



Deep Spiking Neural Networks for Large Vocabulary Automatic Speech Recognition

Jibin Wu^{1*}, Emre Yilmaz¹, Malu Zhang¹, Haizhou Li^{1,2} and Kay Chen Tan³

¹ Department of Electrical and Computer Engineering, National University of Singapore, Singapore, Singapore, ² Faculty for Computer Science and Mathematics, University of Bremen, Bremen, Germany, ³ Department of Computer Science, City University of Hong Kong, Kowloon Tong, Hong Kong

Artificial neural networks (ANN) have become the mainstream acoustic modeling technique for large vocabulary automatic speech recognition (ASR). A conventional ANN features a multi-layer architecture that requires massive amounts of computation. The brain-inspired spiking neural networks (SNN) closely mimic the biological neural networks and can operate on low-power neuromorphic hardware with spike-based computation. Motivated by their unprecedented energy-efficiency and rapid information processing capability, we explore the use of SNNs for speech recognition. In this work, we use SNNs for acoustic modeling and evaluate their performance on several large vocabulary recognition scenarios. The experimental results demonstrate competitive ASR accuracies to their ANN counterparts, while require only 10 algorithmic time steps and as low as 0.68 times total synaptic operations to classify each audio frame. Integrating the algorithmic power of deep SNNs with energy-efficient neuromorphic hardware, therefore, offer an attractive solution for ASR applications running locally on mobile and embedded devices.

OPEN ACCESS

Edited by:

Huajin Tang,
Zhejiang University, China

Reviewed by:

Federico Corradi,
Imec, Netherlands
Juan Pedro Dominguez-Morales,
University of Seville, Spain

*Correspondence:

Jibin Wu
jibin.wu@u.nus.edu

Specialty section:

This article was submitted to
Neuromorphic Engineering,
a section of the journal
Frontiers in Neuroscience

Received: 19 November 2019

Accepted: 24 February 2020

Published: 17 March 2020

Citation:

Wu J, Yilmaz E, Zhang M, Li H and
Tan KC (2020) Deep Spiking Neural
Networks for Large Vocabulary
Automatic Speech Recognition.
Front. Neurosci. 14:199.
doi: 10.3389/fnins.2020.00199

Keywords: deep spiking neural networks, automatic speech recognition, tandem learning, neuromorphic computing, acoustic modeling

1. INTRODUCTION

Automatic speech recognition (ASR) has enabled the voice interface of mobile devices and smart home appliances in our everyday life. The rapid progress in the integration of voice interfaces has been viable on account of the remarkable performance of the ASR systems using artificial neural networks (ANN) for acoustic modeling (Lippmann, 1989; Lang et al., 1990; Hinton et al., 2012; Yu and Deng, 2015). Various ANN architectures, either feedforward or recurrent, have been investigated for modeling the acoustic information preserved in speech signals (Dahl et al., 2012; Graves et al., 2013; Abdel-Hamid et al., 2014).

The performance gains come with immense computational requirements often due to the time-synchronous processing of input audio signals. Several techniques have been proposed to reduce the computational load and memory storage of ANNs by reducing the number of parameters that have to be used for inference (Sainath et al., 2013; Xue et al., 2013; He et al., 2014; Povey et al., 2018). Another common solution, for reducing the processing load, uses a wake word or phrase to control the access to speech recognition services (Zehetner et al., 2014; Sainath and Parada, 2015; Wu M. et al., 2018). Moreover, most devices with voice control rely on cloud-based ASR engines rather than local on-device solutions. The necessity of online processing of speech via cloud computing comes with various concerns, such as data security and processing speed, etc. There have been multiple

efforts to develop on-device ASR solutions in which the speech signal is processed locally using the computational resources of mobile devices (Lei et al., 2013; McGraw et al., 2016).

Alternatively, event-driven models such as spiking neural networks (SNNs) inspired by the human brain have attracted ever-growing attention in recent years. The human brain is remarkably efficient and capable of performing complex perceptual and cognitive tasks. Notably, the adult's brain only consumes about 20 watts to solve complex tasks that are equivalent to the power consumption of a dim light bulb (Laughlin and Sejnowski, 2003). While brain-inspired ANNs have demonstrated great capabilities in many perceptual (He et al., 2016; Xiong et al., 2017) and cognitive tasks (Silver et al., 2017), these models are computationally intensive and memory inefficient to operate as compared to the biological brains. Unlike ANNs, asynchronous and event-driven information processing of SNNs resembles the computing paradigm that observed in the human brains, whereby the energy consumption matches the activity levels of sensory stimuli. Given temporally sparse information transmitted in the surrounding environment, the event-driven computation, therefore, exhibits great computational efficiency than the synchronous computation used in ANNs.

Neuromorphic computing (NC), as a non-von Neumann computing paradigm, mimics the event-driven computation of the biological neural systems with SNN in silicon. The emerging neuromorphic computing architectures (Furber et al., 2012; Merolla et al., 2014; Davies et al., 2018) leverage on the massively parallel, low-power computing units to support spike-based information processing. Notably, the design of co-located memory and computing units effectively circumvents the von Neumann bottleneck of low-bandwidth between memory and the processing units (Monroe, 2014). Therefore, integrating the algorithmic power of deep SNNs with the compelling energy efficiency of NC hardware represents an intriguing solution for pervasive machine learning tasks and always-on applications. Furthermore, growing research efforts are devoted to developing novel non-volatile memory devices for ultra-low-power implementation of biological synapses and neurons (Tang et al., 2019).

Some preliminary work on SNN-based phone classification or small-vocabulary speech recognition systems have been explored in Jim-Shih Liaw and Berger (1998), Näger et al. (2002), Loisel et al. (2005), Holmberg et al. (2005), Kröger et al. (2009), Tavanaei and Maida (2017a,b), Wu et al. (2018a), Wu et al. (2018b), Zhang et al. (2015), Zhang et al. (2019), Bellec et al. (2018), Wu et al. (2019b), and Pan et al. (2018). However, these SNN-based ASR systems are far from the scale and complexity of modern commercialized ANN-based ASR systems. It is mainly due to lacking effective training algorithms for deep SNNs and efficient software toolbox for SNN-based ASR systems.

Due to the discrete and non-differentiable nature of spike generation, the powerful error back-propagation algorithm is not directly applicable to the training of deep SNNs. Recently, considerable research efforts are devoted to addressing this problem and the resulting learning rules can be broadly categorized into the SNN-to-ANN conversion (Cao et al., 2015; Diehl et al., 2015), back-propagation through time with surrogate

gradient (Wu Y. et al., 2018; Neftci et al., 2019; Wu et al., 2019a) and tandem learning (Wu et al., 2019c). Despite several successful attempts on the large-scale image classification tasks with deep SNNs (Rueckauer et al., 2017; Hu et al., 2018; Sengupta et al., 2019; Wu et al., 2019c), their applications to the large-vocabulary continuous ASR (LVCSR) tasks remain unexplored. In this work, we explore an SNN-based acoustic model for LVCSR using a recently proposed tandem learning rule (Wu et al., 2019c) that supports an efficient and rapid inference.

To summarize, the main contributions of this work are threefold:

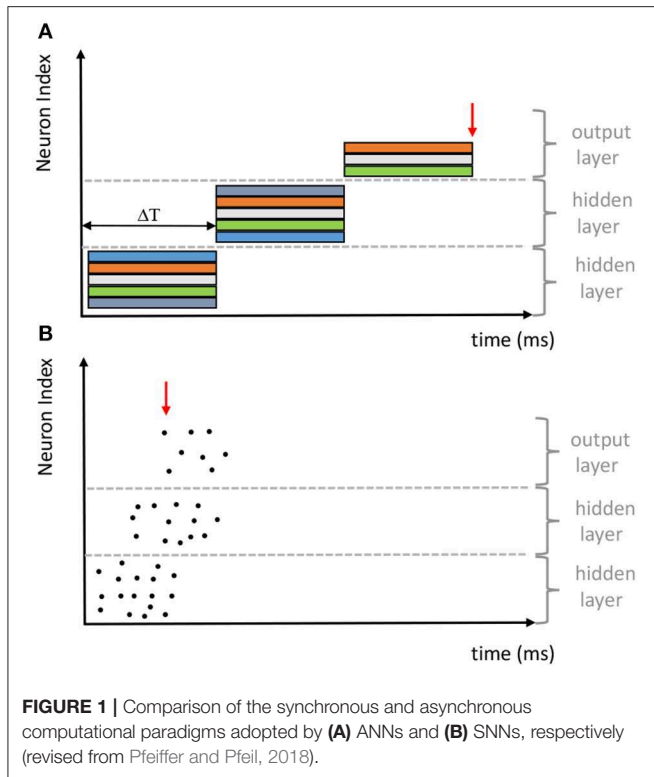
- **Large-Vocabulary Automatic Speech Recognition with SNNs.** We explored the SNN-based acoustic models for large-vocabulary automatic speech recognition tasks. The SNN-based ASR systems achieved competitive accuracy on par with their ANN counterparts across the phone recognition, low-resourced ASR and large-vocabulary ASR tasks. To the best of our knowledge, this is the first work that successfully applied SNNs to the LVCSR task.
- **Toward Rapid and Energy-Efficient Speech Recognition.** Our preliminary study of an SNN-based acoustic model has revealed compelling prospect of rapid inference and unprecedented energy efficiency of a neuromorphic approach. Specifically, SNNs can classify each audio frame accurately with only 10 algorithmic time steps while require as low as 0.68 times total synaptic operations to their ANN counterparts.
- **SNN-Based ASR Toolkit.** We demonstrate that SNN-based acoustic models can be effectively developed in PyTorch and easily integrated into the PyTorch-Kaldi Speech Recognition Toolkit (Ravanelli et al., 2019) for rapid development of SNN-based ASR systems.

The rest of the paper is organized as follows: In section 2, we first give an overview of spiking neural networks, large vocabulary ASR systems, and existing SNN-based ASR systems. In section 3, we introduce the spiking neuron model and the neural coding scheme that converts acoustic features into spike-based representation. We further present a recently introduced tandem learning framework for SNN training and how it is used to train deep SNN-based acoustic models. In section 4, we present experimental results on the learning capability and energy efficiency of SNN-based acoustic models across three different types of recognition tasks including phone recognition, low-resourced and standard large-vocabulary ASR, and compare those to the ANN-based implementations. Finally, a discussion on the experimental findings is given in section 5.

2. FUNDAMENTALS AND RELATED WORK

2.1. Spiking Neural Networks

The third generation spiking neural networks are originally studied as models to describe the information processing in the biological neural networks, wherein the information is communicated and exchanged via stereotypical action potentials or spikes (Gerstner and Kistler, 2002). Neuroscience studies reveal that the temporal structure and frequency of these spike

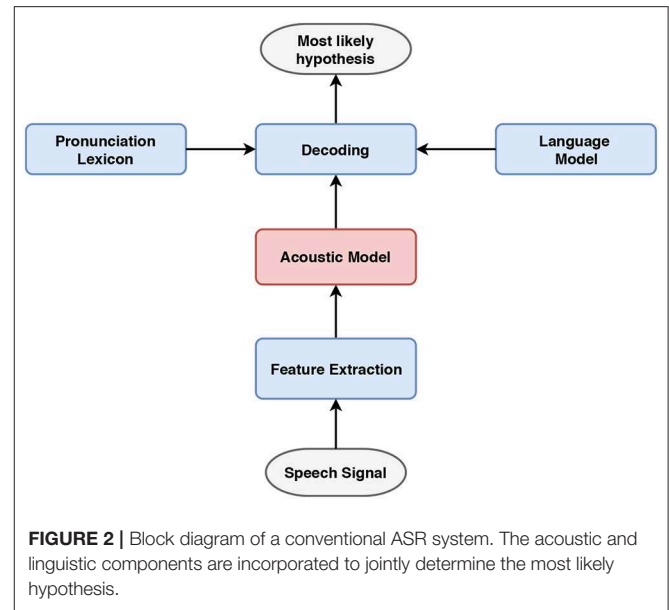


trains are both important information carriers in the biological neural networks. As will be introduced in section 3.1, the spiking neuron operates asynchronously and integrates the synaptic current from its incoming spike trains. An output spike is generated from the spiking neuron whenever its membrane potential crosses the firing threshold, and this output spike will be propagated to the connected neurons via the axon.

Motivated by the same connectionism principle, SNNs share the same network architectures, either feedforward or recurrent, with the conventional ANNs that use analog neurons. As shown in **Figure 1**, the early classification decision can be made from the SNN since the generation of the first output spike. However, the quality of the classification decision is typically improved over time with more evidence accumulated. It differs significantly from the synchronous information processing of the conventional ANNs, where the output layer needs to wait until all preceding layers are fully updated. Therefore, despite information is transmitted and processed at a speed that is several orders of magnitude slower in neural substrates than signal processing in modern transistors, biological neural systems can perform complex tasks rapidly. For more overviews about SNNs and their applications, we refer readers to Pfeiffer and Pfeil (2018) and Tavanaei et al. (2019).

2.2. Large Vocabulary Automatic Speech Recognition

As shown in **Figure 2**, conventional ASR systems use acoustic and linguistic information preserved in three distinct components to convert speech signals to the corresponding text: (1) an acoustic model for preserving the statistical



representations of different speech units, e.g., phones, from speech features, (2) a language model for assigning probabilities to the co-occurring word sequences and (3) a pronunciation lexicon for mapping the phonetic transcriptions to orthography. These resources are jointly used to determine the most likely hypothesis in the decoding stage.

Acoustic modeling can be achieved by using various statistical models such as Gaussian Mixture Models (GMM) for assigning frame-level phone posteriors in conjunction with a Hidden Markov Model (HMM) for duration modeling (Yu and Deng, 2015). More recently, ANN-based approaches have become the standard acoustic models providing state-of-the-art performance across a wide spectrum of ASR tasks (Hinton et al., 2012). Together with numerous ANN architectures explored for acoustic modeling, several end-to-end ANN architectures have been proposed for directly mapping speech features to text with optional use of the other linguistic components (Graves and Jaitly, 2014; Chan et al., 2016; Watanabe et al., 2017).

The probabilistic definition of acoustic modeling becomes more evident via the Bayesian formulation of the speech recognition task. Given a target speech signal that segmented into T overlapped frames, the resulting frame-wise features can be represented as $\mathbf{O} = [\mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_T]$. An ASR system assigns the probability $P(\mathbf{W}|\mathbf{O})$ to all possible word sequences $\mathbf{W} = [w_1, w_2, \dots]$, and the word sequence $\hat{\mathbf{W}}$ with the highest probability is the recognized output,

$$\hat{\mathbf{W}} = \arg \max_{\mathbf{W}} P(\mathbf{W}|\mathbf{O}) \quad (1)$$

The probability $P(\mathbf{W}|\mathbf{O})$ can be decomposed into two parts by applying the Bayes' rule as below,

$$\hat{\mathbf{W}} = \arg \max_{\mathbf{W}} \frac{P(\mathbf{O}|\mathbf{W})P(\mathbf{W})}{P(\mathbf{O})} \quad (2)$$

$P(\mathbf{O})$ can be omitted as it does not depend on \mathbf{W} . This results in

$$\hat{\mathbf{W}} = \arg \max_{\mathbf{W}} P(\mathbf{O}|\mathbf{W})P(\mathbf{W}) \quad (3)$$

which formally defines the theoretical foundation of the conventional ASR systems. $P(\mathbf{W})$ is the prior probability of the word sequence \mathbf{W} and this probability is provided by the language model which is trained on a large written corpus of the target language. $P(\mathbf{O}|\mathbf{W})$ is the *likelihood* of the observed feature sequence \mathbf{O} given the word sequence \mathbf{W} , and this probability is associated with the acoustic model. The acoustic model captures the information about the acoustic component of speech signals, aiming to classify different acoustic units accurately. Traditionally, each phone in the phonetic alphabet is modeled using multiple three-state HMM models for different preceding and following phonetic context (triphone) (Lee, 1990). The emission probability of these HMM states are shared (tied) among different models to reduce the number of model parameters (Hwang and Huang, 1993). The output layer of the ANN-based acoustic model is designed accordingly and trained to assign these frame-level tied triphone HMM state (senone) probabilities (Dahl et al., 2012). The output layer uses the softmax function to normalize the output into a probability distribution. These values are scaled with the prior probabilities of each class, obtained from the training data, to determine the likelihood values. These likelihood values are later combined with the probabilities assigned by the language model during the decoding stage so as to find the most likely hypothesis.

Speech features, used as the inputs to the acoustic model, describe the spectrotemporal dynamics of the speech signal and discriminate among different phones in the target language. Mel-frequency cepstral coefficients (MFCC) (Davis and Mermelstein, 1980) features are commonly used in conjunction with the GMM-HMM acoustic model. The MFCC features are extracted by (1) performing short-time Fourier transform, (2) applying triangular Mel-scaled filter banks to calculate the power at each Mel frequency in log domain (FBANK) and (3) performing a discrete cosine transform to decorrelate the FBANK features. The third step is often skipped and FBANK features are often used when training ANN-based acoustic models since these models can handle correlation among features. In this work, we incorporate deep SNNs for acoustic modeling instead of the conventional ANNs and compare their ASR performance in different ASR scenarios including phone recognition, low-resourced and standard large vocabulary ASR. The ASR performance obtained using popular speech features have been reported to explore the impact of the feature representation space and its dimensionality for SNN-based acoustic models.

2.3. Speech Recognition With Spiking Neural Network

SNNs are well-suited for representing and processing spatial-temporal signals, they hence possess great potentials for speech recognition tasks. Tavanaei and Maida (2017a,b) proposed SNN-based feature extractors to extract discriminative features

from the raw speech signal using unsupervised spiking-timing-dependent plasticity (STDP) rule. While connecting these SNN-based feature extractors with Support Vector Machine (SVM) or Hidden Markov Model (HMM) classifiers, competitive classification accuracies were demonstrated on the isolated spoken digit recognition task. Wu et al. (2018a,b) introduced a SOM-SNN framework for environmental sound and speech recognition. In this framework, the biological-inspired self-organizing map (SOM) is utilized for feature representation, which maps frame-based acoustic features into a spike-based representation that is both sparse and discriminative. The temporal dynamic of the speech signal is further handled by the SNN classifier. Zhang et al. (2019) presented a fully SNN-based speech recognition framework, wherein the spectral information of consecutive frames are encoded with threshold coding and subsequently classified by the SNN that is trained with a novel membrane potential-driven aggregate-labeling learning algorithm.

Recurrent network of spiking neurons (RSNNs) exhibit greater memory capacity than the aforementioned feedforward frameworks. They can capture long temporal information that are useful for speech recognition tasks. In Zhang et al. (2015), Zhang et al. presented a spiking liquid-state machine (LSM) speech recognition framework which is attractive for low-power very-large-scale-integration (VLSI) implementation. Bellec et al. recently demonstrated state-of-the-art phone recognition accuracy on the TIMIT dataset by adding neuronal adaptation mechanism to the vanilla RSNNs (Bellec et al., 2018). It is the first time that RSNNs approaching the performance of LSTM networks (Greff et al., 2016) on the speech recognition task. These preliminary works on the SNN-based ASR systems are however limited to the phone classification or small vocabulary isolated spoken digit recognition tasks. In this work, we apply deep SNNs to LVCSR tasks and demonstrate competitive accuracies over the ANN-based ASR systems.

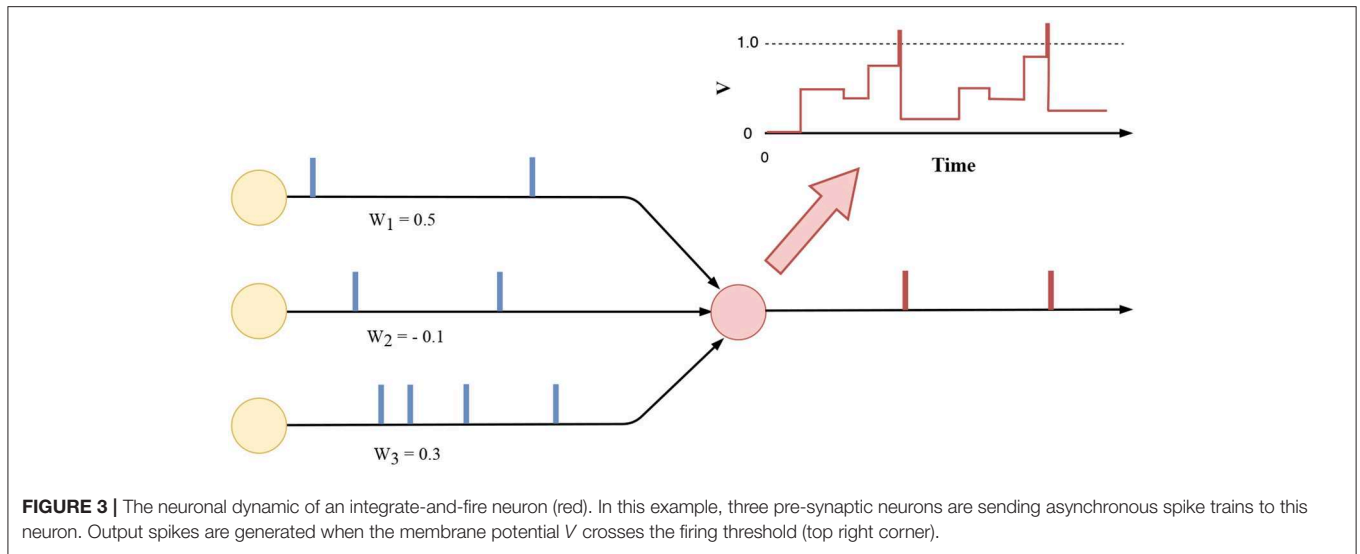
3. METHODS

3.1. Spiking Neuron Model

As shown in **Figure 4**, the frame-based features are first extracted and input into the SNN-based acoustic models. Given the short temporal duration of segmented frames and the slow variation of speech signals, these features are typically assumed to be stationary over the short time-period of segmented frames. In this work, we use the integrate-and-fire (IF) neuron model with reset by subtraction scheme (Rueckauer et al., 2017), which can effectively process these stationary frame-based features with minimal computational costs. Although IF neurons do not emulate rich temporal dynamics of biological neurons, they are however ideal for working with the neural representation that employed in this work, where spike timings play an insignificant role.

At each time step t of a discrete-time simulation, with a total number of time steps N_s , the incoming spikes to neuron j at layer l are transduced into synaptic current as follows

$$z_j^l(t) = \sum_i w_{ji}^{l-1} \cdot \theta_i^{l-1}(t) + b_j^l \quad (4)$$



where $\theta_i^{l-1}(t)$ indicates the occurrence of an input spike from afferent neuron i at time step t . In addition, the w_{ji}^{l-1} denotes the synaptic weight that connects presynaptic neuron i from layer $l - 1$. Here, b_j^l can be interpreted as a constant injecting current. As shown in **Figure 3**, neuron j integrates the input current $z_j^l(t)$ into its membrane potential $V_j^l(t)$ as per Equation (5). The $V_j^l(0)$ is reset and initialized to zero for every new frame-based feature input. Without loss of generality, a unitary membrane resistance is assumed here. An output spike is generated whenever $V_j^l(t)$ crosses the firing threshold ϑ (Equation 6), which we set to a value of 1 for all the experiments by assuming that all synaptic weights are normalized with respect to the ϑ .

$$V_j^l(t) = V_j^l(t - 1) + z_j^l(t) - \vartheta \cdot \theta_j^l(t - 1) \quad (5)$$

$$\theta_j^l(t) = \Theta(V_j^l(t) - \vartheta) \text{ with } \Theta(x) = \begin{cases} 1, & \text{if } x \geq 0 \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

According to Equations (4) and (5), the free aggregated membrane potential of neuron j (no firing) in layer l can be expressed as

$$V_j^{lf} = \sum_i w_{ji}^{l-1} \cdot c_i^{l-1} + b_j^l \cdot N_s \quad (7)$$

where c_i^{l-1} is the input spike count from pre-synaptic neuron i at layer $l - 1$ as per Equation (8).

$$c_i^{l-1} = \sum_{t=1}^{N_s} \theta_i^{l-1}(t). \quad (8)$$

The V_j^{lf} summarizes the aggregate membrane potential contributions of the incoming spikes from pre-synaptic neurons while ignoring their temporal structures. As will be explained in the tandem learning framework section, this intermediate quantity links the SNN layers to the coupled ANN layers for parameter optimization.

3.2. Neural Coding Scheme

SNNs process information transmitted via spike trains, therefore, special mechanisms are required to encode the continuous-valued feature vectors into spike trains and decode the classification results from the activity of output neurons. To this end, we adopt the spiking neural encoding scheme that proposed in Wu et al. (2019c). This encoding scheme first transforms frame-based input feature vector X^0 (e.g., MFCC or FBANK features), where $X^0 = [x_1^0, x_2^0, \dots, x_n^0]^T$, through a weighted layer of rectified linear unit (ReLU) neurons as follows

$$V_j^{0f}(0) \equiv a_j^0 = \rho(\sum_i w_{ji}^0 \cdot x_i^0 + b_j^0) \quad (9)$$

where w_{ji}^0 is the strength of the synaptic connection between the input x_i^0 and ReLU neuron j . The b_j^0 is the corresponding bias term of the neuron j , and $\rho(\cdot)$ denotes the ReLU activation function. The free aggregate membrane potential $V_j^{0f}(0)$ is defined to be equal to the activation value a_j^0 of the ReLU neuron j . We distribute this quantity over the encoding time window N_s and represent it via spike trains as per Equations (10) and (11).

$$\theta_j^0(t) = \Theta(V_j^{0f}(t - 1) - \vartheta) \quad (10)$$

$$V_j^{0f}(t) = V_j^{0f}(t - 1) - \vartheta \cdot \theta_j^0(t) \quad (11)$$

Altogether, the spike train s^0 and spike count c^0 that output from the neural encoding layer can be represented as follows

$$s^0 = \{\theta^0(1), \dots, \theta^0(N_s)\} \quad (12)$$

$$c^0 = \sum_{t=1}^{N_s} \theta^0(t) \quad (13)$$

This encoding layer performs weighted transformation inside an end-to-end learning framework. It transforms the original

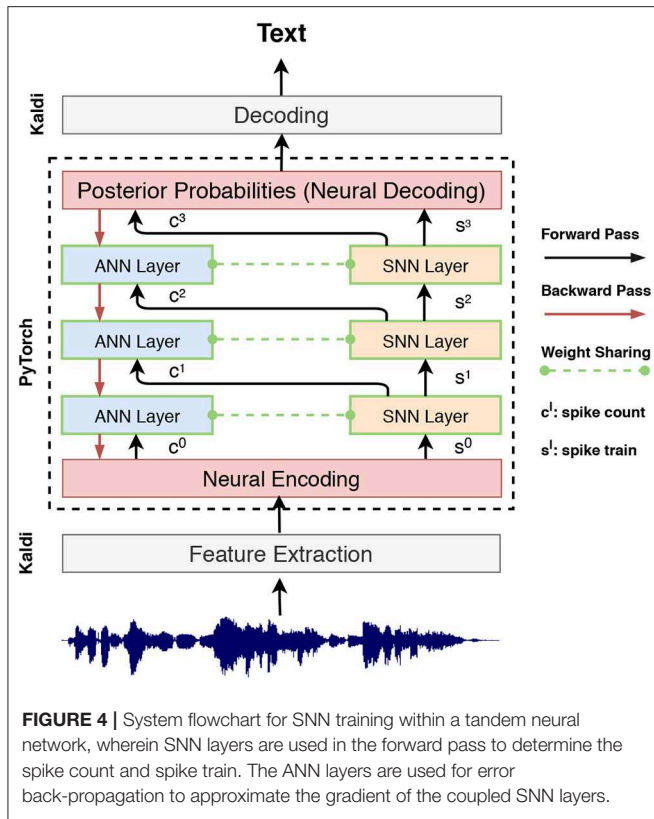


FIGURE 4 | System flowchart for SNN training within a tandem neural network, wherein SNN layers are used in the forward pass to determine the spike count and spike train. The ANN layers are used for error back-propagation to approximate the gradient of the coupled SNN layers.

input representation to match the size of the encoding time window N_s and represents the transformed information via spike counts. This encoding scheme is beneficial for rapid inference since the input information can be effectively encoded within a short encoding window. Start from this neural encoding layer, as shown in **Figure 4**, we input the spike count c^l and s^l to subsequent ANN and SNN layers for tandem learning.

To ensure smooth learning with high precision error gradients derived at the output layer, we use the free aggregate membrane potential of output spiking neurons for neural decoding. Considering that the dimensionality of input feature vectors and output classes are much smaller than that of hidden layers, the computation required will be limited when deploying these two layers onto the edge devices.

3.3. Tandem Learning for Training Deep SNNs

Here, we present a recently proposed SNN learning rule, under the tandem neural network configuration, that exploits a connection between the activation value of ANN neurons and the spike count of IF neurons. As the input features are effectively encoded as spike counts, the temporal structure of the spike trains carries negligible information. The effective non-linear transformation of SNN layers therefore can be summarized as

$$c_j^l = f(s^{l-1}; w_j^{l-1}, b_j^l) \quad (14)$$

where $f()$ denotes the transformation performed by spiking neurons. However, due to the state-dependent nature of spike generation, it is not viable to determine an analytical expression from s^{l-1} to c_j^l directly. Therefore, we simplify the spike generation process by assuming the resulting synaptic currents from s^{l-1} are evenly distributed over the encoding time window. As such, the interspike interval can be determined as follows

$$ISI_j^l = \rho \left(\frac{\vartheta}{V_j^{lf} / N_s} \right) = \rho \left(\frac{\vartheta}{(\sum_i w_{ji}^{l-1} c_i^{l-1} + b_j^l \cdot N_s) / N_s} \right) \quad (15)$$

Hence, the approximated “spike count” a_j^l can be derived according to

$$a_j^l = \frac{N_s}{ISI_j^l} = \frac{1}{\vartheta} \cdot \rho \left(\sum_i w_{ji}^{l-1} c_i^{l-1} + b_j^l \cdot N_s \right) \quad (16)$$

Given a unitary firing threshold ϑ , a_j^l can be effectively determined from an ANN layer of ReLU neurons by setting the spike count c_i^{l-1} as the input and the aggregated constant injecting current $b_j^l \cdot N_s$ as the bias term. This simplification of spike generation process allows the spike-train level error gradients to be approximated from the ANN layer. Wu et al. (2019c) have revealed that the cosine distances between the approximated ‘spike count’ a^l and the actual SNN output spike count c^l are exceedingly small in a high dimensional space, suggesting high quality error gradients can be approximated from the coupled ANN layers.

Based on this formulation, we constructed tandem neural networks as shown in **Figure 4**. During the activation forward propagation, the SNN layers are used to determine the exact spike representation which then propagate the aggregate spike counts and spike trains to the subsequent ANN and SNN layers, respectively. This interlaced layer structure ensures the information that forward propagated to the coupled ANN and SNN layers are synchronized. It worth noting that the ANN is just an auxiliary structure to facilitates the training of SNN, while only SNN is used during inference. The details of this tandem learning rule are provided in the Algorithm 1.

3.4. SNN-Based Acoustic Modeling

To train the deep SNN-based acoustic models, which is the main contribution of this work, several popular speech features have been extracted from the training recordings as described in section 2.2. Before being fed into the SNNs, these input speech features are contextualized by splicing multiple frames so as to exploit more temporal context information. Before training the SNN-based acoustic model, alignments of the speech features with the target senone labels are obtained using a conventional GMM-HMM-based ASR system similar to that described in Dahl et al. (2012). These frame-level alignments enable the training of the deep SNN acoustic model with the tandem learning approach. During the training, the deep SNN learns to map input speech features to posterior probabilities of senones (cf.

Algorithm 1: Pseudo Codes For The Tandem Learning Rule

Input: Input frame-based feature vectors X^0 , target label Y , network parameters w , neural encoding window size N_s

Output: Updated network parameters w

Forward Pass:

$c^0, s^0 = \text{Neural_Encoding}(X^0)$

for layer $l = 1$ to $N-1$ **do**

 // State Update of the ANN Layer

$a^l = \text{ANN.layer}[l].\text{forward}(c^{l-1}, w^{l-1}) *$

for $t = 1$ to N_s **do**

 // State Update of the SNN Layer

$s^l[t] = \text{SNN.layer}[l].\text{forward}(s^{l-1}[t], w^{l-1})$

 // Update the Spike Count

$c^l = \sum_{t=1}^{N_s} s^l[t]$

 /* Neural Decoding with the Aggregate Membrane Potential */

$output = \text{ANN.layer}[N].\text{forward}(c^{N-1}, w^{N-1})$

Loss: $E = \text{LossFunction}(Y, output)$

Backward Pass:

$\frac{\partial E}{\partial a^N} = \text{LossGradient}(Y, output)$

for layer $l = N-1$ to 1 **do**

 // Gradient Update through the ANN Layer

$\frac{\partial E}{\partial a^{l-1}}, \frac{\partial E}{\partial w^{l-1}} = \text{ANN.layer}[l].\text{backward}(\frac{\partial E}{\partial a^l}, c^{l-1}, w^{l-1})$

Update parameters of the ANN layer based on the calculated gradients.

Copy the updated parameters to the corresponding SNN layer.

Note:

* For inference, state updates are performed on the SNN layers entirely.

section 2.2) by passing the input speech frames through multiple layers of spiking neurons.

During the inference phase, the acoustic scores provided by the trained SNN model are combined with the information stored in the language model and pronunciation lexicon. It is a common practice to use the weighted finite state transducers (WFST) (Mohri et al., 2002) as a unified representation of different ASR resources for creating the search graph containing possible hypotheses. The main motivation for using the WFST-based decoding is: (1) the straightforward composition of different ASR resources for constructing a mapping from HMM states to word sequences and (2) the existence of efficient search algorithms operating on WFST that speed up the decoding process. As a result of the search process, the most likely hypotheses are found and stored in the form of a lattice. The ASR output is chosen based on the weighted sum of the acoustic and language model scores belonging to hypotheses in the lattice. For further details of the WFST-based decoding approach used

in this work, we refer the reader to Povey et al. (2012). In the following sections, we describe the ASR experiments conducted to evaluate the recognition performance of the proposed SNN-based acoustic modeling in several recognition scenarios.

3.5. Training and Evaluation

3.5.1. Datasets

The performance of the proposed SNN-based acoustic models is investigated in three different ASR tasks: (1) phone recognition using the TIMIT corpus (Garofolo et al., 1993), (2) low-resourced ASR task using the FAME code-switching Frisian-Dutch corpus (Yilmaz et al., 2016a) and (3) standard large-vocabulary continuous ASR task using the Librispeech corpus (Panayotov et al., 2015). All speech data used in the experiments has a sampling frequency of 16 kHz.

The train, development and test sets of the standard TIMIT corpus contain 3,696, 400, and 192 utterances from 462, 50, and 24 speakers, respectively. Each utterance is phonetically transcribed using a phonetic alphabet consisting of 48 phones in total. The training data of the FAME corpus comprises of 8.5 and 3 h of broadcast speech from Frisian and Dutch speakers, respectively. The training utterances are spoken by 382 speakers in total. This bilingual dataset contains Frisian-only and Dutch-only utterances as well as mixed utterances with inter-sentential, intra-sentential and intra-word code-switching (Myers-Scotton, 1989). The development and test sets consist of 1 h of speech from Frisian speakers and 20 min of speech from Dutch speakers each. The total number of speakers is 61 in the development set and 54 in the test set.

The Librispeech corpus contains 1,000 h of reading speech in total collected from audiobooks. This publicly available corpus¹ has been considered as a popular benchmark for ASR algorithms with multiple training and testing settings. In the ASR experiments, we train acoustic models using the 100 (train_clean_100) and 360 (train_clean_360) h of speech and apply these models to the clean development (dev_clean) and test (test_clean) sets. Further details about this corpus can be found in Panayotov et al. (2015).

3.5.2. Implementation Details

All ASR experiments are performed using the PyTorch-Kaldi ASR toolkit (Ravanelli et al., 2019). This recently introduced toolkit inherits the flexibility of PyTorch toolkit (Paszke et al., 2019) for ANN-based acoustic model development and the efficiency of Kaldi ASR toolkit (Povey et al., 2011). We implement the SNN tandem learning rule in PyTorch and integrate it into the PyTorch-Kaldi toolkit for training the proposed SNN-based acoustic models (cf. **Figure 4**). The PyTorch implementation of the described SNN acoustic models is public available online². For the baseline ANN models, the standard multi-layer perceptron recipes are used. The Kaldi toolkit is used for obtaining the initial alignments, feature extraction, graph creation, and decoding.

¹www.openslr.org/resources/12

²https://github.com/deepspike/snn-for-asr

For all recognition scenarios, ANNs and SNNs are constructed with 4 hidden layers and 2,048 hidden units each using the ReLU activation function. Each fully-connected layer is followed by a batch normalization layer and a dropout layer with a drop probability of 10% to prevent overfitting. We train these models using various popular speech features including the 13-dimensional Mel-frequency cepstral coefficient (MFCC) feature, 23-dimensional Mel-filterbank (FBANK) feature, and higher resolution 40-dimensional MFCC and FBANK features. We further extract feature space maximum likelihood linear regression (FMLLR) (Gales, 1998) features to explore the impact of speaker-dependent features. All features include the deltas and delta-deltas; mean and variance normalization are applied before the splicing. The time context size is set to 11 frames by concatenating 5 frames preceding and following. All features are encoded within a short time window of 10-time steps for SNN simulations.

The neural network training is performed by mini-batch Stochastic Gradient Descent (SGD) with an initial learning rate of 0.08 and a minibatch size of 128. The learning rate is halved if the improvement is less than a preset threshold of 0.001. The final acoustic models of the TIMIT and FAME corpora are obtained after 24 training epochs, while the models of the Librispeech corpus are trained for 12 epochs.

For the TIMIT and Librispeech ASR tasks, we follow the same language model (LM) and pronunciation lexicon preparation pipeline as provided in the corresponding Kaldi recipes³. The smallest 3-gram LM (`tgsmall`) of the Librispeech corpus is used to create the graph for the decoding stage. The details of the LM and lexicon used in the FAME recognition task are given in Yilmaz et al. (2018).

3.5.3. Evaluation Metrics

3.5.3.1. ASR performance

The phone recognition on the TIMIT corpus is reported in terms of the phone error rate (PER). The word recognition accuracies on the FAME and Librispeech corpora are reported in terms of word error rate (WER). Both metrics are calculated as the ratio of all recognition errors (insertion, deletion, and substitution) and the total number of phones or words in the reference transcriptions.

3.5.3.2. Energy efficiency: counting synaptic operations

To compare the energy efficiency of ANN and its equivalent SNN implementation, we follow the convention from NC community and compute the total synaptic operations $SynOps$ that required to perform a certain task (Merolla et al., 2014; Rueckauer et al., 2017; Sengupta et al., 2019). For ANN, the total synaptic operations [Multiply-and-Accumulate (MAC)] per classification is defined as follows

$$SynOps = \sum_{l=1}^L f_{in}^l \cdot N_l \quad (17)$$

where f_{in}^l denotes the number of fan-in connections to each neuron in layer l , and N_l refers to the number of neurons in layer l . In addition, L denotes the total number of network layers. Hence, given a particular network configuration, the total synaptic operations required per classification is a constant number that jointly determined by f_{in}^l and N_l .

While for SNN, as per Equation (18), the total synaptic operations (Accumulate (AC)) required per classification are correlated with the spiking neurons' firing rate, the number of fan-out connections f_{out} to neurons in the subsequent layer as well as the simulation time window N_s .

$$SynOps = \sum_{t=1}^{N_s} \sum_{l=1}^{L-1} \sum_{j=1}^{N_l} f_{out,j}^l \cdot s_j^l(t) \quad (18)$$

where $s_j^l(t)$ indicates whether a spike is generated by neuron j of layer l at time instant t .

4. RESULTS

4.1. Phone Recognition on TIMIT Corpus

We report the PER on the development and test sets of TIMIT corpus in **Table 1**, with numbers in bold being the best performance given by the speaker-independent features. ASR performances of other state-of-the-art systems using various ANN and SNN architectures are given in the upper panel for reference purposes. As the results shown in **Table 1**, the proposed SNN-based acoustic models are applicable to different speech features and provide comparable or slightly worse ASR performance than the ANNs with the same network structure. In particular, the ANN system trained with the standard 13-dimensional FBANK feature achieves the best PER of 16.9% (18.5%) on the development (test) set. The equivalent SNN system using the same feature achieves slightly worse PER of 17.3% (18.7%) on the development (test) set. Although the state-of-the-art ASR systems (Ravanelli et al., 2018) give approximately 1% lower PER than the proposed SNN-based phone recognition system, it is largely credit to the longer time context explored by the recurrent Li-GRU model.

It worth mentioning that phone recognition is still a challenging task for spiking neural networks. To the best of our knowledge, only one recent work with recurrent spiking neural networks (Bellec et al., 2019) demonstrates some promising test results on this corpus with a PER of 26.4%. In contrast, our system has achieved significantly lower PER compared to this preliminary study of SNN-based acoustic modeling. However, these results are not directly comparable since the proposed system incorporates both an acoustic and a language model during decoding unlike the system described in Bellec et al. (2019).

The experimental results on the TIMIT phone recognition task can be considered as an initial indicator of the compelling prospects of the SNN-based acoustic modeling. Given that the phone recognition task on TIMIT corpus is simplistic compared to the modern LVCSR tasks, we further compare the

³<https://github.com/kaldi-asr/kaldi/tree/master/egs/timit>;
<https://github.com/kaldi-asr/kaldi/tree/master/egs/librispeech>

TABLE 1 | PER (%) on the TIMIT development and test sets.

Features	AM		Test	
	Li-GRU (Ravanelli et al., 2018)		RSNN (Bellec et al., 2019)	
MFCC (13-dim.)	16.7		26.4	
FBANK (13-dim.)	15.8		–	
FMLLR	14.9*		–	

Features	AM		Test	
	Dev		Test	
	ANN	SNN	ANN	SNN
MFCC (13-dim.)	17.1	17.8	18.5	19.1
FBANK (13-dim.)	16.9	17.3	18.5	18.7
MFCC (40-dim.)	17.3	18.2	18.7	19.8
FBANK (40-dim.)	16.9	17.8	17.9	19.1
FMLLR	15.8	16.5	17.2	17.4

The upper panel reports the results of various ANN and SNN architectures from the literatures, and the lower panel presents the results achieved by the ANN and SNN models in this work (AM, acoustic model, *the best result to date). The best results given by the speaker-independent features at each column are marked in bold.

ANN and SNN performance on newer corpora designed for LVCSR experiments.

4.2. Low-Resourced ASR on FAME Corpus

In this section, we apply the SNN-based ASR systems to the low-resourced ASR scenario. As summarized in **Table 2**, the word recognition results on the FAME corpus are reported separately for monolingual Frisian (fy), monolingual Dutch (nl) and code-switched (cs) utterances. The overall performance (all) is also included in the rightmost column. Given that 8.5 h Frisian and 3 h of Dutch speech is used during the training phase, we can compare the ASR performance on different subsets, i.e., fy, nl and cs, to identify the variations in the ASR performance for different levels of low-resourcedness. We omit the results on the development set as they follow a similar pattern to the results on the test set.

In this scenario, the SNN acoustic models consistently provide lower WERs than the ANN models for all speech features. Systems with the FBANK features provide lower WERs than those using MFCC features, which is in line with our observations on the TIMIT corpus. The best performance on the test set is obtained using SNN models trained on 40-dimensional FBANK features with an overall WER of 36.9%. In contrast, the ANN model provides a WER of 39.0% for the same setting, which is relatively 5.4% worse than the SNN model. Moreover, the SNN-based acoustic models achieve a relative improvement of 4.7%, 5.2% and 8.2% on the fy, nl and cs subsets of the test set, respectively. These steady improvements in the recognition accuracies highlight the effectiveness of the SNN-based acoustic modeling in scenarios with limited training data compared to the conventional ANN models. The improved ASR performance with SNNs, in the low-resourced setting, may credit to the noisy weight updates derived by the tandem learning framework. It has been recognized that introducing noises into the training

TABLE 2 | WERs (%) achieved on the monolingual and mixed segments of the FAME test set.

Speech features	AM	fy	nl	cs	All
		# of Frisian words # of Dutch words	10,753 0	0 3,475	1,798 306
FBANK (40-dim.)	Kaldi-ANN (Yilmaz et al., 2016b)	32.4	39.7	49.9	36.2
MFCC (40-dim.)	TDNN-LSTM (Yilmaz et al., 2018)	31.5	39.5	47.9	35.2
MFCC (13-dim.)	ANN	34.6	50.0	49.9	39.9
MFCC (13-dim.)	SNN	33.8	45.3	47.9	38.2
FBANK (13-dim.)	ANN	34.3	47.5	48.1	39.0
FBANK (13-dim.)	SNN	33.1	44.3	46.5	37.3
MFCC (40-dim.)	ANN	35.2	48.4	51.7	40.2
MFCC (40-dim.)	SNN	33.7	44.2	46.9	37.7
FBANK (40-dim.)	ANN	34.4	46.3	49.8	39.0
FBANK (40-dim.)	SNN	32.8	43.9	45.7	36.9
FMLLR	ANN	31.2	42.1	47.2	35.7
FMLLR	SNN	31.5	39.5	46.6	35.2

The upper panel summarizes the number of words from each language subset. The middle panel provides the results of state-of-the-art ANN architectures (Yilmaz et al., 2016b, 2018) for reference purposes and the lower panel presents the results achieved by the ANN and SNN models in this work (AM, acoustic model). The best results given by the speaker-independent features at each column are marked in bold.

stage improves the generalization capability of ANN-based ASR systems (Yin et al., 2015). As a result, the noisy training of the tandem learning is expected to improve the recognition performance in low-resourced scenarios. Further investigation on the impact of this noisy training procedure remains as future work.

4.3. LVCSR Experiments on Librispeech Corpus

In the final set of ASR experiments, we train acoustic models using the official 100 and 360-h training subsets of the Librispeech corpus to compare the recognition performance of ANN and SNN models in a standard LVCSR scenario. As the results given in the middle panel of **Table 3**, for 100 h of training data, the ANN systems perform marginally better than the corresponding SNN systems across all different speech features. The absolute WER differences range from 0.1% to 0.6%. These marginal performance degradations of the SNN models is likely due to the reduced representation power of using discrete spike counts. Nevertheless, these results are promising even when comparing to the state-of-the-art ASR systems using more complex ANN architectures as provided in the upper panel of **Table 3**.

It worth noting that both ANN and SNN systems can take benefit of an increased amount of training data. When increasing the training data from 100 to 360 h, the WERs of the best SNN models reduced from 10.0% (10.3%) to 9.2% (9.4%) for the development (test) sets, respectively. To the best of our knowledge, it is the first time that SNN-based acoustic models

TABLE 3 | WER (%) achieved on the Librispeech development and test sets.

		Train - 100 h			
		Dev		Test	
AM					
Kaldi†	p-norm ANN	9.2		9.7	
PyTorch-Kaldi†	Li-GRU	–		8.6	
Features	AM	ANN	SNN	ANN	SNN
MFCC		10.3	10.5	10.6	10.9
FBANK		9.6	10.0	10.2	10.6
MFCC (40-dim.)		9.5	10.1	10.0	10.6
FBANK (40-dim.)		9.6	10.2	10.1	10.3
FMLLR		9.2	9.3	9.7	9.9
		Train - 360 h			
		Dev		Test	
Features	AM	ANN	SNN	ANN	SNN
MFCC		9.2	9.9	9.6	10.3
FBANK		8.6	9.7	9.1	10.0
MFCC (40-dim.)		8.6	9.2	8.9	9.4
FBANK (40-dim.)		8.5	9.4	8.9	9.7
FMLLR		8.4	9.2	8.8	9.7

The upper panel gives the results, with 100-h of training data, reported at the Github repo of Kaldi and PyTorch-Kaldi. The middle and lower panel present the results achieved by ANN and SNN models in this work using 100-h and 360-h of training data, respectively. The best results given by the speaker-independent features in the middle and lower panel are marked in bold. (AM, acoustic model, †: reported at Github repo).

have achieved comparable results over the ANN models for LVCSR tasks. These results suggest that SNNs are potentially good candidates for acoustic modeling.

4.4. Energy Efficiency of SNN-Based ASR Systems

In addition to the promising modeling capability, the SNN-based ASR systems can achieve unprecedented performance gain when implemented on the low-power neuromorphic chips. In this section, we shed light on this prospect by comparing the energy efficiency of ANN- and SNN-based acoustic models. Given that data movements are the most energy-consuming operations for data-driven AI applications, we calculate the average synaptic operations on 5 randomly chosen utterances from the TIMIT corpus and report the ratio of average synaptic operations required per feature classification [SynOps(SNN)/SynOps(ANN)]. To investigate the effect of different feature representations, we repeat our analysis on the 40-dimensional MFCC, FBANK, and FMLLR features as summarized in Table 4 and Figure 5.

Taking advantage of the short encoding time window ($N_s = 10$), the sparse neuronal activities are observed for all network layers as shown in Figure 5. Among the three features explored in this experiment, it is interesting to note the FMLLR feature achieves the lowest average spike rate. It is likely due to the

TABLE 4 | Comparison of the computational costs between SNN and ANN.

Utterance Index	1	2	3	4	5	Avg. SynOps Ratio
Num. of frames	474	287	274	268	223	
MFCC (40-dim.)	1.71	1.73	1.76	1.71	1.68	1.72
FBANK (40-dim.)	1.08	1.08	1.14	1.09	1.10	1.10
FMLLR	0.67	0.68	0.71	0.66	0.67	0.68

The ratio of their required total synaptic operations [SynOps(SNN)/SynOps(ANN)] is reported. It worth mentioning that ANNs use more costly MAC operations than the AC operations used in the SNNs.

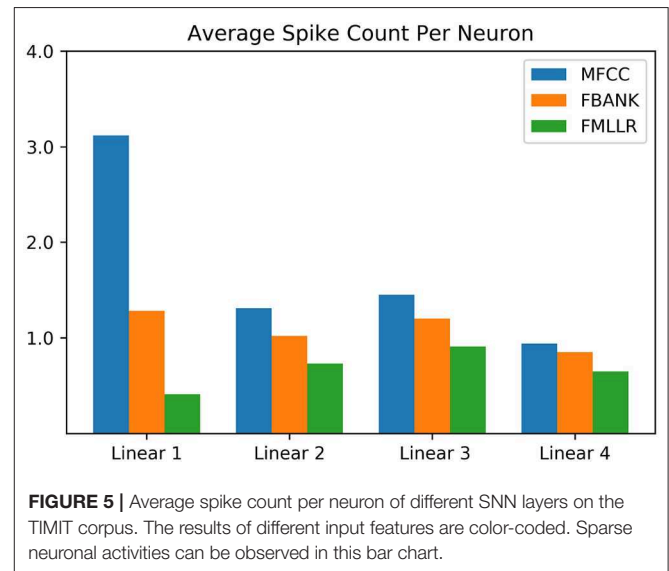


FIGURE 5 | Average spike count per neuron of different SNN layers on the TIMIT corpus. The results of different input features are color-coded. Sparse neuronal activities can be observed in this bar chart.

more discriminative nature of the speaker-dependent feature, while it worth to note that the FMLLR feature is not always available in all ASR scenarios. As provided in Table 4, the SNN implementations taking MFCC, FBANK and FMLLR input features require 1.72, 1.10, and 0.68 times synaptic operations to their ANN counterparts, respectively. Although the average number of synaptic operations required for SNNs that using MFCC and FBANK features are slightly higher than the ANNs, the AC operations performed on SNNs are much cheaper than the MAC operations required for ANNs. One recent study on the Global Foundry 28 nm process has revealed that MAC operations are 14 times more costly than AC operations and requires 21 times more chip area (Rueckauer et al., 2017). This study provides some good indicators for the potential energy and chips area savings that can be received from deploying SNNs onto the emerging neuromorphic chips for inference (Merolla et al., 2014; Davies et al., 2018). While the actual energy savings for SNN-based acoustic models are dependent on the chip architectures and materials used, which is beyond the scope of this work.

5. DISCUSSION

The remarkable progress in the automatic speech recognition systems has revolutionized the human-computer interface. The

rapid growing demands of ASR services have raised concerns on computational efficiency, real-time performance, and data security, etc. It, therefore, motivates novel solutions to address all those concerns. As inspired by the event-driven computation that observed in the biological neural systems, we explore using brain-inspired spiking neural networks for large vocabulary ASR tasks. For this purpose, we proposed a novel SNN-based ASR framework, wherein the SNN is used for acoustic modeling and map the frame-level features into a set of acoustic units. These frame-level outputs will further integrate the word-level information from the corresponding language model to find the most likely word sequence corresponding to the input speech signal.

5.1. Superior Speech Recognition Performance With SNNs

The phone and word recognition experiments on the well-known TIMIT and Librispeech benchmarks have demonstrated the promising modeling capacity of SNN acoustic models and their applicability to different input features. These preliminary results have shown that the recognition performance of SNNs is either comparable or slightly worse than the ANNs with the same network architecture on the TIMIT and Librispeech benchmarks. A possible reason for this performance degradation is the reduced representation power of the discrete neural representation (i.e., spike counts) as compared to the continuous floating-point representation of the ANNs (Wu et al., 2019c). This performance gap could potentially be closed by extending the encoding window N_s of SNNs. Moreover, the recognition performance of ANN and SNN models in a low-resourced scenario is also investigated. In this scenario, the SNN acoustic models outperform the conventional ANNs that could be attributed to the noisy training of the tandem learning framework, wherein error gradients of the SNN layers are approximated from the coupled ANN layers.

The neural encoding scheme adopted in this work allows input features to be encoded inside a short encoding time window for rapid processing by SNNs. It is attractive for the time-synchronous ASR tasks that require real-time performance. The preliminary study of the energy efficiency on the TIMIT corpus reveals attractive energy and chip area savings, as compared to the equivalent ANNs, can be achieved when deploying the offline trained SNNs onto neuromorphic chips. The recent study of a keyword spotting task on the Loihi neuromorphic research chip (Blouw et al., 2019) has also demonstrated the compelling energy savings, real-time performance and good scalability of emerging NC architectures over conventional low-power AI chips designed for ANNs.

5.2. Development of SNN-Based ASR Systems

The active development of open-source software toolkits plays a significant role in the rapid progress of ASR research, instances include the Kaldi (Povey et al., 2011) and ESPnet (Watanabe et al., 2018). In this work, we demonstrate that state-of-the-art SNN acoustic models can be easily developed

in PyTorch and integrated into the PyTorch-Kaldi Speech Recognition Toolkit (Ravanelli et al., 2019). This software toolkit integrates the efficiency of Kaldi and the flexibility of PyTorch, therefore, it can support the rapid development of SNN-based ASR systems.

5.3. Future Directions

The recurrent neural networks have shown great modeling capability for temporal signals by exploring long temporal context information in the input signals (Graves and Jaitly, 2014). As future work, we will explore the recurrent networks of spiking neurons for large-vocabulary ASR tasks to further improve the recognition performance.

The substantial research efforts are devoted to reducing the computational cost and memory footprint of ANNs during inference, instances include network compression (Han et al., 2015), network quantization (Courbariaux et al., 2016; Zhou et al., 2016) and knowledge distillation (Hinton et al., 2015). While the computational paradigm underlying the efficient biological neural networks is fundamentally different from ANNs and hence fosters enormous potentials for neuromorphic computing architectures. Furthermore, grounded on the same connectionism principle, the information of both ANN and SNN are encoded in the network connectivity and connection strength. Therefore, SNN can also take benefits from these early research works on the network compression and quantization of ANNs to further reduce its memory footprint and computation cost (Deng et al., 2019).

The event-driven silicon cochlea audio sensors (Liu et al., 2014) are designed to mimic the functional mechanism of human cochlea and transform input audio signals into spiking events. Given temporally sparse information is transmitted in the surrounding environment, these sensors have shown greater coding efficiency than conventional microphone sensors (Liu et al., 2019). There are some interesting preliminary ASR studies explore the input spiking events captured by these silicon cochlea sensors (Acharya et al., 2018; Anumula et al., 2018). Additionally, Dominguez-Morales et al. (2018) have proposed a fully SNN-based framework for voice commands recognition, wherein the event-driven silicon cochlea audio sensor is directly interfaced with the SpiNNaker neuromorphic processor through the Address-Event Representation protocol (AER). Notably, a buffering layer is introduced to ensure real-time performance. However, the scale of the ASR tasks explored in these studies is relatively small comparing to modern ASR benchmarks due to the limited availability of event-based ASR corpora. Pan et al. (2020) recently proposed an efficient and perceptually motivated auditory neural encoding scheme to encode the large-scale ASR corpora collected by microphone sensors into spiking events. With this encoding scheme, approximately 50% spiking events can be reduced with negligible interference to the perceptual quality of inputs audio signals. Taking benefits from these earlier research on the neuromorphic auditory front-end, we are expecting to further improve the energy efficiency of SNN-based ASR systems.

The promising initial results demonstrated by the SNN-based large vocabulary ASR systems in this work is the

first step toward a myriad opportunities for the integration of state-of-the-art ASR engines into mobile and embedded devices with power restrictions. In the long run, the SNN-based ASR systems are expected to take benefits from ever-growing research on novel neuromorphic auditory front-end, SNN architectures, neuromorphic computing architectures and ultra-low-power non-volatile memory devices to further improve the computing performance.

DATA AVAILABILITY STATEMENT

Publicly available datasets were analyzed in this study. This data can be found here: <https://catalog.ldc.upenn.edu/LDC93S1>; <https://www.openslr.org/12>; <https://repository.uibn.ru.nl/handle/2066/162244>.

REFERENCES

- Abdel-Hamid, O., Mohamed, A., Jiang, H., Deng, L., Penn, G., and Yu, D. (2014). Convolutional neural networks for speech recognition. *IEEE/ACM Trans. Audio Speech Lang. Process.* 22, 1533–1545. doi: 10.1109/TASLP.2014.2339736
- Acharya, J., Patil, A., Li, X., Chen, Y., Liu, S.-C., and Basu, A. (2018). A comparison of low-complexity real-time feature extraction for neuromorphic speech recognition. *Front. Neurosci.* 12:160. doi: 10.3389/fnins.2018.00160
- Anumula, J., Neil, D., Delbruck, T., and Liu, S. C. (2018). Feature representations for neuromorphic audio spike streams. *Front. Neurosci.* 12:23. doi: 10.3389/fnins.2018.00023
- Bellec, G., Salaj, D., Subramoney, A., Legenstein, R., and Maass, W. (2018). “Long short-term memory and learning-to-learn in networks of spiking neurons,” in *Advances in Neural Information Processing Systems* (Montréal, QC), 787–797.
- Bellec, G., Scherr, F., Subramoney, A., Hajek, E., Salaj, D., Legenstein, R., et al. (2019). A solution to the learning dilemma for recurrent networks of spiking neurons. *bioRxiv [Preprint]*. doi: 10.1101/738385. Available online at: <https://graz.pure.elsevier.com/en/publications/a-solution-to-the-learning-dilemma-for-recurrent-networks-of-spik>
- Blouw, P., Choo, X., Hunsberger, E., and Eliasmith, C. (2019). “Benchmarking keyword spotting efficiency on neuromorphic hardware,” in *Proceedings of the 7th Annual Neuro-inspired Computational Elements Workshop* (Albany, NY: ACM), 1.
- Cao, Y., Chen, Y., and Khosla, D. (2015). Spiking deep convolutional neural networks for energy-efficient object recognition. *Int. J. Comput. Vis.* 113, 54–66. doi: 10.1007/s11263-014-0788-3
- Chan, W., Jaitly, N., Le, Q. V., and Vinyals, O. (2016). “Listen, attend and spell: A neural network for large vocabulary conversational speech recognition,” in *Proceedings of the ICASSP* (Shanghai: IEEE), 4960–4964. doi: 10.1109/ICASSP.2016.7472621
- Courbariaux, M., Hubara, I., Soudry, D., El-Yaniv, R., and Bengio, Y. (2016). Binarized neural networks: training deep neural networks with weights and activations constrained to +1 or -1. *arXiv:1602.02830*.
- Dahl, G. E., Yu, D., Deng, L., and Acero, A. (2012). Context-dependent pre-trained deep neural networks for large-vocabulary speech recognition. *IEEE Trans. Audio Speech Lang. Process.* 20, 30–42. doi: 10.1109/TASL.2011.2134090
- Davies, M., Srinivasa, N., Lin, T. H., China, G., Cao, Y., Choday, S. H., et al. (2018). Loihi: a neuromorphic manycore processor with on-chip learning. *IEEE Micro* 38, 82–99. doi: 10.1109/MM.2018.112130359
- Davis, S., and Mermelstein, P. (1980). Comparison of parametric representations for monosyllabic word recognition in continuously spoken sentences. *IEEE Trans. Acoust.* 28, 357–366. doi: 10.1109/TASSP.1980.1163420
- Deng, L., Wu, Y., Hu, Y., Liang, L., Li, G., Hu, X., et al. (2019). Comprehensive snn compression using admm optimization and activity regularization. *arXiv:1911.00822*.

AUTHOR CONTRIBUTIONS

JW and EY designed and conducted all the experiments. All authors contributed to the results interpretation and writing.

FUNDING

This research was supported by Programmatic Grant No. A1687b0033 from the Singapore Government’s Research, Innovation and Enterprise 2020 plan (Advanced Manufacturing and Engineering domain). JW was also partially supported by the Zhejiang Lab’s International Talent Fund for Young Professionals and the Zhejiang Lab (No.2019KC0AB02). HL was also partially supported by U Bremen Excellence Chairs program (2019–2022), Germany.

- Diehl, P. U., Neil, D., Binas, J., Cook, M., Liu, S. C., and Pfeiffer, M. (2015). “Fast-classifying, high-accuracy spiking deep networks through weight and threshold balancing,” in *2015 International Joint Conference on Neural Networks (IJCNN)* (Killarney), 1–8.
- Dominguez-Morales, J. P., Liu, Q., James, R., Gutierrez-Galan, D., Jimenez-Fernandez, A., Davidson, S., et al. (2018). “Deep spiking neural network model for time-variant signals classification: a real-time speech recognition approach,” in *2018 International Joint Conference on Neural Networks (IJCNN)* (Rio de Janeiro), 1–8.
- Furber, S. B., Lester, D. R., Plana, L. A., Garside, J. D., Painkras, E., Temple, S., et al. (2012). Overview of the SpiNNaker system architecture. *IEEE Trans. Comput.* 62, 2454–2467. doi: 10.1109/TC.2012.142
- Gales, M. J. F. (1998). Maximum likelihood linear transformations for hmm-based speech recognition. *Comput. Speech Lang.* 12, 75–98. doi: 10.1006/csla.1998.0043
- Garofolo, J. S., Lamel, L. F., Fisher, W. M., Fiscus, J. G., Pallett, D. S., Dahlgren, N. L., et al. (1993). *TIMIT Acoustic-Phonetic Continuous Speech Corpus (LDC93S1)*. Philadelphia, PA: Linguistic Data Consortium.
- Gerstner, W., and Kistler, W. M. (2002). *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge: Cambridge University Press.
- Graves, A., and Jaitly, N. (2014). “Towards end-to-end speech recognition with recurrent neural networks,” in *Proceedings of the 31st International Conference on Machine Learning (ICML)* (Beijing: ACM), 1764–1772.
- Graves, A., Mohamed, A., and Hinton, G. (2013). “Speech recognition with deep recurrent neural networks,” in *Proceedings of the ICASSP* (Vancouver, BC), 6645–6649.
- Greff, K., Srivastava, R. K., Koutnik, J., Steunebrink, B. R., and Schmidhuber, J. (2016). Lstm: a search space odyssey. *IEEE Trans. Neural Netw. Learn. Syst.* 28, 2222–2232. doi: 10.1109/TNNLS.2016.2582924
- Han, S., Mao, H., and Dally, W. J. (2015). Deep compression: compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv:1510.00149*.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). “Deep residual learning for image recognition,” in *Proceedings of the IEEE CVPR* (Las Vegas, NV), 770–778. doi: 10.1109/CVPR.2016.90
- He, T., Fan, Y., Qian, Y., Tan, T., and Yu, K. (2014). “Reshaping deep neural network for fast decoding by node-pruning,” in *Proceedings of the ICASSP* (Vancouver, BC), 245–249. doi: 10.1109/ICASSP.2014.6853595
- Hinton, G., Deng, L., Yu, D., Dahl, G. E., Mohamed, A.-R., Jaitly, N., et al. (2012). Deep neural networks for acoustic modeling in speech recognition: the shared views of four research groups. *IEEE Signal Process. Mag.* 29, 82–97. doi: 10.1109/MSP.2012.2205597
- Hinton, G., Vinyals, O., and Dean, J. (2015). Distilling the knowledge in a neural network. *arXiv:1503.02531*.
- Holmberg, M., Gelbart, D., Ramacher, U., and Hemmert, W. (2005). “Automatic speech recognition with neural spike trains,” in *INTERSPEECH* (Lisbon), 1253–1256.

- Hu, Y., Tang, H., Wang, Y., and Pan, G. (2018). Spiking deep residual network. *arXiv:1805.01352*.
- Hwang, M.-Y., and Huang, X. (1993). Shared-distribution hidden markov models for speech recognition. *IEEE Trans. Speech Audio Process.* 1, 414–420. doi: 10.1109/89.242487
- Kröger, B. J., Kannampuzha, J., and Neuschaefer-Rube, C. (2009). Towards a neurocomputational model of speech production and perception. *Speech Commun.* 51, 793–809. doi: 10.1016/j.specom.2008.08.002
- Lang, K. J., Waibel, A. H., and Hinton, G. E. (1990). A time-delay neural network architecture for isolated word recognition. *Neural Netw.* 3, 23–43. doi: 10.1016/0893-6080(90)90044-L
- Laughlin, S. B., and Sejnowski, T. J. (2003). Communication in neuronal networks. *Science* 301, 1870–1874. doi: 10.1126/science.1089662
- Lee, K. (1990). Context-independent phonetic hidden markov models for speaker-independent continuous speech recognition. *IEEE Trans. Acoust.* 38, 599–609. doi: 10.1109/29.52701
- Lei, X., Senior, A. W., Gruenstein, A., and Sorensen, J. S. (2013). “Accurate and compact large vocabulary speech recognition on mobile devices,” in *Proceedings of the INTERSPEECH* (Vancouver, BC), 662–665.
- Liaw, J. S., and Berger, T. W. (1998). “Robust speech recognition with dynamic synapses,” in *IEEE International Joint Conference on Neural Networks Proceedings (IJCNN)*, Vol. 3 (Anchorage, AK), 2175–2179.
- Lippmann, R. P. (1989). Review of neural networks for speech recognition. *Neural Comput.* 1, 1–38. doi: 10.1162/neco.1989.1.1.1
- Liu, S.-C., Rueckauer, B., Ceolini, E., Huber, A., and Delbruck, T. (2019). Event-driven sensing for efficient perception: vision and audition algorithms. *IEEE Signal Process. Mag.* 36, 29–37. doi: 10.1109/MSP.2019.2928127
- Liu, S. C., van Schaik, A., Minch, B. A., and Delbruck, T. (2014). Asynchronous binaural spatial audition sensor with 2x64x4 channel output. *IEEE Trans. Biomed. Circuits Syst.* 8, 453–464. doi: 10.1109/TBCAS.2013.2281834
- Loiselle, S., Rouat, J., Pressnitzer, D., and Thorpe, S. (2005). “Exploration of rank order coding with spiking neural networks for speech recognition,” in *IEEE International Joint Conference on Neural Networks (IJCNN)*, Vol. 4 (Montréal, QC), 2076–2080.
- McGraw, I., Prabhavalkar, R., Alvarez, R., Arenas, M. G., Rao, K., Rybach, D., et al. (2016). “Personalized speech recognition on mobile devices,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (Shanghai), 5955–5959.
- Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., et al. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science* 345, 668–673. doi: 10.1126/science.1254642
- Mohri, M., Pereira, F., and Riley, M. (2002). Weighted finite-state transducers in speech recognition. *Comput. Speech Lang.* 16, 69–88. doi: 10.1006/csla.2001.0184
- Monroe, D. (2014). Neuromorphic computing gets ready for the (really) big time. *Commun. ACM* 57, 13–15. doi: 10.1145/2601069
- Myers-Scotton, C. (1989). Codeswitching with English: types of switching, types of communities. *World Englishes* 8, 333–346. doi: 10.1111/j.1467-971X.1989.tb00673.x
- Näger, C., Storck, J., and Deco, G. (2002). Speech recognition with spiking neurons and dynamic synapses: a model motivated by the human auditory pathway. *Neurocomputing* 44–46, 937–942. doi: 10.1016/S0925-2312(02)00494-0
- Neftci, E. O., Mostafa, H., and Zenke, F. (2019). Surrogate gradient learning in spiking neural networks. *arXiv: 1901.09948*.
- Pan, Z., Chua, Y., Wu, J., Zhang, M., Li, H., and Ambikairajah, E. (2020). An efficient and perceptually motivated auditory neural encoding and decoding algorithm for spiking neural networks. *Front. Neurosci.* 13:1420. doi: 10.3389/fnins.2019.01420
- Pan, Z., Li, H., Wu, J., and Chua, Y. (2018). “An event-based cochlear filter temporal encoding scheme for speech signals,” in *2018 International Joint Conference on Neural Networks (IJCNN)* (Rio de Janeiro: IEEE), 1–8.
- Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. (2015). “Librispeech: an ASR corpus based on public domain audio books,” in *Proceedings of the ICASSP* (South Brisbane, QLD), 5206–5210.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). “PyTorch: an imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems* 32, eds H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Vancouver, BC: Curran Associates, Inc.), 8026–8037.
- Pfeiffer, M., and Pfeil, T. (2018). Deep learning with spiking neurons: opportunities & challenges. *Front. Neurosci.* 12:774. doi: 10.3389/fnins.2018.00774
- Povey, D., Cheng, G., Wang, Y., Li, K., Xu, H., Yarmohammadi, M., et al. (2018). “Semi-orthogonal low-rank matrix factorization for deep neural networks,” in *Proceedings of the Interspeech* (Hyderabad), 3743–3747. doi: 10.21437/Interspeech.2018-1417
- Povey, D., Ghoshal, A., Boulianne, G., Burget, L., Glembek, O., Goel, N., et al. (2011). “The kaldi speech recognition toolkit,” in *IEEE ASRU* (Hawaii).
- Povey, D., Hannemann, M., Boulianne, G., Burget, L., Ghoshal, A., Janda, M., et al. (2012). “Generating exact lattices in the WFST framework,” in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (Kyoto), 4213–4216.
- Ravanelli, M., Brakel, P., Omologo, M., and Bengio, Y. (2018). Light gated recurrent units for speech recognition. *IEEE Trans. Emerg. Top. Comput. Intell.* 2, 92–102. doi: 10.1109/TETCI.2017.2762739
- Ravanelli, M., Parcollet, T., and Bengio, Y. (2019). “The pytorch-kaldi speech recognition toolkit,” in *Proceedings of the ICASSP* (Brighton, UK), 6465–6469.
- Rueckauer, B., Lungu, I. A., Hu, Y., Pfeiffer, M., and Liu, S. C. (2017). Conversion of continuous-valued deep networks to efficient event-driven networks for image classification. *Front. Neurosci.* 11:682. doi: 10.3389/fnins.2017.00682
- Sainath, T. N., Kingsbury, B., Sinthwani, V., Arisoy, E., and Ramabhadran, B. (2013). “Low-rank matrix factorization for deep neural network training with high-dimensional output targets,” in *Proceedings of the ICASSP* (Vancouver, BC), 6655–6659.
- Sainath, T. N., and Parada, C. (2015). “Convolutional neural networks for small-footprint keyword spotting,” in *Proceedings of the INTERSPEECH* (Dresden), 1478–1482.
- Sengupta, A., Ye, Y., Wang, R., Liu, C., and Roy, K. (2019). Going deeper in spiking neural networks: VGG and residual architectures. *Front. Neurosci.* 13:95. doi: 10.3389/fnins.2019.00095
- Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., et al. (2017). Mastering the game of go without human knowledge. *Nature* 550:354. doi: 10.1038/nature24270
- Tang, J., Yuan, F., Shen, X., Wang, Z., Rao, M., He, Y., et al. (2019). Bridging biological and artificial neural networks with emerging neuromorphic devices: fundamentals, progress, and challenges. *Adv. Mater.* 31:1902761. doi: 10.1002/adma.201902761
- Tavanaei, A., Ghodrati, M., Kheradpisheh, S. R., Masquelier, T., and Maida, A. (2019). Deep learning in spiking neural networks. *Neural Netw.* 111, 47–63. doi: 10.1016/j.neunet.2018.12.002
- Tavanaei, A., and Maida, A. (2017a). “Bio-inspired multi-layer spiking neural network extracts discriminative features from speech signals,” in *International Conference on Neural Information Processing* (Guangzhou: Springer), 899–908.
- Tavanaei, A., and Maida, A. (2017b). A spiking network that learns to extract spike signatures from speech signals. *Neurocomputing* 240, 191–199. doi: 10.1016/j.neucom.2017.01.088
- Watanabe, S., Hori, T., Karita, S., Hayashi, T., Nishitoba, J., Unno, Y., et al. (2018). Espnet: End-to-end speech processing toolkit. *arXiv: 1804.00015*. doi: 10.21437/Interspeech.2018-1456
- Watanabe, S., Hori, T., Kim, S., Hershey, J. R., and Hayashi, T. (2017). Hybrid CTC/attention architecture for end-to-end speech recognition. *IEEE J. Sel. Top. Signal Process.* 11, 1240–1253. doi: 10.1109/JSTSP.2017.2763455
- Wu, J., Chua, Y., and Li, H. (2018a). “A biologically plausible speech recognition framework based on spiking neural networks,” in *2018 International Joint Conference on Neural Networks (IJCNN)* (Rio de Janeiro), 1–8.
- Wu, J., Chua, Y., Zhang, M., Li, G., Li, H., and Tan, K. C. (2019c). A tandem learning rule for efficient and rapid inference on deep spiking neural networks. *arXiv:1907.01167*.
- Wu, J., Chua, Y., Zhang, M., Li, H., and Tan, K. C. (2018b). A spiking neural network framework for robust sound classification. *Front. Neurosci.* 12:836. doi: 10.3389/fnins.2018.00836
- Wu, J., Chua, Y., Zhang, M., Yang, Q., Li, G., and Li, H. (2019a). “Deep spiking neural network with spike count based learning rule,” in *2019 International Joint Conference on Neural Networks (IJCNN)* (Budapest: IEEE), 1–6.

- Wu, J., Pan, Z., Zhang, M., Das, R. K., Chua, Y., and Li, H. (2019b). "Robust sound recognition: a neuromorphic approach," in *Proceedings of the Interspeech 2019* (Graz), 3667–3668.
- Wu, M., Panchapagesan, S., Sun, M., Gu, J., Thomas, R., Prasad Vitaladevuni, S. N., et al. (2018). "Monophone-based background modeling for two-stage on-device wake word detection," in *Proceedings of the ICASSP* (Calgary, AB), 5494–5498.
- Wu, Y., Deng, L., Li, G., Zhu, J., and Shi, L. (2018). Direct training for spiking neural networks: faster, larger, better. *arXiv:1809.05793*.
- Xiong, W., Droppo, J., Huang, X., Seide, F., Seltzer, M. L., Stolcke, A., et al. (2017). Toward human parity in conversational speech recognition. *IEEE/ACM Trans. Audio Speech Lang. Process.* 25, 2410–2423. doi: 10.1109/TASLP.2017.2756440
- Xue, J., Li, J., and Gong, Y. (2013). "Restructuring of deep neural network acoustic models with singular value decomposition," in *Proceedings of the Interspeech* (Lyon), 2365–2369.
- Yilmaz, E., Andringa, M., Kingma, S., Van der Kuip, F., Van de Velde, H., Kampstra, F., et al. (2016a). "A longitudinal bilingual Frisian-Dutch radio broadcast database designed for code-switching research," in *Proceedings of the LREC* (Portorož), 4666–4669.
- Yilmaz, E., van den Heuvel, H., and van Leeuwen, D. (2016b). "Code-switching detection using multilingual DNNs," in *2016 IEEE Spoken Language Technology Workshop (SLT)* (San Diego, CA), 610–616.
- Yilmaz, E., Van den Heuvel, H., and Van Leeuwen, D. A. (2018). "Acoustic and textual data augmentation for improved ASR of code-switching speech," in *Proceedings of the INTERSPEECH* (Hyderabad), 1933–1937.
- Yin, S., Liu, C., Zhang, Z., Lin, Y., Wang, D., Tejedor, J., et al. (2015). Noisy training for deep neural networks in speech recognition. *EURASIP J. Audio Speech Music Process.* 2015:2. doi: 10.1186/s13636-014-0047-0
- Yu, D., and Deng, L. (2015). *Automatic Speech Recognition: A Deep Learning Approach. Signals and Communication Technology*. Lisbon: Springer.
- Zehetner, A., Haggmuller, M., and Pernkopf, F. (2014). "Wake-up-word spotting for mobile systems," in *Proceedings of the EUSIPCO*, 1472–1476.
- Zhang, M., Wu, J., Chua, Y., Luo, X., Pan, Z., Liu, D., et al. (2019). "Mpd-al: an efficient membrane potential driven aggregate-label learning algorithm for spiking neurons," in *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33 (Hawaii), 1327–1334. doi: 10.1609/aaai.v33i01.33011327
- Zhang, Y., Li, P., Jin, Y., and Choe, Y. (2015). A digital liquid state machine with biologically inspired learning and its application to speech recognition. *IEEE Trans Neural Netw. Learn. Syst.* 26, 2635–2649. doi: 10.1109/TNNLS.2015.2388544
- Zhou, S., Wu, Y., Ni, Z., Zhou, X., Wen, H., and Zou, Y. (2016). Dorefanet: training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv:1606.06160*.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Wu, Yilmaz, Zhang, Li and Tan. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.