

---

# Towards efficient end-to-end speech recognition with biologically-inspired neural networks

---

**Thomas Bohnstingl\***  
IBM Research, Zurich  
Graz University of Technology

**Ayush Garg**  
IBM Research, Zurich  
ETH Zurich

**Stanisław Woźniak**  
IBM Research, Zurich

**George Saon**  
IBM Research AI, Yorktown Heights

**Evangelos Eleftheriou†**  
IBM Research, Zurich

**Angeliki Pantazi**  
IBM Research, Zurich

## Abstract

Automatic speech recognition (ASR) is a capability which enables a program to process human speech into a written form. Recent developments in artificial intelligence (AI) have led to high-accuracy ASR systems based on deep neural networks, such as the recurrent neural network transducer (RNN-T). However, the core components and the performed operations of these approaches depart from the powerful biological counterpart, i.e., the human brain. On the other hand, the current developments in biologically-inspired ASR models, based on spiking neural networks (SNNs), lag behind in terms of accuracy and focus primarily on small scale applications. In this work, we revisit the incorporation of biologically-plausible models into deep learning and we substantially enhance their capabilities, by taking inspiration from the diverse neural and synaptic dynamics found in the brain. In particular, we introduce neural connectivity concepts emulating the axo-somatic and the axo-axonic synapses. Based on this, we propose novel deep learning units with enriched neuro-synaptic dynamics and integrate them into the RNN-T architecture. We demonstrate for the first time, that a biologically realistic implementation of a large-scale ASR model can yield competitive performance levels compared to the existing deep learning models. Specifically, we show that such an implementation bears several advantages, such as a reduced computational cost and a lower latency, which are critical for speech recognition applications.

## 1 Introduction

Speech is one of the most natural ways for human beings to communicate and to interact, which led several research groups into developing speech recognition applications using computer systems. However, even if the origin of the sound signal is the same for the humans and the computers, the transformation, the feature extraction and the further processing differ significantly. Particularly, the later often utilizes mathematical abstractions, such as the Hidden Markov Model (HMM) or domain-specific models [1; 2; 3; 4]. More recently, AI researchers have increasingly resorted to deep learning approaches [5; 6; 7; 8]. Especially three kinds of architectures have been widely adopted in the literature [9; 10]: architectures based solely on recurrent neural networks (RNN), such as the RNN transducer [11], architectures based on RNNs with the attention mechanism, such as the listen attend and spell model (LAS) [12], and recently also transformer-based models, such as the conformer [13]. While the transformer-based and attention-based models potentially provide higher accuracy, they require the entire audio signal to compute the transcript and hence have a high latency.

---

\*Correspondence to boh@zurich.ibm.com

† Currently with Axelera AI, Zurich

In contrast, the RNN-T architecture can produce a transcript while receiving the audio signal and thus exhibits a lower latency, allows for real-time transcription and is even commercially deployed, for example in mobile devices [14]. Recently, researchers have also tried to merge the RNN-T and the LAS model in order to leverage the benefits of both [15; 16; 17].

Despite their great successes, all the aforementioned models take inspiration from biology only remotely and depart considerably from the human speech recognition system. Researchers working on biologically plausible spiking neural networks (SNNs), have used the leaky integrate-and-fire (LIF) neurons for speech recognition [18; 19; 20]. However, there are several limitations of these works. For example, many of them use simpler network architectures compared to the ones from the ML domain. Moreover, often only parts of the network architecture employ biologically-inspired units and the learning algorithms used are inferior to error-backpropagation [21; 22]. Thus, the performance of those approaches lags behind their ML-based counterparts in terms of accuracy.

In this paper, we address these limitations by proposing an RNN-T architecture that incorporates biologically-inspired neural units. Specifically, we leverage the diverse neuron and synapse types observed in the brain and propose novel extensions of the LIF model with much richer dynamics. We build upon the spiking neural units (SNUs) [23], which allows us to leverage the advanced training capabilities from the ML domain [24; 25] that are essential for speech recognition [26]. We demonstrate that a state-of-the-art network incorporating biologically-inspired units can yield competitive performance levels in a large-scale speech recognition application, while significantly reducing the computational cost as well as the latency. In particular, our key contributions are:

- a novel neural connectivity concept emulating the axo-somatic synapses that enhances the threshold adaptation in biologically-inspired neurons,
- a novel neural connectivity concept emulating the axo-axonic synapses that modulates the output of biologically-inspired neurons,
- a biologically-inspired RNN-T architecture with significantly reduced computational cost and increased throughput, which still provides a competitive performance to the LSTM-based architecture.

## 2 Biologically-inspired models for deep learning

Typically, architectures of biologically-inspired neurons consider only axo-dendritic synapses, in which the output from the pre-synaptic neuron travelling down the axon is modulated by the synaptic weight and arrives at the post-synaptic dendrite. However, neural networks in the brain exhibit a much more complex connectivity structure [27; 28]. In this work, we leverage this diversity and investigate two additional types of synapses, namely the axo-somatic and the axo-axonic synapses, and propose novel biologically-inspired models with richer dynamics. Figure 1a shows an illustration of such neuron models and their connectivity via the various types of synapses.

The simplest biologically-inspired model is based on the LIF dynamics which employs axo-dendritic synapses only, see dark blue color in Figure 1a. It can be abstracted into the form of a recurrent unit using the SNU framework [30]. Figure 1b shows the basic configuration of the SNU. A layer of  $n$  SNUs receiving  $m$  inputs is governed by the following equations:

$$\mathbf{s}^t = \mathbf{g}(\mathbf{W}\mathbf{x}^t + \mathbf{H}\mathbf{y}^{t-1} + d \cdot \mathbf{s}^{t-1} \odot (\mathbf{1} - \mathbf{y}^{t-1})), \quad (1)$$

$$\mathbf{y}^t = \mathbf{h}(\mathbf{s}^t + \mathbf{b}), \quad (2)$$

where  $\mathbf{x}^t \in \mathbb{R}^m$  represents the inputs at time step  $t$ ,  $\mathbf{s}^t \in \mathbb{R}^n$  represents the membrane potentials, the hidden states,  $\mathbf{y}^t \in \mathbb{R}^n$  represents the outputs,  $d \in \mathbb{R}$  represents the constant decay of the membrane potential,  $\mathbf{b} \in \mathbb{R}^n$  represents the trainable firing threshold,  $\mathbf{W} \in \mathbb{R}^{n \times m}$  and  $\mathbf{H} \in \mathbb{R}^{n \times n}$  denote the trainable input and the recurrent weight matrices. Note that although the SNU shares similarities with an RNN, e.g., the term  $\mathbf{W}\mathbf{x}^t + \mathbf{H}\mathbf{y}^{t-1}$  and the activation function  $\mathbf{h}$ , it has a richer dynamics, including a state decay and an internal state reset mechanism, i.e.,  $d \cdot \mathbf{s}^{t-1} \odot (\mathbf{1} - \mathbf{y}^{t-1})$ . As described in [23], the SNU can operate in principle in two different modes, one in which it yields continuous outputs (called sSNU), i.e.  $\mathbf{h} = \sigma$ , where  $\sigma$  denotes the sigmoid function, and one in which it emits discrete spikes, i.e.  $\mathbf{h} = \Theta$ , where  $\Theta$  indicates the Heaviside function. In this work, we focus on the dynamics of the neuronal models, in particular on the different synapse types, and thus mainly consider the sSNU variants.

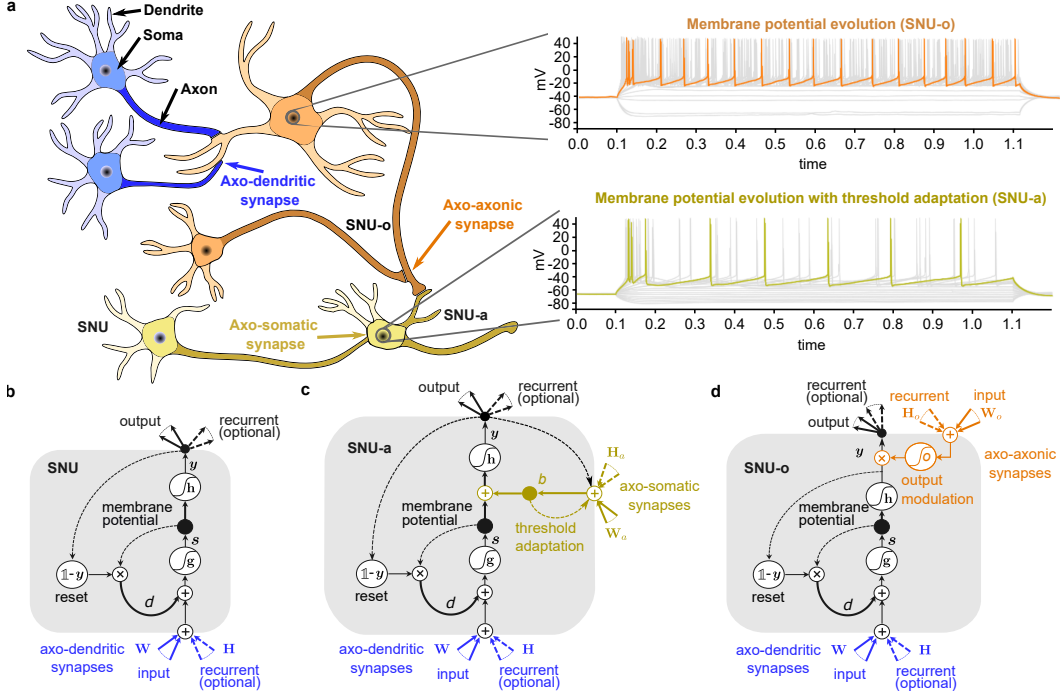


Figure 1: **Illustration of different biological synapse and neuron types and their realization in simulations.** **a** A neural network example composed of three neuron and three synapse types along with recordings of membrane potentials from the Human Brain Atlas [29]. **b** Visualization of SNU modelling the LIF dynamics using solely axo-dendritic synapses. **c** Visualization of the SNU-a modelling the adaptive threshold behavior using axo-somatic synapses, illustrated with yellow color in **a**. **d** Visualization of the SNU-o using axo-axonic synapses, illustrated in orange in **a**.

One well-known property of neurons in the human brain is that they can adapt their firing threshold across a wide variety of timescales [29; 31; 30]. In biology, there are complex mechanisms influencing the firing threshold. In particular, it can be increased or decreased following the neuron’s own dynamics, but also based on the activity of the other neurons [32; 33]. We propose a novel adaptive SNU (SNU-a) that enhances the threshold adaptivity dynamics by emulating the axo-somatic synapses, indicated with yellow color in Figure 1a. A layer of  $n$  SNU-a units is governed by

$$s^t = \mathbf{g}(\mathbf{W}\mathbf{x}^t + \mathbf{H}\mathbf{y}^{t-1} + d \cdot s^{t-1} \odot (\mathbf{1} - \mathbf{y}^{t-1})), \quad (3)$$

$$\mathbf{b}^t = \boldsymbol{\rho} \odot \mathbf{b}^{t-1} + (1 - \boldsymbol{\rho}) \odot (\mathbf{W}_a \mathbf{x}^t + \mathbf{H}_a \mathbf{y}^{t-1}), \quad (4)$$

$$\mathbf{y}^t = \mathbf{h}(s^t + \beta \mathbf{b}^t + \mathbf{b}_0), \quad (5)$$

where  $\mathbf{b}_0 \in \mathbb{R}^n$  represents the trainable baseline threshold,  $\beta \in \mathbb{R}$  represents a constant scaling factor,  $\mathbf{b}^t \in \mathbb{R}^n$  represents the state of the threshold at time  $t$ ,  $\boldsymbol{\rho} \in \mathbb{R}^n$  represents the constant decay of the threshold,  $\mathbf{W}_a \in \mathbb{R}^n \times m$  and  $\mathbf{H}_a \in \mathbb{R}^n \times n$  denote the trainable input and recurrent weight matrices influencing the threshold via axo-somatic synapses. Figure 1c shows an illustration of the SNU-a and the lower right part of Figure 1a shows an example of the membrane potential evolution.

Another type of neural connectivity in the human brain is via axo-axonic synapses, indicated with orange color in Figure 1a. These synapses mediate the release of neurotransmitters from the pre-synaptic to the post-synaptic neuron. Interneurons (INs), e.g. Chandelier INs, that connect via axo-axonic synapses are found in the brain and often act as inhibitors for the post-synaptic neuron [34; 35]. However, it has also been found that such neurons may be excitatory as well [35]. We incorporate the axo-axonic synapses into the SNU framework and propose SNU-o units, with the following equations

$$s^t = \mathbf{g}(\mathbf{W}\mathbf{x}^t + \mathbf{H}\mathbf{y}^{t-1} + d \cdot s^{t-1} \odot (\mathbf{1} - \tilde{\mathbf{y}}^{t-1})), \quad (6)$$

$$\tilde{\mathbf{y}}^t = \mathbf{h}(s^t + \mathbf{b}^t), \quad (7)$$

$$\mathbf{y}^t = \tilde{\mathbf{y}}^t \odot \mathbf{o}(\mathbf{W}_o \mathbf{x}^t + \mathbf{H}_o \mathbf{y}^{t-1} + \mathbf{b}_o^t) \quad (8)$$

where  $\tilde{y}^t$  represents the unmodulated output of the neuron driving the neural reset,  $y^t$  represents the modulated output of the neuron propagating to other connected neurons. The trainable parameters  $\mathbf{W}_o \in \mathbb{R}^{n \times m}$  and  $\mathbf{H}_o \in \mathbb{R}^{n \times n}$  denote the input and recurrent weight matrices and  $\mathbf{b}_o^t \in \mathbb{R}^n$  denotes the trainable bias term that all three modulate the neuronal output via axo-axonic connections. Note that in our simulations we use the sigmoid function as the activation for the output modulation in Equation 8, i.e.,  $\mathbf{o} = \sigma$ , to mimic the inhibitory character of these synapses. However, by using a different activation function  $\mathbf{o}$ , the outputs can also be modulated in an excitatory manner. Figure 1d shows an illustration of the SNU-o with the output modulation and the upper right part of Figure 1a shows the connectivity motif along with an example of the membrane potential evolution.

As these novel units provide different dynamics than those of the predominantly used LSTM units, we include a more detailed comparison of their dynamics in the supplementary material 1. Note that although the subsequently presented models use only a single type of synapses per neuron, multiple synapse types can be combined, resulting in further enriched neural dynamics.

### 3 RNN-T with biologically-inspired dynamics

The network architecture of RNN-T, illustrated in Figure 2, consists of two main subnetworks, the encoding network and the prediction network, whose outputs are combined to ultimately form the speech transcript. More details of the network architecture are presented in supplementary material 2.

In this work, we redesign the RNN-T architecture, using the novel units introduced in Section 2. In order to do so, we follow a three-step approach. First, the units are introduced into the prediction network, replacing the LSTMs, while the encoding network remains composed of LSTM units. In a second step, we incorporate the sSNU variants into the encoding network only, i.e., the prediction network remains composed of LSTMs. Finally, the sSNU variants are integrated into both network parts and thus the full RNN-T architecture is composed of sSNU units. The encoding network carries out a different task than the prediction network and hence different units might be better suited. Note that a composition of diverse units and synapse types is also observed in the brain.

Speech recognition requires a large amount of compute resources. Thus, reducing the computational cost is of paramount importance. Our proposed units not only reflect the biological inspirations, but

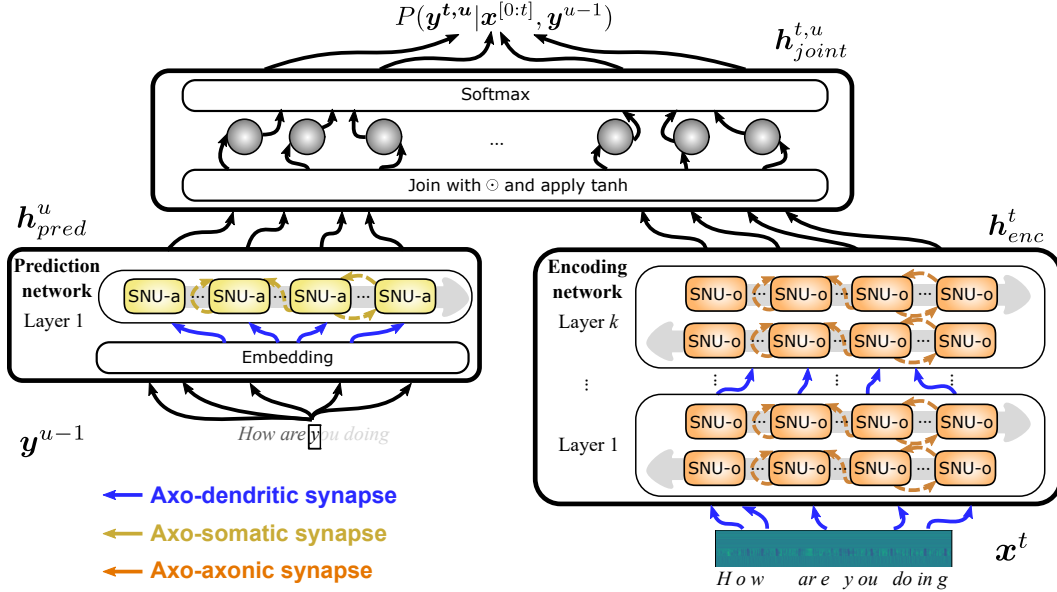


Figure 2: **Illustration of the RNN transducer architecture.** An input vector  $x^t$ , containing MFCC features of a speech signal, is processed by the encoding network, here represented with sSNU-o units. The prediction network, corresponding to a language model, here represented with sSNU-a units, enhances the predictions based on the past output sequence of the RNN-T, i.e.  $y^{u-1}$ . The joint network forms the final predictions based on the outputs from both subnetworks.

Table 1: RNN acronyms used in the result tables and details on the trainable parameters.

RNN <i>Suffix</i>	Comment	Thr.	Axo-dendritic	Axo-somatic	Axo-axonic
sSNU	Feedforward	$\mathbf{b}$	$\mathbf{W}$		
sSNU <i>R</i>	Recurrent	$\mathbf{b}$	$\mathbf{W}, \mathbf{H}$		
sSNU-a	Adaptive thr. feedforward	$\mathbf{b}_o$	$\mathbf{W}$		
sSNU-a <i>R</i>	Adaptive thr. recurrent	$\mathbf{b}_o$	$\mathbf{W}, \mathbf{H}$		
sSNU-a <i>Ra</i>	Adaptive thr. axo-somatic recurrent	$\mathbf{b}_o$	$\mathbf{W}, \mathbf{H}$	$\mathbf{H}_a$	
sSNU-o	Output modulating feedforward	$\mathbf{b}$	$\mathbf{W}$		$\mathbf{W}_o, \mathbf{b}_o$
sSNU-o <i>R</i>	Output modulating recurrent	$\mathbf{b}$	$\mathbf{W}, \mathbf{H}$		$\mathbf{W}_o, \mathbf{H}_o, \mathbf{b}_o$

additionally are simpler in terms of the number of gates and parameters than LSTMs. Hence, they provide the potential to drastically reduce the computational cost and the transcription time.

## 4 Experiments and results

The simulations were performed using the Switchboard speech corpus comprising roughly 300 hours of English telephone conversations. In order to be comparable to the state-of-the-art literature, we closely followed the preprocessing, the evaluation, as well as the network architecture setup presented in [26]. An LSTM-based version of this RNN-T architecture achieves a word error rate (WER) of 12.7%, which we consider as our baseline. We investigated various configurations of the sSNU variants, listed with their abbreviations and trainable parameters in Table 1. In addition, supplementary material 3 provides details about the dataset and its preprocessing, as well as hyperparameter settings.

**Comparison of the word error rate.** The first part of Table 2 summarizes the results of the RNN-T architecture, where the sSNU variants were integrated into the prediction network. Several configurations, e.g., sSNU *R* (12.4% WER), sSNU-a *R* (12.0% WER) and sSNU-o *R* (12.4% WER), outperform the LSTM-based variant. In a second step, we evaluated the performance of the encoding network composed of sSNU-a and sSNU-o units. The best performance is achieved with sSNU-o *R* with 14.7% WER compared to 12.7% WER of the LSTM baseline. Finally, the sSNU units are integrated into both subnetworks so that the RNN-T is solely based on sSNUs, see the last part of Table 2. Consistent with the prior two cases, the RNN-T architecture achieved competitive performance with 16.0% WER and 14.9% WER.

**Evaluation of the inference time.** The inference time is a critical metric for speech recognition that depends on the model architecture and the efficiency of the neural units. In our case, this efficiency is largely dominated by the number of multiplications, including vector-matrix and scalar multiplications, and the number of activations that the units require to compute. Thus, we accumulated the total

Table 2: Performance comparison of the RNN-T network for different configurations of the prediction and the encoding network

Enc. RNN	Pred. RNN	WER (%)	# Param.	(%)	# Multipl.	(%)	$t_{inf}$ (s)	(%)
LSTM	LSTM	12.7	2.39M	100	2.39M	100	2.78	100
	sSNU	15.1	8.45k	< 1	9.2k	< 1	2.71	97
	sSNU <i>R</i>	12.4	0.60M	25	0.60M	25	2.76	99
	sSNU-a	12.1	8.45k	< 1	11.52k	< 1	2.73	98
	sSNU-a <i>R</i>	12.0	0.60M	25	0.60M	25	2.78	100
	sSNU-o	12.6	16.90k	< 1	17.66k	< 1	2.75	98
	sSNU-o <i>R</i>	12.4	1.20M	50	1.20M	50	2.76	99
LSTM	LSTM	12.7	54.20M	100	54.20M	100	2.78	100
sSNU-a <i>Ra</i>		25.2	18.47M	34	18.50M	34	1.23	44
sSNU-o		23.2	17.27M	32	17.28M	32	1.22	44
sSNU-o <i>R</i>		14.7	27.10M	50	27.11M	50	1.76	63
LSTM	LSTM	12.7	56.59M	100	56.58M	100	2.78	100
sSNU-o <i>R</i>	sSNU-a <i>R</i>	16.0	27.70M	49	27.71M	49	1.64	59
sSNU-o <i>R</i>	sSNU-o <i>R</i>	14.9	28.30M	50	28.30M	50	1.66	60




Figure 3: **Comparison of the transcription time of various architectures.** The data points represent the mean time taken to decode utterances of different lengths. The markers indicate the various architectures used and the error bars depict the standard deviation over 10 executions.

time, running on CPUs, for all computations involved during the inference of a single utterance. The last three columns in Table 2 show the number of parameters, the number of multiplications as well as the average time taken for the greedy decoding of an utterance with  $T = 388$  input frames ( $\sim 7.8$ s of audio input).

When incorporating the sSNUs into the prediction network, the number of multiplications of this subnetwork is reduced substantially, see the first part of Table 2. However, this is not fully reflected in the transcription time, because this subnetwork contributes only a small fraction to the total computational cost of the RNN-T. As one can see in the second part of Table 2, when the LSTM units of the encoding network are replaced with sSNUs, the total number of parameters is reduced by more than 50% and this results in a more than  $\sim 40\%$  reduction of the inference time. A similar trend can be observed in the final part of Table 2, where the RNN-T is solely composed of sSNUs. In this case, the RNN-T achieves a competitive WER, albeit using 50% fewer parameters and experiencing a  $\sim 40\%$  reduce inference time.

Figure 3 depicts a comparison of the transcription time to decode utterances of various lengths for selected models. To ensure reproducibility and consistency, the examples were chosen such that the transcripts of all models were identical, and the timing results were averaged over 10 repetitions. A general observation is that the time taken to transcribe utterances increases proportionally to the utterance length. The RNN-T architecture composed of LSTM units is indicated with red dots and is the slowest among the evaluated models. Depending on the sSNU variant used, the inference time can be reduced. For example, the RNN-T using sSNU-o units for the encoding and sSNU-o units for the prediction network (6x640 sSNU-o R, 1x768 sSNU-o R), has an approximately 40% reduced inference time.

## 5 Conclusions

In this work we combine insights from neuroscience with a state-of-the-art network architecture from machine learning and apply it to the task of speech recognition. We resort to the diverse types of synapses present in the brain and enhance the dynamics of the commonly used LIF neuron model with a threshold adaptation mechanism based on axo-somatic synapses as well as with an output modulating mechanism based on the axo-axonic synapses. We successively integrate those novel units into the RNN-T architecture and not only demonstrate that end-to-end speech recognition with biologically-inspired units is possible, but also that they achieve a competitive performance level of 14.9% WER, compared to 12.7% WER of the LSTM-based RNN-T. Moreover, the introduced units are simpler than LSTMs, so that the computational costs, as well as the inference time can be reduced by 50% and 40% respectively. Finally, it is worth mentioning that biology provides an abundance of mechanisms, which could further enrich the dynamics of neurons, and that are not yet covered in practical neural networks. Our simulation results indicate that such effects can potentially bolster the performance of biologically-inspired neural units and are therefore an essential step going forward.

## References

- [1] J. Baker, “The dragon system—an overview,” *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 23, no. 1, pp. 24–29, 1975.
- [2] M. Gales and S. Young, “The application of hidden markov models in speech recognition,” *Found. Trends Signal Process.*, vol. 1, p. 195–304, Jan. 2007.
- [3] M. Benzeghiba, R. De Mori, O. Deroo, S. Dupont, T. Erbes, D. Jouvet, L. Fissore, P. Laface, A. Mertins, C. Ris, R. Rose, V. Tyagi, and C. Wellekens, “Automatic speech recognition and speech variability: A review,” *Speech Communication*, vol. 49, pp. 763–786, Oct 2007.
- [4] D. Yu and L. Deng, *Automatic Speech Recognition*. London, England, UK: Springer, 2015.
- [5] A. Graves, D. Eck, N. Beringer, and J. Schmidhuber, “Biologically plausible speech recognition with LSTM neural nets,” in *Biologically Inspired Approaches to Advanced Information Technology* (A. J. Ijspeert, M. Murata, and N. Wakamiya, eds.), (Berlin, Heidelberg), pp. 127–136, Springer Berlin Heidelberg, 2004.
- [6] B.-H. Juang and L. R. Rabiner, “Automatic speech recognition—a brief history of the technology development,” *Georgia Institute of Technology. Atlanta Rutgers University and the University of California. Santa Barbara*, vol. 1, p. 67, 2005.
- [7] W. Xiong, J. Droppo, X. Huang, F. Seide, M. Seltzer, A. Stolcke, D. Yu, and G. Zweig, “The Microsoft 2016 conversational speech recognition system,” *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5255–5259, Mar 2017.
- [8] T.-S. Nguyen, S. Stueker, and A. Waibel, “Super-Human Performance in Online Low-latency Recognition of Conversational Speech,” *arXiv*, Oct 2020.
- [9] D. Wang, X. Wang, and S. Lv, “An Overview of End-to-End Automatic Speech Recognition,” *Symmetry*, vol. 11, p. 1018, Aug 2019.
- [10] G. Rao, “A Comparison of End-to-End Speech Recognition Architectures in 2021,” January 2021. [Online; accessed via <https://www.assemblyai.com/blog/a-survey-on-end-to-end-speech-recognition-architectures-in-2021> on 6. May 2021].
- [11] A. Graves, “Sequence Transduction with Recurrent Neural Networks,” *arXiv*, Nov 2012.
- [12] W. Chan, N. Jaitly, Q. Le, and O. Vinyals, “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4960–4964, 2016.
- [13] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu, and R. Pang, “Conformer: Convolution-augmented Transformer for Speech Recognition,” in *Proc. Interspeech 2020*, pp. 5036–5040, 2020.
- [14] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang, Q. Liang, D. Bhatia, Y. Shangguan, B. Li, G. Pundak, K. C. Sim, T. Bagby, S.-y. Chang, K. Rao, and A. Gruenstein, “Streaming End-to-end Speech Recognition For Mobile Devices,” *arXiv*, Nov 2018.
- [15] T. Sainath, R. Pang, D. Rybach, Y. He, R. Prabhavalkar, W. Li, M. Visontai, Q. Liang, T. Strohman, Y. Wu, I. McGraw, and C.-C. Chiu, “Two-pass end-to-end speech recognition,” pp. 2773–2777, 09 2019.
- [16] K. Hu, T. N. Sainath, R. Pang, and R. Prabhavalkar, “Deliberation Model Based Two-Pass End-to-End Speech Recognition,” *arXiv*, Mar 2020.
- [17] K. Hu, R. Pang, T. N. Sainath, and T. Strohman, “Transformer based deliberation for two-pass speech recognition,” in *2021 IEEE Spoken Language Technology Workshop (SLT)*, pp. 68–74, 2021.

- [18] J. Dennis, Q. Yu, H. Tang, H. D. Tran, and H. Li, “Temporal coding of local spectrogram features for robust sound recognition,” in *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 803–807, 2013.
- [19] J. Wu, Y. Chua, and H. Li, “A biologically plausible speech recognition framework based on spiking neural networks,” in *2018 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, 2018.
- [20] L. Dong and B. Xu, “CIF: Continuous Integrate-and-Fire for End-to-End Speech Recognition,” *arXiv*, May 2019.
- [21] R. Güttig and H. Sompolinsky, “The tempotron: a neuron that learns spike timing–based decisions - Nature Neuroscience,” *Nat. Neurosci.*, vol. 9, pp. 420–428, Mar 2006.
- [22] C. Shi, T. Wang, J. He, J. Zhang, L. Liu, and N. Wu, “DeepTempo: A Hardware-Friendly Direct Feedback Alignment Multi-Layer Tempotron Learning Rule for Deep Spiking Neural Networks,” *IEEE Trans. Circuits Syst. II*, vol. 68, pp. 1581–1585, Mar 2021.
- [23] S. Woźniak, A. Pantazi, T. Bohnstingl, and E. Eleftheriou, “Deep learning incorporating biologically inspired neural dynamics and in-memory computing,” *Nat. Mach. Intell.*, vol. 2, pp. 325–336, Jun 2020.
- [24] D. P. Kingma and J. Ba, “Adam: A Method for Stochastic Optimization,” *arXiv*, Dec 2014.
- [25] I. Loshchilov and F. Hutter, “Decoupled Weight Decay Regularization,” *arXiv*, Nov 2017.
- [26] G. Saon, Z. Tueske, D. Bolanos, and B. Kingsbury, “Advancing RNN Transducer Technology for Speech Recognition,” *arXiv*, Mar 2021.
- [27] A. Howard, G. Tamas, and I. Soltesz, “Lighting the chandelier: new vistas for axo-axonic cells,” *Trends Neurosci.*, vol. 28, pp. 310–316, Jun 2005.
- [28] M. Bear, B. Connors, and M. A. Paradiso, *Neuroscience: Exploring the brain*. Jones & Bartlett Learning, LLC, 2020.
- [29] A. I. for Brain Science, “Allen Human Brain Atlas,” 2010.
- [30] S. Woźniak, A. Pantazi, T. Bohnstingl, and E. Eleftheriou, “Deep learning incorporating biologically-inspired neural dynamics,” *arXiv*, Dec 2018.
- [31] G. Bellec, D. Salaj, A. Subramoney, R. Legenstein, and W. Maass, “Long short-term memory and Learning-to-learn in networks of spiking neurons,” in *NeurIPS*, pp. 787–797, 2018.
- [32] B. Fontaine, J. L. Peña, and R. Brette, “Spike-Threshold Adaptation Predicted by Membrane Potential Dynamics In Vivo,” *PLoS Comput. Biol.*, vol. 10, p. e1003560, Apr 2014.
- [33] C. Huang, A. Resnik, T. Celikel, and B. Englitz, “Adaptive Spike Threshold Enables Robust and Temporally Precise Neuronal Encoding,” *PLoS Comput. Biol.*, vol. 12, p. e1004984, Jun 2016.
- [34] R. Tremblay, S. Lee, and B. Rudy, “Gabaergic interneurons in the neocortex: from cellular properties to circuits,” *Neuron*, vol. 91, no. 2, pp. 260–292, 2016.
- [35] G. Fishell and B. Rudy, “Mechanisms of Inhibition within the Telencephalon: “Where the Wild Things Are”,” *Annu. Rev. Neurosci.*, vol. 34, pp. 535–567, Jun 2011.