Recurrent Neural Networks: a promising approach to predict spike times of a single neuron in comparison to biophysical models

submitted to

Prof. Dr. Ulf Leser, Prof. Dr. Susanne Schreiber

Humboldt-Universität zu Berlin

Institute for Computer Science, Institute for Theoretical Biology

by

Mark-Dominik Schatz (545097)

Abstract

As large-scale, complex network modeling projects are flourishing in the field of computational neuroscience [1], it is more and more important to design single-neuron models that capture qualitative features of real neurons and are quantitatively accurate in silico representations of those. Recent years have seen substantial effort being put into the development of algorithms for the systematic evaluation and optimization of biophysical neuron models with respect to electrophysiological data [2,3]. In this bachelor thesis, we will train a Recurrent Neural Net (RNN) on neuron recordings using methods such as Sparse Teacher Forcing and compare its predictive power to that of other neuron models provided by the community of the single-neuron modeling competition[4]. We will use the benchmarks to try to outperform the given best performing methods of the biological neuron models.

1. Motivation

1.1 Subject of the study

Recent years have seen many large-scale network modeling projects in computational neuroscience [1]. These projects require the design of single-neuron models that can accurately capture real neuron activity. Single-neuron models can be divided into two categories: detailed biophysical models (Hodgkin and Huxley 1952, HH-type models and variants thereof [5]), and simple phenomenological models [26]. The Quantitative Single-Neuron Modeling Competition [4] focuses on constructing and studying biophysically realistic models of single-neuron electrical activity, which can be used to reproduce a variety of neuronal behaviors.

This year's competition focuses on developing and studying biophysically realistic models of single-neuron electrical activity, which can be used to reproduce a variety of neuronal behaviors. The competition is divided into four individual challenges, focusing on single-compartment and multi-compartmental models, where each challenge has its own dataset and evaluation criteria. The competition allows for the evaluation of models on a standardized set of tests, giving insight into how well these models can predict neuron behavior. We focus on the first Challenge A, which will be to predict the timing of somatic action potentials or to predict the neuronal response upon stimulation by ramp and step currents [16].

Spikes or action potentials are prominent features in most brain neurons, through which neural information can be efficiently processed and reliably transmitted, and eventually form the basis underlying how the brain works [19]. Previous studies have proposed many computational models to characterize the firing dynamics of neurons, among which HH-type models are the most famous [25]. To ensure the accuracy of spike timing calculation, the step size on which one numerically iterate has to be very small, such as 0.01 ms. Therefore, the simulation cost of the HH-type model is very high, especially when modeling large-scale brain networks.

Time series prediction problems are a difficult type of predictive modeling problem. In contrast to high dimensional data, time series add the complexity of the temporal sequence of dependence among the input variable. Due to the good mapping ability of
an artificial neural network (ANN), we want to propose an ANN-based method.
To predict the occurrence of spikes and their times by using only the data provided by the competition, we use a type of neural network designed to handle sequence dependence called recurrent neural network (RNN) [33].
We hope to drastically reduce the simulation cost and compare those results with the given best model at the current time.

As of our knowledge, the approach of using an ANN instead of a biophysical model with real neuron data is rather new and is knowingly used for parameter fitting in the Macke Lab[3].

1.2 Goals of this study

In this thesis, we want to know if we can get a better fit to neuron data as a biophysical model, and thus better spike time prediction. We want to abandon the physical assumptions which are in the biophysical model equations, and use an agnostic model instead i.e an ANN. An ANN is parameterized by weights, biases, structure and the activation functions. The question to be resolved will be whether a fitted ANN architecture can match or outperform a fitted biophysical neuron model in terms of spike time prediction.

To achieve this, we want to compare and modify different methods on the challenge data, which are already used for reconstructing nonlinear dynamical systems. Methods such as an RNN model using tractable dendritic RNNs for reconstructing nonlinear dynamical systems (dendPLRNN) [28]. An RNN model using data-driven forecasting (DDF), which are reduced-dimension models relative to HH models as they are built on and forecast only observables [35].
Then compare them against the benchmarks of the winner of the Quantitative Single-Neuron Modeling Competition which were the adaptive exponential integrate-and-fire model (AdEx) and the Autoregressive exogenous model (ARX) [12].

2. Background

Artificial Neural Networks (ANNs) are inspired by the biological neurons found in the brain as they process information and generate output in response to input [27]. However, there are distinct differences between the two types of neurons. Biological neurons usually have dendrites that receive input from other neurons, a cell body that processes the input, and an axon that carries the processed information to other neurons. In contrast, ANNs are composed of neurons with no dendrites and no axons, and instead, receive input through weighted connections and generate output through an activation function. Additionally, biological neurons transmit information by firing action potentials, which allows for spike-rate and spike-time encodings, while single units in ANNs usually do not spike, but have an output analogous to the firing rate of a real neuron allowing only for spike-rate encoding.

Single-neuron recordings are recordings of the electrical signals of a single neuron. This allows scientists to measure the electrical activity of a neuron and learn more about its function and behavior. By measuring the action potentials of a single neuron and analyzing its characteristics, we can gain insight into how neurons communicate with each other, how they operate under different conditions, and how they may be involved in certain diseases [34].

The Quantitative Single-Neuron Modeling Competition has been created to benchmark the performance of single neuron models, and the results of this benchmarking suggest that models combining features of standard leaky integrate-and-fire models with a second variable reflecting adaptation, refractoriness, or a dynamic threshold are the best performers [11].

In Fig. 1a, we show a cortical pyramidal neuron which was simulated by injection of randomly fluctuating currents of various amplitudes. The fluctuating currents were injected and voltage responses were recorded at the soma.

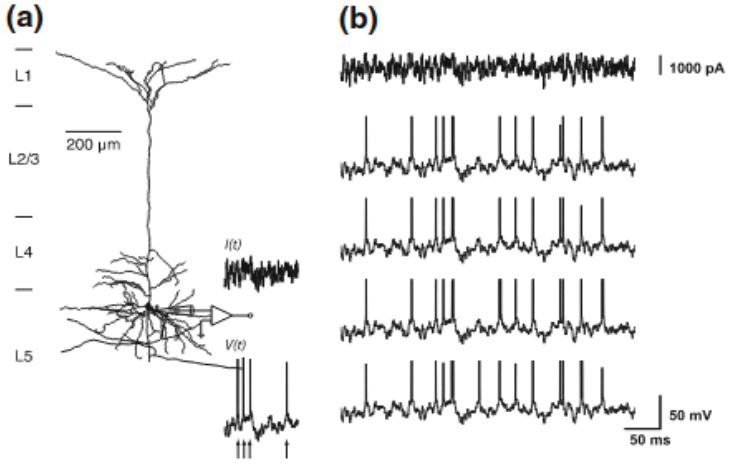

Fig 1. Challenge A. The goal of the challenge is to predict the timing of somatic action potentials with +-2ms precision, given the injected current as a function of time. The data were low-pass filtered at 2.5 kHz with sampling frequency twice the filter frequency[23].

In Fig 1b, we show a sample of injected current (the top) and waveforms in order to evaluate the intrinsic reliability of the cell in terms of the timing of output spikes. The neuron membrane depolarizes and fires action potentials or 'spikes' as a response to this current input.

To predict the occurrence of spikes and their times by using only the data provided by the competition, we want to use methods that already are providing different approaches to the estimation of parameters of the underlying dynamical systems.
To achieve this we want to use a specific algorithm used in the dendritic, piecewise linear recurrent neural network (dendPLRNN) approach, called sparsely forced backpropagation through time (BPTT) [30].
Another approach for comparison is data-driven forecasting (DDF) on neurons [35]. DDF is to be viewed as a way to capture observed aspects of neurobiological data in contrast to the usual method of creating a detailed biophysical model, often of the Hodgkin-Huxley (HH) type.

2.1 What is the approach of this thesis?

All code is written in Python using Pytorch as the main deep learning framework.

All backpropagation in Pytorch is implemented by automatically differentiating the operations in the forward pass of the network, and adding explicit operations for computing the gradient at each point in the network [29].

The data consists of 13 repetitions of a 60s-stimulation protocol. The training set is the first 39 seconds, the test set is the last 21 seconds. In the challenge, we have to predict the spike timing of a regular spiking L5 pyramidal cell responding to in-vivo-like current injection and compare the predicted output spike times to the outputs of the competitors.

The input of our RNN approach using the DDF model and the dendPLRNN model will be the same as the competition. We input the injected current waveform in pA (peak amplitude) for both the training and test phases. For training, we use the voltage trace of each repetition for the first 39s in mV (millivolts). For testing, we predict the spike times of the remaining 21s in each repetition (both the training and test sets are using the same current and voltage trace input managed in a time points file). The spike time is defined as the time at which the voltage recording is crossing 0 mV.

The measure of the performance is based on the Gamma coincidence factor [31]. To evaluate this quantity, we calculate the number of coincidences between the spikes in the data spike train (the binary waveform) one at a time (the gold standard) and the spike train of our LSTM model. This number is calculated by counting the number of target spikes for which we can find at least one model spike within 4ms. We subtract the expected number of coincidences that a Poisson spike train with the same average frequency would give and divide it by the number of spikes in the data and model spike trains. The coincidence factor, which is calculated through the overlapping of spikes from our model and the data subtracted by the number of coincidences, will be our prediction identifier [32]. We will calculate the gamma coincidence factor for each repetition between the predicted spike train and the observed spike train. The performance is the averaged coincidence factor. Furthermore, the standard error of the mean is computed with the bootstrapping method.

As described in the Challenge, our model will outperform the competitor models if the performance is one standard deviation above the performance of all other competitors.

References

[1] Izhikevich, E. M., & Edelman, G. M. (2008). Large-scale model of mammalian thalamocortical systems. *Proceedings of the National Academy of Sciences*, *105*(9), 3593–3598. https://doi.org/10.1073/pnas.0712231105 [Accessed 01 Dec. 2022].

[2] Markram, H. (2006). The Blue Brain Project. *Nature Reviews Neuroscience*, *7*(2), 153–160. https://doi.org/10.1038/nrn1848 [Accessed 01 Dec. 2022]

[3] Gonçalves, P. J., Lueckmann, J.-M., Deistler, M., Nonnenmacher, M., Öcal, K., Bassetto, G., Chintaluri, C., Podlaski, W. F., Haddad, S. A., Vogels, T. P., Greenberg, D. S., & Macke, J. H. (2020, September 17). *Training deep neural density estimators to identify mechanistic models of neural dynamics*. eLife Sciences Publications, Ltd. https://elifesciences.org/articles/56261 [Accessed 01 Dec. 2022]

[4] Jolivet, R., Schürmann, F., Berger, T. K., & Roth, A. (2008, December 1). *The quantitative single-neuron modeling competition*. Springer Verlag.

https://www.researchgate.net/publication/23476245_The_Quantitative_Single-Neuron_Modeling_Competition [Accessed 01 Dec. 2022]

[5] Camera, G. L., Rauch, A., Lüscher, H.-R., Senn, W., & Fusi, S. (2004). Minimal models of adapted neuronal response to in vivo–like input currents. *Neural Computation*, *16*(10), 2101–2124. https://doi.org/10.1162/0899766041732468 [Accessed 01 Dec. 2022]

[6] Ramirez, J.-M. (2003). Faculty Opinions recommendation of Alternative to hand-tuning conductance-based models: Construction and analysis of databases of model neurons. *Faculty Opinions – Post-Publication Peer Review of the Biomedical Literature*. https://doi.org/10.3410/f.1015987.199365 [Accessed 01 Dec. 2022]

[7] Brillinger, D. R., & Segundo, J. P. (2011). Empirical examination of the threshold model of neuron firing. In *Selected Works of David Brillinger* (pp. 569–576). Springer New York. http://dx.doi.org/10.1007/978-1-4614-1344-8_32 [Accessed 01 Dec. 2022]

[8] Kim, & Panda. (2021). Visual explanations from spiking neural networks using inter-spike intervals. *Scientific Reports*, *11*(1), 1–14. https://doi.org/10.1038/s41598-021-98448-0 [Accessed 01 Dec. 2022]

[9] Williams, & Zipser. (1989). A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, *1*(2), 270–280. https://doi.org/10.1162/neco.1989.1.2.270 [Accessed 01 Dec. 2022]

[10] *A ten-minute introduction to sequence-to-sequence learning in Keras*. (n.d.). Retrieved January 25, 2023, from https://blog.keras.io/a-ten-minute-introduction-to-sequence-to-sequence-learning-in-keras.html [Accessed 01 Dec. 2022]

[11] Jolivet, R., Kobayashi, R., Rauch, A., Naud, R., Shinomoto, S., & Gerstner, W. (2008). A benchmark test for a quantitative assessment of simple neuron models. *Journal of Neuroscience Methods*, *169*(2), 417–424. https://doi.org/10.1016/j.jneumeth.2007.11.006 [Accessed 01 Dec. 2022]

[12] Brette, R., & Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of Neurophysiology*, *94*(5), 3637–3642. https://doi.org/10.1152/jn.00686.2005 [Accessed 01 Dec. 2022]

[13] Xie, Z., He, F., Fu, S., Sato, I., Tao, D., & Sugiyama, M. (2021). Artificial neural variability for deep learning: On overfitting, noise memorization, and catastrophic forgetting. *Neural Computation*, *33*(8), 2163–2192. https://doi.org/10.1162/neco_a_01403 [Accessed 01 Dec. 2022]

[14] Mikhaeil, J. M., Monfared, Z., & Durstewitz, D. (2021, October 14). *On the difficulty of learning chaotic dynamics with RNNs*. arXiv.Org. https://arxiv.org/abs/2110.07238 [Accessed 01 Dec. 2022]

[15] Feng, Y., Gu, S., Guo, D., Yang, Z., & Shao, C. (2021). Guiding teacher forcing with seer forcing for neural machine translation. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. http://dx.doi.org/10.18653/v1/2021.acl-long.223 [Accessed 01 Dec. 2022]

[16] Stuart, G. J., & Sakmann, B. (1994). Active propagation of somatic action potentials into neocortical pyramidal cell dendrites. *Nature*, *367*(6458), 69–72. https://doi.org/10.1038/367069a0 [Accessed 01 Dec. 2022]

[17] Galovič, M. (2018, November 2). Time series forecasting with rnns. *Towards Data Science*. https://towardsdatascience.com/time-series-forecasting-with-rnns-ff22683bbbb0 [Accessed 01 Jan. 2023]

[18] Onnen, H. (2021, November 1). Python RNN: Recurrent neural networks for time series forecasting. *Towards Data Science*. https://towardsdatascience.com/temporal-loops-intro-to-recurrent-neural-networks-for-time-series-forecasting-in-python-b0398963dc1f [Accessed 01 Jan. 2023]

[19]Teka, W., Marinov, T. M., & Santamaria, F. (2014). Neuronal spike timing adaptation described with a fractional leaky integrate-and-fire model. *PLOS Computational Biology*, *10*(3). https://doi.org/10.1371/journal.pcbi.1003526 [Accessed 15 Jan. 2023]

[20]Alonso, L. M., & Marder, E. (2020, June 2). *Temperature compensation in a small rhythmic circuit*. eLife Sciences Publications, Ltd. https://elifesciences.org/articles/55470 [Accessed 15 Jan. 2023]

[21] Wong, W. (2019, October 15). What is Teacher Forcing? *Towards Data Science*. https://towardsdatascience.com/what-is-teacher-forcing-3da6217fed1c [Accessed 20 Jan. 2023]

[22] Liang, Wu, & Gu. (2016, January 1). *Single neuron models*. Springer Netherlands. https://link.springer.com/chapter/10.1007/978-94-017-7393-5_3 [Accessed 20 Jan. 2023]

[23] Rauch, A., Camera, G. L., Luscher, H.-R., & Fusi, S. (2003, October 1). *Neocortical pyramidal cells respond as integrate-and-fire neurons to in vivo-like input currents*. American Physiological Society. https://www.researchgate.net/publication/7276100_Neocortical_Pyramidal_Cells_Respond_as_Integrate-and-Fire_Neurons_to_In_Vivo-Like_Input_Currents [Accessed 25 Jan. 2023]

[24] Lamb, A., Goyal, A., Zhang, Y., Zhang, S., Courville, A., & Bengio, Y. (2016, October 27). *Professor forcing: A new algorithm for training recurrent networks*. arXiv.Org. https://arxiv.org/abs/1610.09038 [Accessed 25 Jan. 2023]

[25] HODGKIN, A. L., & HUXLEY, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal of Physiology*, *117*(4), 500–544. https://doi.org/10.1113/jphysiol.1952.sp004764 [Accessed 25 Jan. 2023

[26] Lenk. (2011, January 1). *A simple phenomenological neuronal model with inhibitory and excitatory synapses*. Springer Berlin Heidelberg. https://link.springer.com/content/pdf/10.1007/978-3-642-25020-0_30.pdf?pdf=inline%20link [Accessed 01 Mar. 2023]

[27] Mohamed. (2018, January 1). *Brain-Inspired machine learning algorithm: Neural network optimization*. Springer International Publishing. https://link.springer.com/chapter/10.1007/978-3-319-75714-8_6 [Accessed 01 Mar. 2023]

[28] Durstewitz, M. B., Florian Hess, Jonas M. Mikhaeil, Leonard Bereska, Zahra Monfared, Po-Chen Kuo, Daniel. (n.d.). *Tractable dendritic rnns for reconstructing nonlinear dynamical systems*. https://arxiv.org/pdf/2207.02542.pdf [Accessed 01 Mar. 2023]

[29] *A Gentle Introduction to torch.autograd — PyTorch Tutorials 2.1.0+cu121 documentation*. (n.d.). Retrieved Feb 3, 2023, from https://pytorch.org/tutorials/beginner/blitz/autograd_tutorial.html [Accessed 01 Mar. 2023]

[30] Neural machine translation with attention. (n.d.). *TensorFlow*. Retrieved June 9, 2023, from https://www.tensorflow.org/text/tutorials/nmt_with_attention#text_preprocessing [Accessed 01 Mar. 2023]

[31] Drescher, A., Yoho, M., & Landsberger, S. (2018). Gamma–gamma coincidence in neutron activation analysis. *Journal of Radioanalytical and Nuclear Chemistry*, *318*(1), 527–532. https://doi.org/10.1007/s10967-018-6033-8 [Accessed 01 May. 2023]

[32] Jolivet, R., J., T., & Gerstner, W. (2003, January 1). *The spike response model: A framework to predict neuronal spike trains*. Springer Berlin Heidelberg. https://link.springer.com/chapter/10.1007/3-540-44989-2_101 [Accessed 01. Jun. 2023]

[33] Trinh, T. H., Dai, A. M., Luong, M.-T., & Le, Q. V. (2018, March 1). *Learning longer-term dependencies in rnns with auxiliary losses*. arXiv.Org. https://arxiv.org/abs/1803.00144 [Accessed 01. Aug 2023]

[34] Geisler, W. S., Albrecht, D. G., Salvi, R. J., & Saunders, S. S. (1991). Discrimination performance of single neurons: Rate and temporal-pattern information. *Journal of Neurophysiology*, *66*(1), 334–362. https://doi.org/10.1152/jn.1991.66.1.334 [Accessed 01. Jun 2023]

[35] Clark, Fuller, Platt, & Abarbanel. (2022). Reduced-Dimension, biophysical neuron models constructed from observed data. *Neural Computation*, *34*(7), 1545–1587. https://doi.org/10.1162/neco_a_01515 [Accessed 01. Aug 2023]