Delayed Synapses: An LSM Model for Studying Aspects of Temporal Context in Memory

Predrag Jakimovski and Hedda R. Schmidtke

Pervasive Computing Systems, TecO,
Vincenz-Prießnitz-Strasse 3, 76131 Karlsruhe, Germany
{predrag.jakimovski,hedda.schmidtke}@kit.edu
http://www.teco.edu

Abstract. Spiking neural networks are promising candidates for representing aspects of cognitive context in human memory. We extended the liquid state machine model with time-delayed connections from liquid neurons to the readout unit to better capture context phenomena. We performed experiments in the area of spoken language recognition for studying two aspects of context dependency: influence of memory and temporal context. For the experiments, we derived a test data set from the well-known Brody-Hopfield test set to which we added varying degrees of Gaussian noise. We studied the influence of temporal context with a further specially designed test set. We found that the temporal context encoded in the pattern to be recognized was recognized better with our delayed synapses than without. Our experiments shed light on how context serves to integrate information and to increase robustness in human signal processing.

Keywords: Spiking neural networks (SNN), liquid state machine (LSM), context-dependency, memory acquirements, fading memory (FM)

1 Introduction

Artificial neural networks (ANN) have been proposed as a means to model the human ability to successfully adapt to, and react in a complex environment. However, ANNs have the disadvantage of not modeling the information encoded temporally in the signal. A further class of neural networks are spiking neural networks (SNN), which mimic the behavior of biological neural networks. A liquid state machine consists of a SNN with a readout unit implemented by perceptrons. The readout unit of the LSM proposed by [12] interprets only the current snapshot of the liquid states. In this paper, the concept of an LSM is extended, so that not only the current liquid state can be taken into account but also past liquid states. The extension of the approach is realized using time-delayed connections from liquid neurons to the readout unit.

The structure of the article is as follows. We first introduce the classical model and our extension for modeling memory (Sect. 2). We then present our experimental setup and the test-bed we implemented, and discuss the results (Sect. 3).

2 Liquid State Machine Model

The human brain processes information from the environment in real-time and immediately delivers a "meaningful" response. The liquid state machine (LSM) was developed as a model to simulate properties of a biological brain and to mimic its behavior for real-time processing [10]. A liquid state machine consists of a spiking neural network (SNN) with a readout unit implemented by sigmoid neural networks. The approach is based on the idea that a recurrent neural network that is sufficiently dynamically complex and becomes excited by a continuous stream of information $u(\cdot)$, such as a stream of auditory input, retains this information for a longer duration while processing it. The complex network becomes an information *medium*. The spatio-temporal stream is transformed into a high-dimensional spatial pattern $x(t)$, also called the *liquid state*, from which information can be retrieved. The metaphor refers to water into which stones are thrown. As the stones hit the water splashes and ripples on the surface are created, from which an observer can infer the location where the stone where thrown. The role of the observer is played by a readout unit, which can be implemented by a simple machine learning mechanism, such as a neural network [3].

Formally an LSM M is a filter L^M. Mathematically, L^M is an operator, which when applied to a function u, yields an intermediate result x: $x(t) = (L^M u)(t)$. The output $x(t)$ depends on $u(t)$ and, in a non-linear manner, on the previous inputs $u(s)$. A readout function f^M transforms the output $x(t)$ into a target output $y(t)$ so that $y(t) = f(x(t))$. The classical architecture of the LSM model is shown in figure 1.

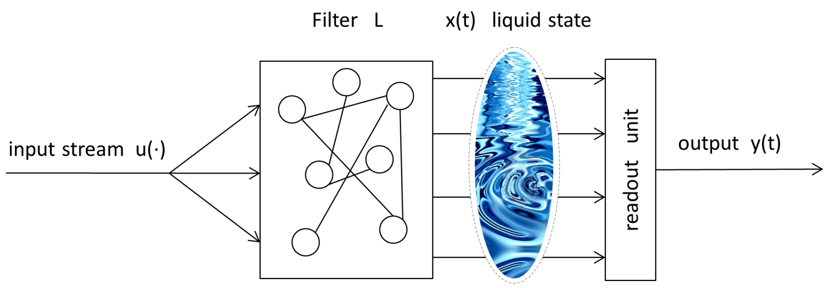

Fig. 1. The flow-graph shows the principal processing of a conventional LSM: A low-dimensional input stream $u(t)$ excites filter L, consisting of a recurrent neural network (RNN); A readout unit is interpreting immediately the liquid state $x(t)$ generated by L. Inside of the dotted ellipse indicates water circles and represents the current liquid state.

In general the readout unit f is memoryless, i. e. the former liquid states $x(s)$, $s < t$, are not considered at t. Therefore, the chosen type of liquid filter L

must have the ability to retain the past information into the liquid state $x(t)$. In this way the readout unit f is enabled to generate adequate correct output $y(t)$. This will be only possible if the liquid filter L fulfills the *separation property* (SP) [10]: For any input stream u, v the liquid filter L is reacting sufficiently different, i. e. the liquid filter L is producing different liquid states $x(t)$ by $u(\cdot) \neq v(\cdot)$.

A liquid filter L which exhibits the SP, features also the property of *fading memory (FM)* [2]. That means that the liquid filter L is retaining information of an input stream u over a fixed time span. In the modeling of information processing, the fading memory property is essential because novel incoming information can only be processed as novel information if previous information fades.

In the realization of the LSM, the liquid state $x(t)$ is modeled by a n-dimensional vector. The dimension of the vector is fixed by the number of neurons used in the liquid filter. For the neuron-type, the classical Leaky-Integrate-and-Fire-neuron (cf. [10] for an introduction) is chosen, which fires a spike at a fixed time point t_i and produces a discrete time series of spikes called spike train. The spike activity of every single neuron in a time window determines after a low-pass filtering the n components of the liquid vector. The n-dimensional vector is called a *snapshot* of the liquid state $x(t_i)$. In practice, a time window of length with $30ms$ up to $150ms$ is chosen based on experimental data acquired from human behavior experiments [4].

In comparison to its simple neural network readout unit, the LSM achieves better classification results. The advantage of the LSM is its specific way of integrating temporal context: everything within a certain time window is retained, but everything outside is removed as the fading memory property requires. The question then is how memories, as they are retained for a longer duration, i.e. on a coarser temporal granularity, can be integrated with the model. While we cannot offer an answer to this question on the neurobiological side, we propose the idea of *delayed synapses* to model integration on a coarser level of temporal granularity, i.e. for time windows of a longer duration [13].

The challenge is to integrate memory functionality into the model, in order to acquire context. To realize this we employ so-called delayed synapses, which set back the emitting spikes of a neuron N_i by using different time delays. In particular all delayed synapses belonging to a neuron N_i are leading to the readout unit directly. In this way past information that occurred before the current time window of the LSM is provided to the readout unit. When the readout unit captures the current liquid state $x(t)$, it captures also the past simultaneously. In Figure 2, an operating LSM with delayed synapses is shown.

3 Experiments

To give proof to our approach we initially deployed the Brody-Hopfield benchmark to show the capabilities of the LSM. The benchmark itself originates in speech recognition and was used to test SNN models by Brody and Hopfield [5, 6] for the first time. The speech benchmark consists of 500 audio files, recorded

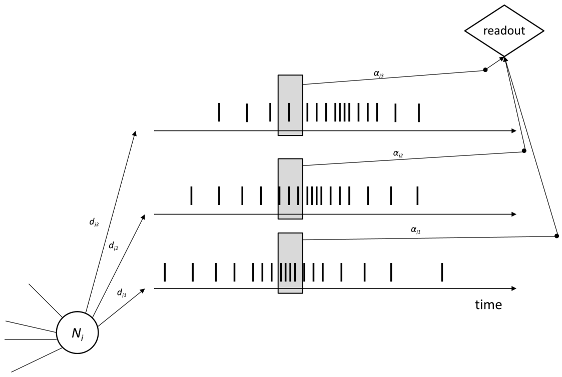

Fig. 2. Illustration of the readout event. The three readout synapses delay the spike emitted by the liquid-neuron N_i by d_{i_1}, d_{i_2} and d_{i_3} respectively, so that the readout unit can evaluate liquid-states with regard to the time. During training phase of the readout unit the weights α_{i_1}, α_{i_2} and α_{i_3} will be adjusted by means of the p-delta learning rule.

by five female speakers speaking the digits “zero” to “nine”. The task here is to recognize the spoken digits successfully. For testing our approach we transformed the audio files into spatio-temporal spike patterns by using auditory models of the human ear. It is assumed that the inner hair cells of the ear transmit the electromagnetic pulses called spikes onto the auditory nerve directly.

To build up the evaluation platform a generator based on work in [11] for creating SNN was implemented. SNN with 135 neurons and 40 input neurons were created and for the readout unit of the LSM an one-layer ANN was deployed, called parallel perceptron which was developed by Auer [1].

3.1 Test with noisy Brody-Hopfield digits

So far, the whole framework of the LSM were integrated and tested by the Brody-Hopfield benchmark. Here, the digits ”one” and ”four” were used to train and test the LSM and white noise was added to increase recognition difficulty. The new data set consists of 500 samples partitioned into randomly chosen 300 samples for training and 200 samples for the test set. During the training phase 10-fold cross-validation was applied. The evaluation was performed by two randomly generated LSMs, i. e. *lsm_0adv2.net* and *lsm_0adv8.net* were chosen. The readout interval for the parallel perceptron was set to $[150, 180]$ms.

The tests yield success rates of 85.5% and 71.0%, respectively. Training the readout unit without any liquid filter resulted in 62.5% success. The evaluation results of our model extension are listed in table 1. The setting for each trial is fixed by the number of delayed synapses and the delay. For example, the set-up with 3 delayed synapses and 10ms delay means three links from each liquid neuron to the readout unit delays the emitting liquid neuron in 10ms, 20ms

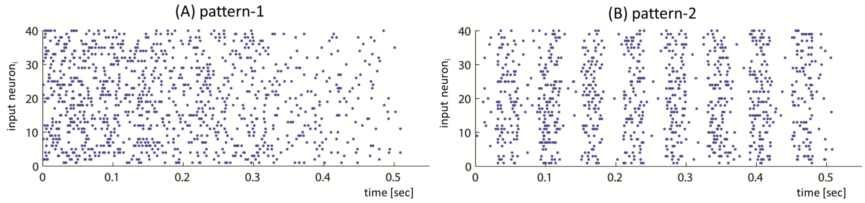

Fig. 3. Two patterns in the ArtBM-Benchmark.

and 30ms, respectively. In this way context information from the past can be captured by the readout unit.

number of synapses	*lsm_0adv2.net*						*lsm_0adv8.net*					
	5ms	10ms	15ms	20ms	30ms	50ms	5ms	10ms	15ms	20ms	30ms	50ms
1 del. syn.	84.5	87.0	87.0	86.0	86.5	86.0	72.0	73.0	73.0	74.0	76.5	73.5
2 del. syn.	86.0	87.0	88.5	88.5	87.5	86.5	72.0	74.0	77.5	74.0	77.0	83.0
3 del. syn.	87.0	87.5	88.5	88.0	85.5	87.0	71.5	77.0	75.0	75.5	80.5	88.0
5 del. syn.	87.5	89.0	87.5	86.0	89.5	–	74.5	74.0	76.0	81.5	91.5	–
7 del. syn.	88.0	88.0	86.0	87.5	–	–	76.5	76.0	81.5	93.0	–	–
10 del. syn.	89.5	86.0	89.0	–	–	–	74.0	81.0	92.0	–	–	–
$\frac{1}{n}\sum_{i=1}^{n} R_i$	87.08	87.42	87.75	87.2	87.25	86.5	73.42	75.83	79.17	79.60	81.38	81.5

Table 1. Performance results of trails with different number of applied delayed synapses.

The evaluation described previously revealed a conventional LSM provides in general better performance than a trained parallel perceptron without a preceding liquid filter. As table 1 shows, the model extension using time delayed synapses provides better performance than the classical LSM. However, through the evaluation new issues came up: Why is there a large performance difference between the investigated LSMs? Why does the performance not rise steadily by increasing the number of delayed synapses? Further on, why does the performance vary of the chosen LSMs differently strong? In the following this issues are discussed.

3.2 Further investigation results

Investigating the separation property (SP) of the engaged LSMs by using constructed input streams with known distances to each other which are calculated by l^2-norm, the output trajectories indicate the separability during time (not depicted here). These explained the strong different performance of the LSMs, but did not show the proof of real improvement. In order to do this, we designed an artificial benchmark for binary classification called ArtBM. Here, the samples consisting of spatio-temporally spike patterns were constructed in the manner that spatial information in the signal is suppressed and temporal encoded. The spatial and temporal structure of the samples is best explained by Fig 3. The entire benchmark consist of 1500 samples divided into 1000 randomly chosen samples for training and 500 test samples. For the evaluation setting the read-out interval was fixed to $[510, 540]$ms. The evaluation of the classical approach resulted in 59.6% and 81.4% for the two chosen LSMs, respectively. The results of the model extension are listed in table 2. Here, the model extension yielded almost better performance by choosing different evaluation settings. The best result with 99.8% success was achieved by choosing 5 delayed synapses and a time delay of 20ms. In this trial only 1 of 500 test samples was classified false.

lsm_0adv2.net	1 delayed synapse	2 delayed synapses	3 delayed synapses	4 delayed synapses	5 delayed synapses
10ms	64.6	72.6	83.0	92.0	94.8
20ms	76.2	93.0	93.8	93.8	96.2
30ms	86.8	89.8	93.6	93.8	91.6
60ms	85.4	92.8	90.2	86.6	85.8
90ms	79.8	82.2	83.0	88.2	90.4
180ms	78.2	90.4	90.4	–	–

lsm_0adv8.net	1 delayed synapse	2 delayed synapses	3 delayed synapses	4 delayed synapses	5 delayed synapses
10ms	81.4	83.4	89.8	97.2	99.2
20ms	84.0	98.4	98.6	98.6	99.8
30ms	91.8	98.6	98.6	99.2	98.8
60ms	94.8	98.0	98.4	99.2	99.0
90ms	82.4	89.8	91.6	82.8	94.4
180ms	84.6	71.8	71.8	–	–

Table 2. Evaluation results of our LSM model extension

4 Conclusion

Spiking neural networks are promising candidates for studying aspects of temporal context in human memory. In this paper, we extended the liquid state machine model with time-delayed connections from liquid neurons to the read out unit to capture and study the influence of temporal contexts of different durations. We performed two types of experiments in the area of spoken language recognition. For the first experiments, we derived a test data set from the well-known Brody-Hopfield test set to which we added varying degrees of Gaussian noise. We studied the influence of temporal context with a specially designed test set. We found that the temporal context encoded in the pattern was recognized better with our delayed synapses than without. Our experiments shed light on how context serves to integrate information and to increase robustness in human signal processing.

LSM are just one possible approach in which it is possible to model context-dependency. In future works we plan to generalize our approach to context-dependency to generalize it to alternative approaches, such as the "echo state" approach [7, 8]. In particular, we hope to find an implementation for our context-dependent pattern recognition that is similarly powerful as LSMs and also applicable as a tool for real-time pattern recognition.

Acknowledgments. This research is supported by the project Polytos, in the BMBF-funded excellence cluster "Printed Organic Circuits and Chips"[1].

References

1. Auer, P., Burgsteiner, H.M., Maass, W.: The p-delta learning rule for parallel perceptrons. Science (2002)
2. Boyd, S., Chua, L.: Fading memory and the problem of approximating nonlinear operators with voltera series. IEEE Trans. on Circuits and Systems 32, 1150–1161 (1985)
3. Feldbusch, F., Kaiser, F.: Simulation of spiking neural nets with INspiRE. IEEE Conference on Systems, Man and Cybernetics SMC (2005)
4. Gerstner, W., Kistler, W.: Spiking Neuron Models. Cambridge University Press (2002)
5. Hopfield, J.J., Brody, C.D.: What is a moment? cortical sensory integration over a brief interval. Proc. Natl. Acad. Sci. USA 97(25), 13919–13924 (2000)
6. Hopfield, J.J., Brody, C.D.: What is a moment? transient synchrony as a collective mechanism for spatiotemporal integration. Proc. Nat. Acad. Sci. USA 98(3), 1282–1287 (2001)
7. Jaeger, H.: Adaptive nonlinear system identification with echo state networks. Advances in Neural Information Processing Systems 15, 593–600 (2003)
8. Jaeger, H., Haas, H.: Harnessing nonlinearity: Predicting chaotic systems and saving energy in wireless communication. Science 304, 78–80 (2004)
9. Maass, W., Joshi, P., Sontag, E.: Computational aspects of feedback in neural circuits. PLOS Computational Biology (2006)
10. Maass, W.: Liquid state machines: Motivation, theory and applications. World Scientific Review Volume 189 (2010)
11. Maass, W., Markram, H.: On the computational power of recurrent circuits of spiking neurons. In: Electronic Colloquium on Computational Complexity. vol. 22 (2001)
12. Maass, W., Natschläger, T., Markram, H.: Real-time computing without stable states: A new framework for neural computation based on perturbations. Neural Computation 14(11), 2531–2560 (2002)
13. Schmidtke, H.R.: Granularity as a parameter of context. In: Dey, A.K., Kokinov, B.N., Leake, D.B., Turner, R.M. (eds.) International Conference on Modeling and Using Context. LNCS, vol. 3554, pp. 450–463. Springer (2005)

[1] http://www.innovationlab.de/en/research/excellence-cluster-forum-organic-electronics/research-projects/printed-organic-circuits-and-chips-polytos/