investigate biological system dynamics.

Acknowledgments

This work was supported by grants from the National Institutes of Health (1R01DC020892) and the National Science Foundation (2323241, GRFP).

References

Bazhenov, M., Huerta, R., and Smith, B. H. (2013). A Computational Framework for Understanding Decision Making through Integration of Basic Learning Rules. *The Journal of Neuroscience*, 33(13):5686–5697.

Bazhenov, M., Stopfer, M., Rabinovich, M. I., Huerta, R., Abarbanel, H. D. I., Sejnowski, T. J., Gilles Laurent, and Laurent, G. (2001). Model of Transient Oscillatory Synchronization in the Locust Antennal Lobe. *Neuron*, 30(2):553–567. MAG ID: 2148452798.

Chen, J.-Y., Marachlian, E., Assisi, C., Huerta, R., Smith, B. H., Locatelli, F., and Bazhenov, M. (2015). Learning Modifies Odor Mixture Processing to Improve Detection of Relevant Components. *Journal of Neuroscience*, 35(1):179–197.

Eccles, J. C., Fatt, P., and Koketsu, K. (1954). Cholinergic and inhibitory synapses in a pathway from motor-axon collaterals to motoneurones. *The Journal of Physiology*, 126(3):524–562.

Kim, R., Li, Y., and Sejnowski, T. J. (2019). Simple framework for constructing functional spiking recurrent neural networks. *Proceedings of the National Academy of Sciences*, 116(45):22811–22820. Publisher: Proceedings of the National Academy of Sciences.

Laurent, G. (2002). Olfactory network dynamics and the coding of multidimensional signals. *Nature Reviews Neuroscience*, 3(11):884–895. Number: 11 Publisher: Nature Publishing Group.

Laurent, G., Seymour-Laurent, K. J., and Johnson, K. (1993). Dendritic excitability and a voltage-gated calcium current in locust nonspiking local interneurons. *Journal of Neurophysiology*, 69(5):1484–1498.

MacLeod, K. and Laurent, G. (1996). Distinct Mechanisms for Synchronization and Temporal Patterning of Odor-Encoding Neural Assemblies. *Science*, 274(5289):976–979.

Mozer, M. (1995). A focused backpropagation algorithm for temporal pattern recognition. *Complex Systems*, 3.

Patel, M., Rangan, A. V., and Cai, D. (2009). A large-scale model of the locust antennal lobe. *Journal of Computational Neuroscience*, 27(3):553–567.

Saha, D., Sun, W., Li, C., Nizampatnam, S., Padovano, W., Chen, Z., Chen, A., Altan, E., Lo, R., Barbour, D. L., and Raman, B. (2017). Engaging and disengaging recurrent inhibition coincides with sensing and unsensing of a sensory stimulus. *Nature Communications*, 8(1):15413.

Song, H. F., Yang, G. R., and Wang, X.-J. (2016). Training Excitatory-Inhibitory Recurrent Neural Networks for Cognitive Tasks: A Simple and Flexible Framework. *PLOS Computational Biology*, 12(2):e1004792. Publisher: Public Library of Science.

Wilson, R. I. (2011). Understanding the functional consequences of synaptic specialization: insight from the Drosophila antennal lobe. *Current Opinion in Neurobiology*, 21(2):254–260.