



OPEN ACCESS

EDITED BY

Daniela Gandolfi,
University of Modena and Reggio Emilia, Italy

REVIEWED BY

Thomas Wennekers,
University of Plymouth, United Kingdom
Timothée Masquelier,
Centre National de la Recherche Scientifique
(CNRS), France

*CORRESPONDENCE

Alexandre Bittar
✉ abittar@idiap.ch

RECEIVED 14 June 2024

ACCEPTED 28 August 2024

PUBLISHED 25 September 2024

CITATION

Bittar A and Garner PN (2024) Exploring neural oscillations during speech perception via surrogate gradient spiking neural networks. *Front. Neurosci.* 18:1449181. doi: 10.3389/fnins.2024.1449181

COPYRIGHT

© 2024 Bittar and Garner. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). 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.

Exploring neural oscillations during speech perception via surrogate gradient spiking neural networks

Alexandre Bittar^{1,2*} and Philip N. Garner²

¹Idiap Research Institute, Audio Inference, Martigny, Switzerland, ²École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

Understanding cognitive processes in the brain demands sophisticated models capable of replicating neural dynamics at large scales. We present a physiologically inspired speech recognition architecture, compatible and scalable with deep learning frameworks, and demonstrate that end-to-end gradient descent training leads to the emergence of neural oscillations in the central spiking neural network. Significant cross-frequency couplings, indicative of these oscillations, are measured within and across network layers during speech processing, whereas no such interactions are observed when handling background noise inputs. Furthermore, our findings highlight the crucial inhibitory role of feedback mechanisms, such as spike frequency adaptation and recurrent connections, in regulating and synchronizing neural activity to improve recognition performance. Overall, on top of developing our understanding of synchronization phenomena notably observed in the human auditory pathway, our architecture exhibits dynamic and efficient information processing, with relevance to neuromorphic technology.

KEYWORDS

neural oscillations, spiking neural networks, speech recognition, brain-inspired computing, deep learning, surrogate gradient, spike-frequency adaptation, neuromorphic computing

1 Introduction

In the field of speech processing technologies, the effectiveness of training deep artificial neural networks (ANNs) with gradient descent has led to the emergence of many successful encoder-decoder architectures for automatic speech recognition (ASR), typically trained in an end-to-end fashion over vast amounts of data (Gulati et al., 2020; Baevski et al., 2020; Li et al., 2021; Radford et al., 2023). Despite recent efforts (Brodbeck et al., 2024; Millet et al., 2022; Millet and King, 2021; Magnuson et al., 2020) toward understanding how these ANN representations can compare with speech processing in the human brain, the cohesive integration of the fields of deep learning and neuroscience remains a challenge. Nonetheless, spiking neural networks (SNNs), a type of artificial neural network inspired by the biological neural networks in the brain, present an interesting convergence point of the two disciplines. Although slightly behind in terms of performance compared to ANNs, SNNs have recently achieved concrete progress (Hammouamri et al., 2024; Sun et al., 2023; Bittar and Garner, 2022a; Yin et al., 2021) on speech command recognition tasks

using the surrogate gradient method (Neftci et al., 2019) which allows them to be trained via gradient descent. Further work has also shown that they can be used to define a spiking encoder inside a hybrid ANN-SNN end-to-end trainable architecture on the more challenging task of large vocabulary continuous speech recognition (Bittar and Garner, 2022b). Their successful inclusion into contemporary deep learning ASR frameworks offers a promising path to bridge the existing gap between deep learning and neuroscience in the context of speech processing. This integration not only equips deep learning tools with the capacity to engage in speech neuroscience but also offers a scalable approach to simulate spiking neural dynamics, which supports the exploration and testing of hypotheses concerning the neural mechanisms and cognitive processes related to speech. This investigation of complex brain functions via physiologically inspired networks aligns with the work of Pulvermüller et al. (2021); Henningsen-Schomers and Pulvermüller (2022); Pulvermüller (2023), who applied biological constraints to large-scale simulations of language learning in SNNs. We complement their approach by training on more realistic speech data, albeit at the cost of some simplifications.

In neuroscience, various neuroimaging techniques such as electroencephalography (EEG) can detect rhythmic and synchronized postsynaptic potentials that arise from activated neuronal assemblies. These give rise to observable neural oscillations, commonly categorized into distinct frequency bands: delta (0.5–4 Hz), theta (4–8 Hz), alpha (8–13 Hz), beta (13–30 Hz), low-gamma (30–80 Hz), and high-gamma (80–150 Hz) (Buzsáki, 2006). It is worth noting that while these frequency bands provide a useful framework, their boundaries are not rigidly defined and can vary across studies. Nevertheless, neural oscillations play a crucial role in coordinating brain activity and are implicated in cognitive processes such as attention (Fries et al., 2001; Jensen and Colgin, 2007; Womelsdorf and Fries, 2007; Vinck et al., 2013), memory (Kucewicz et al., 2017), sensory perception (Başar et al., 2000; Senkowski et al., 2007), and motor function (MacKay, 1997; Ramos-Murguialday and Birbaumer, 2015). Of particular interest is the phenomenon of cross-frequency coupling (CFC) which reflects the interaction between oscillations occurring in different frequency bands (Jensen and Colgin, 2007; Jirsa and Müller, 2013). As reviewed in Abubaker et al. (2021), many studies have demonstrated a relationship between CFC and working memory performance (Tort et al., 2009; Axmacher et al., 2010). In particular phase-amplitude coupling (PAC) between theta and gamma rhythms appears to support memory integration (Buzsáki and Moser, 2013; Backus et al., 2016; Hummos and Nair, 2017), preservation of sequential order (Reddy et al., 2021; Colgin, 2013; Itskov et al., 2008) and information retrieval (Mizuseki et al., 2009). In contrast, alpha-gamma coupling commonly manifests itself as a sensory suppression mechanism during selective attention (Foxe and Snyder, 2011; Banerjee et al., 2011), inhibiting task-irrelevant brain regions (Jensen and Mazaheri, 2010) and ensuring controlled access to stored knowledge (Klimesch, 2012). Finally, beta oscillations are commonly associated with cognitive control and top-down processing (Engel et al., 2001).

In the context of speech perception, numerous investigations have revealed a similar oscillatory hierarchy, where the temporal organization of high-frequency signal amplitudes in the gamma

range is orchestrated by low-frequency neural phase dynamics, specifically in the delta and theta ranges (Canolty et al., 2006; Ghitza, 2011; Giraud and Poeppel, 2012; Hyafil et al., 2015; Attaheri et al., 2022). These three temporal scales—delta, theta and gamma—naturally manifest in speech and represent specific perceptual units. In particular, delta-range modulation (1–2 Hz) corresponds to perceptual groupings formed by lexical and phrasal units, encapsulating features such as the intonation contour of an utterance. Modulation within the theta-range aligns with the syllabic rate (4 Hz) around which the acoustic envelope consistently oscillates. Finally, (sub)phonemic attributes, including formant transitions that define the fine structure of speech signals, correlate with higher modulation frequencies (30–50 Hz) within the low-gamma range. The close correspondence between the perception of (sub)phonemic, syllabic and phrasal attributes on one hand, and the manifestation of gamma, theta and delta neural oscillations on the other, was notably emphasized in Giraud and Poeppel (2012). These different levels of temporal granularity inherent to speech signals therefore appear to be processed in a hierarchical fashion, with the intonation and syllabic contour encoded by earlier neurons guiding the encoding of phonemic features by later neurons. Some insights about how phoneme features end up being encoded in the temporal gyrus were given in Mesgarani et al. (2014). Drawing from recent research (Bonhage et al., 2017) on the neural oscillatory patterns associated with the sentence superiority effect, it is suggested that such low-frequency modulation may facilitate automatic linguistic chunking by grouping higher-order features into packets over time, thereby contributing to enhanced sentence retention. The engagement of working memory in manipulating phonological information enables the sequential retention and processing of speech sounds for coherent word and sentence representations. Additionally, alpha modulation has also been shown to play a role in improving auditory selective attention (Strauß et al., 2014b,a; Wöstmann et al., 2017), reflecting the listener's sensitivity to acoustic features and their ability to comprehend speech (Obleser and Weisz, 2012).

Computational models (Hyafil et al., 2015; Hovsepian et al., 2020) have shown that theta oscillations can indeed parse speech into syllables and provide a reliable reference time frame to improve gamma-based decoding of continuous speech. These approaches (Hyafil et al., 2015; Hovsepian et al., 2020) implement specific models for theta and gamma neurons along with a distinction between inhibitory and excitatory neurons. The resulting networks are then optimized to detect and classify syllables with very limited numbers of trainable parameters (10–20). In contrast, this work proposes to utilize significantly larger end-to-end trainable multi-layered architectures (400k–20M trainable parameters) where all neuron parameters and synaptic connections are optimized to predict sequences of phoneme/subword probabilities, that can subsequently be decoded into words. By avoiding constraints on theta or gamma activity, the approach allows us to explore which forms of CFC naturally arise when solely optimizing the decoding performance. Even though the learning mechanism is not biologically plausible, we expect that a model with sufficiently realistic neuronal dynamics and satisfying ASR performance should reveal similarities with the human brain. We divide our analysis in two parts,

- Architecture:** As a preliminary analysis, we conduct hyperparameter tuning to optimize the model's architectural parameters. On top of assessing the network's capabilities and scalability, we notably evaluate how the incorporation of spike-frequency adaptation (SFA) and recurrent connections impact the speech recognition performance.
- Oscillations:** We then explore the central aspect of our analysis concerning the emergence of neural oscillations within our model. Each SNN layer is treated as a distinct neuron population, from which spike trains are aggregated into a population signal similar to EEG data. Through intra- and inter-layer CFC analysis, we investigate the presence of significant delta-gamma, theta-gamma, alpha-gamma and beta-gamma PAC. We also investigate how incorporating Dale's law, SFA and recurrent connections affect the synchronization of neural activity.

2 Materials and methods

2.1 Spiking neuron model

Physiologically grounded neuron models such as the well-known Hodgkin and Huxley model (Hodgkin and Huxley, 1952) can be reduced to just two variables (FitzHugh, 1961; Morris and Lecar, 1981). More contemporary models, such as the Izhikevich (Izhikevich, 2003) and adaptive exponential integrate-and-fire (Brette and Gerstner, 2005) models, have similarly demonstrated the capacity to accurately replicate voltage traces observed in biological neurons using just membrane potential and adaptation current as essential variables (Badel et al., 2008). With the objective of incorporating realistic neuronal dynamics into large-scale neural network simulations with gradient descent training, the linear AdLIF neuron model stands out as an adequate compromise between physiological plausibility and computational efficiency. It can be described in continuous time by the following differential equations (Gerstner and Kistler, 2002),

$$\tau_u \dot{u}(t) = -(u(t) - u_{\text{rest}}) - R w(t) + R I(t) - \tau_u \sum_f \delta(t - t^f) \quad (1)$$

$$\tau_w \dot{w}(t) = -w(t) + a(u(t) - u_{\text{rest}}) + \tau_w b \sum_f \delta(t - t^f). \quad (2)$$

The neuron's internal state is characterized by the membrane potential $u(t)$ which linearly integrates stimuli $I(t)$ and gradually decays back to a resting value u_{rest} with time constant $\tau_u \in [3, 25]$ ms. A spike is emitted when the threshold value ϑ is attained, $u(t) \geq \vartheta$, denoting the firing time $t = t^f$, after which the potential is decreased by a fixed amount $\vartheta - u_r$. In the following, we will use $u_r = u_{\text{rest}}$ for simplicity. The second variable $w(t)$ is coupled to the potential with strength a and decay constant $\tau_w \in [30, 350]$ ms, characterizing sub-threshold adaptation. Additionally, $w(t)$ experiences an increase of b after a spike is emitted, which defines spike-triggered adaptation. The differential equations can

be simplified as,

$$\tau_u \dot{u}(t) = -u(t) - w(t) + I(t) - \tau_u \sum_f \delta(t - t^f) \quad (3)$$

$$\tau_w \dot{w}(t) = -w(t) + a u(t) + \tau_w b \sum_f \delta(t - t^f). \quad (4)$$

by making all time-dependent quantities dimensionless with changes of variables,

$$u \rightarrow \frac{u - u_{\text{rest}}}{\vartheta - u_{\text{rest}}}, \quad w \rightarrow \frac{R w}{\vartheta - u_{\text{rest}}} \quad \text{and} \quad I \rightarrow \frac{R I}{\vartheta - u_{\text{rest}}},$$

and redefining neuron parameters as,

$$a \rightarrow R a, \quad b \rightarrow \frac{R b}{\vartheta - u_{\text{rest}}}, \quad \vartheta \rightarrow \frac{\vartheta - u_{\text{rest}}}{\vartheta - u_{\text{rest}}} = 1$$

$$\text{and} \quad u_{\text{rest}} \rightarrow \frac{u_{\text{rest}} - u_{\text{rest}}}{\vartheta - u_{\text{rest}}} = 0.$$

This procedure gets rid of unnecessary parameters such as the resistance R , as well as resting, reset and threshold values, so that a neuron ends up being fully characterized by four parameters: τ_u , τ_w , a and b . As derived in Appendix, the differential equations can be solved in discrete time with step size Δt using a forward-Euler first-order exponential integrator method. After initializing $u_0 = w_0 = s_0 = 0$, and defining $\alpha := \exp\left(-\frac{\Delta t}{\tau_u}\right)$ and $\beta := \exp\left(-\frac{\Delta t}{\tau_w}\right)$, the neuronal dynamics can be solved by looping over time steps $t = 1, 2, \dots, T$ as,

$$u_t = \alpha(u_{t-1} - s_{t-1}) + (1 - \alpha)(I_t - w_{t-1}) \quad (5)$$

$$w_t = \beta(w_{t-1} + b s_{t-1}) + (1 - \beta) a u_{t-1} \quad (6)$$

$$s_t = (u_t \geq 1). \quad (7)$$

Stability conditions for the value of the coupling strength a are derived in Appendix. Additionally, we constrain the values of the neuron parameters to biologically plausible ranges (Gerstner and Kistler, 2002; Augustin et al., 2013),

$$\tau_u \in [3, 25] \text{ ms}, \quad \tau_w \in [30, 350] \text{ ms}, \quad a \in [-0.5, 5], \quad b \in [0, 2]. \quad (8)$$

This AdLIF neuron model is equivalent to a Spike Response Model (SRM) with continuous time kernel functions illustrated in Figure 1. The four neuron parameters τ_u , τ_w , a and b all characterize the shape of these two curves representing the membrane potential response to an input spike and to an emitted spike, respectively. A derivation of the kernel-based SRM formulation is presented in Appendix.

2.1.1 Spiking layers

The spiking dynamics described in Equations 5–7 can be computed in parallel for a layer of neurons. In a multi-layered network, the l -th layer receives as inputs a linear combination of the previous layer's outputs $s^{l-1} \in \{0, 1\}^{B \times T \times N^{l-1}}$ where B , T and N^{l-1} represent the batch size, number of time steps and number of neurons, respectively. Feedback from the l -th layer can also be implemented as a linear combination of its own outputs at the previous time step $s_{t-1}^l \in \{0, 1\}^{N^l}$, so that the overall neuron stimulus I_t^l for neurons in the l -th layer at time step t is computed as,

$$I_t^l = W^l s_t^{l-1} + V^l s_{t-1}^l. \quad (9)$$

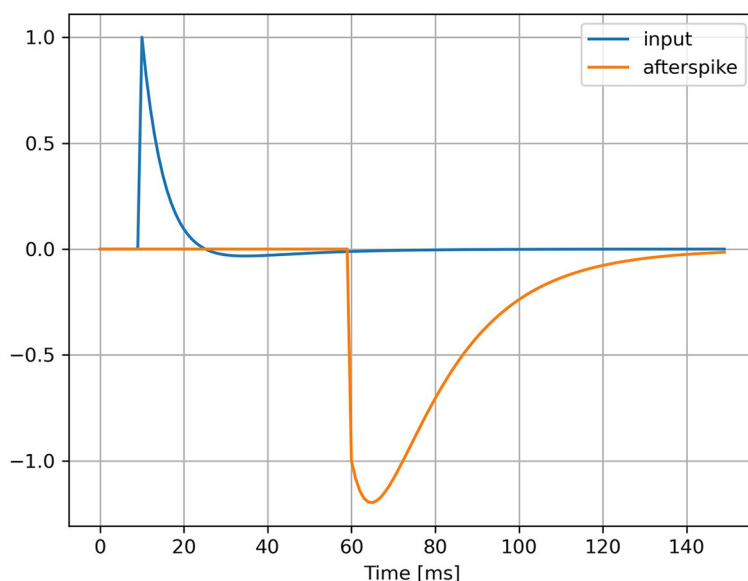


FIGURE 1
Kernel functions of AdLIF neuron model. Membrane potential response to an input pulse at $t = 10$ ms (in blue) and to an emitted spike at $t = 60$ ms (in orange). The neuron parameters are $\tau_u = 5$ ms, $\tau_w = 30$ ms, $a = 0.5$ and $b = 1.5$.

Here the feedforward $W^l \in \mathbb{R}^{N^{l-1} \times N^l}$ and feedback connections $V^l \in \mathbb{R}^{N^l \times N^l}$ are trainable parameters. Diagonal elements of V^l are set to zero as afterspike self inhibition is already taken into account in Equation 5. While this choice excludes autapses, which are rare but do exist in the brain, it simplifies the model for our study. Additionally, a binary mask can be applied to matrices W^l and V^l to limit the number of nonzero connections. Similarly, a portion of neurons in a layer can be reduced to leaky integrate-and-fire (LIF) dynamics without any SFA by applying another binary mask to the neuron adaptation parameters a and b . Indeed, if $a = b = 0$, the adaptation current vanishes $w_t = 0 \forall t \in \{1, 2, \dots, T\}$ and has no more impact on the membrane potential.

2.1.2 Surrogate gradient method

The main challenge in applying stochastic gradient descent to the derived neuronal dynamics stems from the threshold operation described in Equation 7, which has a derivative of zero everywhere except at the threshold point, where it is undefined. To address this, a surrogate derivative can be manually specified using PyTorch (Paszke et al., 2017), which enables the application of the Back-Propagation Through Time algorithm for training the resulting SNN in a manner similar to recurrent neural network (RNN) training. In this paper, we adopt the boxcar function as our surrogate function. This choice has been proven effective in various contexts (Kaiser et al., 2020; Bittar and Garner, 2022a,b) and requires minimal computational resources as expressed by the derivative definition,

$$\frac{\partial s_t}{\partial u_t} = \begin{cases} 0.5 & \text{if } |u_t - 1| \leq 0.5 \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

2.1.3 Spike frequency regularization

The firing rate $f_{b,n}^l$ of neuron n in layer l when processing utterance b can be calculated in Hz as,

$$f_{b,n}^l = \frac{1}{T_b} \sum_{t=1}^T s_{b,t,n}^l, \quad (11)$$

where T_b is the utterance duration in seconds. We regularize the firing rates of all spiking neurons between $f_{\min} = 0.5$ Hz and $f_{\max} = f_{\text{Nyquist}}$ using the following regularization loss,

$$\mathcal{L}_{\text{reg}} = \frac{1}{BL} \sum_{b=1}^B \sum_{l=1}^L \frac{1}{N^l} \sum_{n=1}^{N^l} \text{ReLU}(f_{\min} - f_{b,n}^l) + \text{ReLU}(f_{b,n}^l - f_{\max}), \quad (12)$$

to discourage neurons from remaining silent or from firing above the Nyquist frequency.

2.2 Overview of the auditory pathway in the brain

Sound waves are initially received by the outer ear and then transmitted as vibrations to the cochlea in the inner ear, where the basilar membrane allows for a representation of different frequencies along its length (Gundersen et al., 1978). Distinct sound frequencies induce localized membrane vibrations that activate adjacent inner hair cells. These specialized sensory cells, covering the entire basilar membrane, release neurotransmitters when activated, stimulating neighboring auditory nerve fibers and initiating the production of action potentials. Tonotopy is maintained through the conversion of mechanical motion into electric signals as each inner hair cell, tuned to a specific frequency,

only affects nearby auditory nerve fibers (Saenz and Langers, 2014). The resulting spike trains then propagate through a multi-layered neural network, ultimately reaching cortical regions associated with higher-order cognitive functions such as speech recognition. Overall, the auditory system is organized hierarchically, with each level contributing to the progressively more sophisticated processing of auditory information.

2.3 Simulated speech recognition pipeline

Our objective is to design a speech recognition architecture that, while sufficiently plausible for meaningful comparisons with neuroscience observations, remains simple and efficient to ensure compatibility with modern deep learning techniques and achieve good ASR performance. We implement the overall waveform-to-phoneme pipeline illustrated in Figure 2 inside the Speechbrain (Ravanelli et al., 2021) framework. We provide a description of each of its components here below.

2.3.1 Feature extractor

Mel filterbank features are extracted from the raw waveform using 80 filters and a 25 ms window with a shift of 2 ms. This procedure down samples the 16 kHz input speech signal to a 500 Hz spectrogram representation with 80 frequency bins.

2.3.2 Auditory CNN

A single-layered two-dimensional convolution module is applied to the 80 extracted Mel features using 16 channels, a kernel size of (7, 7), a padding of (7, 0) and a stride of 1, producing $16 \cdot ((80 - 7) + 1) = 1,184$ output signals with unchanged number of time steps. Layer normalization, drop out on the channel dimension and a Leaky-ReLU activation are then applied. Each produced signal characterizes the evolution over time of the spectral energy across a frequency band of seven consecutive Mel bins.

2.3.3 Auditory nerve fibers

Each 500 Hz output signal from the auditory CNN constitutes the stimulus of a single auditory nerve fiber, which converts the real-valued signal into a spike train. These nerve fibers are modeled as a layer of LIF neurons without recurrent connections and using a single trainable parameter per neuron, $\tau_u \in [3, 25]$ ms, representing the time constant of the membrane potential decay.

2.3.4 Multi-layered SNN

The resulting spike trains are sent to a fully connected multi-layered SNN architecture with 512 neurons in each layer. The proportion of neurons with nonzero adaptation parameters is controlled in each layer so that only a fraction of the neurons are AdLIF and the rest are LIF. Similarly the proportion of nonzero feedforward and recurrent connections is controlled in each layer by applying fixed random binary masks to the weight matrices. Compared to a LIF neuron, an AdLIF neuron has three additional

trainable parameters, $\tau_w \in [30, 350]$ ms, $a \in [-0.5, 5]$ and $b \in [0, 2]$, related to the adaptation variable coupled to the membrane potential.

2.3.5 Spikes to probabilities

The spike trains of the last layer are sent to an average pooling module which down samples their time dimension to 25 Hz. These are then projected to 512 phoneme features using two fully connected (FC) layers with Leaky-ReLU activation. A third FC layer with Log-Softmax activation finally projects them to 40 log-probabilities representing 39 phoneme classes and a blank token as required by connectionist temporal classification (CTC).

2.3.6 Training and inference

The log-probabilities are sent to a CTC loss (Graves et al., 2006) so that the parameters of the complete architecture can be updated through back propagation. Additionally, regularization of the firing rate as defined in Equation 12 is used to prevent neurons from being silent or firing above the Nyquist frequency. At inference, CTC decoding is used to output the most likely phoneme sequence from the predicted log-probabilities, and the phoneme error rate (PER) is computed to evaluate the model's performance.

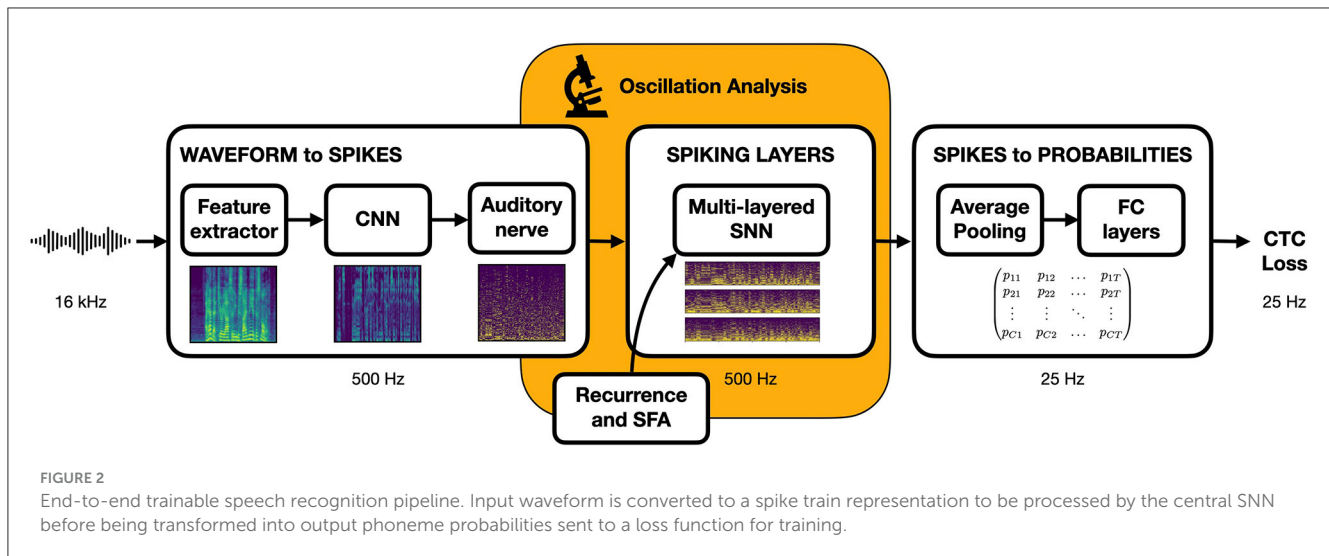
2.4 Physiological plausibility and limitations

2.4.1 Cochlea and inner hair cells

While some of the complex biological processes involved in converting mechanical vibrations to electric neuron stimuli can be abstracted, we assume that the key feature to retain is the tonotopic encoding of sound information. A commonly used metric in neuroscience is the ratio of characteristic frequency to bandwidth, which defines how sharply tuned a neuron is around the frequency it is most responsive to. As detailed in Supplementary Table S6, from measured Q10 values in normal hearing humans reported in Devi et al. (2022), we evaluate that a single auditory nerve fiber should receive inputs from 5 to 7 adjacent frequency bins when using 80 Mel filterbank features. The adoption of a Mel filterbank frontend can be justified by its widespread utilization within deep learning ASR frameworks. Although we do not attempt to directly model cochlear and hair cell processing, we can provide a rough analog in the form of Mel features passing through a trainable convolution module that yields plausible ranges of frequency sensitivity for our auditory nerve fibers.

2.4.2 Simulation time step

Modern ASR systems (Gulati et al., 2020; Radford et al., 2023) typically use a frame period of $\Delta t = 10$ ms during feature extraction, which is then often sub-sampled to 40 ms using a CNN before entering the encoder-decoder architecture. In the brain, typical minimal inter-spike distances imposed by a neuron's absolute refractory period can vary from 0 to 5 ms (Gerstner and Kistler, 2002). We therefore assume that using a time step >5 ms could result in dynamics that are less representative of



biological phenomena. Although using a time step $\Delta t < 1$ ms may yield biologically more realistic simulations, we opt for time steps ranging from 1 to 5 ms to ensure computational efficiency. After the SNN, the spike trains of the last layer are down-sampled to 25 Hz via average pooling on the time dimension. This prevents an excessive number of time steps from entering the CTC loss, which could potentially hinder its decoding efficacy. We use $\Delta t = 5$ ms for most of the hyperparameter tuning to reduce training time, but favor $\Delta t = 2$ ms for the oscillation analysis so that the full gamma range of interest (30–150 Hz) remains below the Nyquist frequency at 250 Hz.

2.4.3 Neuron model

The LIF neuron model is an effective choice for modeling auditory nerve fibers as it accurately represents their primary function of encoding sensory inputs into spike trains. We avoid using SFA and recurrent connections, as they are not prevalent characteristics of nerve fibers. On the other hand, for the multi-layered SNN, the linear AdLIF neuron model with layer-wise recurrent connections stands out as a good compromise between accurately reproducing biological firing patterns and remaining computationally efficient (Bittar and Garner, 2022a; Deckers et al., 2024). Although less popular than the moving threshold formulation by Bellec et al. (2018), recently reviewed in Ganguly et al. (2024), our implementation of SFA using the AdLIF model combines spike-triggered adaptation with subthreshold coupling. Previous work (Bittar and Garner, 2022a) demonstrated that the AdLIF outperforms moving threshold implementations (Yin et al., 2021; Salaj et al., 2021; Shaban et al., 2021; Yin et al., 2020) in speech command recognition tasks. Nevertheless, we will still implement and train an additional model with moving threshold SFA to ensure that our conclusions hold consistently across different SFA models.

2.4.4 Organization in layers

Similarly to ANNs, our simulation incorporates a layered organization, which facilitates the progressive extraction and

representation of features from low-order to higher-order, without the need of concretely defining and distinguishing neuron populations. This fundamental architectural principle aligns with the general hierarchical processing observed in biological brains. However, it oversimplifies the complexities of auditory processing, which extends beyond a straightforward sequential framework. While there is some sort of sequential processing in sub-cortical structures, the levels of processed features are more intricate than a simple hierarchy. This simplification is made to ensure compatibility with deep learning frameworks.

2.4.5 Layer-wise recurrence

While biological efferent pathways in the brain involve complex and widespread connections that span across layers and regions, modeling such intricate connectivity can introduce computational challenges and complexity, potentially hindering training and scalability. By restricting feedback connections to layer-wise recurrence, we simplify the network architecture and enhance compatibility with deep learning frameworks.

2.4.6 Excitatory and inhibitory

In the neuroscience field, neurons are commonly categorized into two types: excitatory neurons, which stimulate action potentials in postsynaptic neurons, and inhibitory neurons, which reduce the likelihood of spike production in postsynaptic neurons. This principle is referred to as Dale's law.

In ANNs, weight matrices are commonly initialized with zero mean and a symmetric distribution, so that the initial number of excitatory and inhibitory connections is balanced. During training, synaptic connections are updated across all layers without enforcing a distinction between excitatory and inhibitory neurons. Dale's law can nevertheless be imposed (Li et al., 2024; Cornford et al., 2021) even if it typically results in slightly reduced performance.

In our baseline model, Dale's law is not applied, so that similarly to standard ANNs, weight matrices are trained without

constraining values to be positive or negative. Additionally, we train separate SNNs with Dale's law to evaluate its impact on neural oscillations. In this setup, half the neurons are excitatory and half are inhibitory, while the auditory nerve fibers are all excitatory.

2.4.7 Delays

In biological neural networks, the propagation time of spikes between neurons introduces delays, primarily due to axonal transmission. To incorporate and assess the impact of these delays on neural oscillations, we additionally train separate SNNs using dilated convolutions in the temporal dimension instead of fully connected feedforward matrices. This approach, introduced in Hammouamri et al. (2024), allows us to introduce controlled delays directly into the network architecture. Based on their configuration for speech command recognition tasks, we use a maximum delay value of 300 ms.

2.4.8 Learning rule

Stochastic gradient descent, though biologically implausible due to its global and offline learning framework, allows us to leverage parallelizable and fast computations to optimize larger-scale neural networks. While this approach facilitates effective training and scaling, it diverges from biologically inspired synaptic plasticity mechanisms, such as those mediated by AMPA and NMDA receptors.

2.4.9 Decoding into phoneme sequences

Although lower PERs could be achieved with a more sophisticated decoder, our primary focus is on analyzing the spiking layers within the encoder. For simplicity, we therefore opt for straightforward CTC decoding, which more directly reflects the encoder's capabilities.

2.4.10 Hybrid ANN-SNN balance

The CNN module in the ASR frontend as well as the ANN module (average pooling and FC layers) converting spikes to probabilities are intentionally kept simple to give most of the processing and representational power to the central SNN on which focuses our neural oscillations analysis.

2.5 Speech processing tasks

The following datasets are used in our study.

- The TIMIT dataset (Garofolo et al., 1993) provides a comprehensive and widely utilized collection of phonetically balanced American English speech recordings from 630 speakers with detailed phonetic transcriptions and word alignments. It represents a standardized benchmark for evaluating ASR model performance. The training, validation and test sets contain 3,696, 400, and 192 sentences, respectively. Utterance durations vary between 0.9 and 7.8 s. Due to its compact size of ~ 5 h of speech data, the

TIMIT dataset is well-suited for investigating suitable model architectures and tuning hyperparameters. It is however considered small for ASR hence the use of Librispeech presented below.

- The Librispeech corpus (Panayotov et al., 2015) contains about 1,000 h of English speech audiobook data read by over 2,400 speakers with utterance durations between 0.8 and 35 s. Given its significantly larger size, we only train a few models selected on TIMIT to confirm that our analysis holds when scaling to more data.
- The Google Speech Commands dataset (Warden, 2018) contains one-second audio recordings of 35 spoken commands such as “yes,” “no,” “stop,” “go,” “left,” “right” “up”. The training, validation and testing splits contain ~ 85 , 10, and 5k examples, respectively. It is used to test whether similar CFCs arise when simply recognizing single words instead of phoneme or subword sequences.

When evaluating an ASR model, the error rate signifies the proportion of incorrectly predicted words or phonemes. The count of successes in a binary outcome experiment, such as ASR testing, can be effectively modeled using a binomial distribution. In the context of trivial priors, the posterior distribution of the binomial distribution follows a beta distribution. By leveraging the equal-tailed 95% credible intervals derived from the posterior distribution, we establish error bars, yielding a range of $\pm 0.8\%$ for the reported PERs on the TIMIT dataset, about $\pm 0.2\%$ for the reported word error rates on LibriSpeech, and about $\pm 0.4\%$ for the reported accuracy on Google Speech Commands.

2.6 Analysis methods

2.6.1 Hyperparameter tuning

Before reporting results on the oscillation analysis, we investigate the optimal architecture by tuning some relevant hyperparameters. All experiments are run using a single NVIDIA GeForce RTX 3090 GPU. On top of assessing their respective impact on the error rate, we test if more physiologically plausible design choices correlate with better performance. Here is the list of the fixed parameters that we do not modify in our reported experiments:

- number of Mel bins: 80
- Mel window size: 25 ms
- auditory CNN kernel size (7, 7)
- auditory CNN stride: (1, 1)
- auditory CNN padding: (7, 0)
- average pooling size: $40/\Delta t$
- number of phoneme FC layers: 2
- number of phoneme features: 512
- dropout: 0.15
- activation: LeakyReLU

where for CNN attributes of the form (n_t, n_f) , n_t and n_f correspond to time and feature dimensions, respectively. The tunable parameters are the following,

- filter bank hop size controlling the SNN time step in ms: {1, 2, 5}
- number of auditory CNN channels (filters): {8, 16, 32, 64, 128}
- number of SNN layers: {1, 2, 3, 4, 5, 6, 7}
- neurons per SNN layer: {64, 128, 256, 512, 768, 1024, 1536, 2048}
- proportion of neurons with SFA: [0, 1]
- feedforward connectivity: [0, 1]
- recurrent connectivity: [0, 1]

While increasing the number of neurons per layer primarily impacts memory requirements, additional layers mostly extend training time.

2.6.2 Population signal

In the neuroscience field, EEG stands out as a widely employed and versatile method for studying brain activity. By placing electrodes on the scalp, this non-invasive technique measures the aggregate electrical activity resulting from the synchronized firing of neurons within a specific brain region. An EEG signal therefore reflects the summation of postsynaptic potentials from a large number of neurons operating in synchrony. The typical sampling rate for EEG data is commonly in the range of 250–1,000 Hz which matches our desired simulation time steps. With our SNN, we do not have EEG signals but directly the individual spike trains of all neurons in the architecture. In order to perform similar population-level analyses, we sum the binary spike trains $s_b^l \in \{0, 1\}^{T \times N^l}$ emitted by all neurons in a specific layer l for a single utterance b as follows,

$$p_{b,t}^l = \sum_{n=1}^{N^l} s_{b,t,n}^l \quad (13)$$

Before performing the PAC analysis, the resulting population activity signal $p_{b,t}^l$ is then normalized over the time dimension with a mean of 0 and a standard deviation of 1, yielding the normalized population signal $\hat{p}_{b,t}^l$ defined as,

$$\hat{p}_{b,t}^l = \frac{p_{b,t}^l - \mu_b^l}{\sigma_b^l}, \quad (14)$$

where μ_b^l is the mean and σ_b^l is the standard deviation of $p_{b,t}^l$ over the time dimension.

2.6.3 Phase-amplitude coupling

Using finite impulse response band-pass filters, the obtained population signals are decomposed into different frequency ranges. We study CFC in the form of PAC both within a single population and across layers. This technique assesses whether a relationship exists between the phase of a low frequency signal and the envelope (amplitude) of a high frequency signal. As recommended in Hülsemann et al. (2019), we implement both the modulation index (Tort et al., 2008) and mean vector length (Canolty et al., 2006) metrics to quantify the observed amount of PAC. For each measure type, the observed coupling value is compared to a distribution of 10,000 surrogates to assess the significance.

Surrogate couplings are computed by disrupting the temporal order of the amplitude time series while preserving its overall characteristics. Specifically, the amplitude time series is permuted by cutting it at a random point and reversing the order of the two segments. This method, as discussed by Hülsemann et al. (2019), maintains all inherent properties of the original data except for the temporal relationship between phase angle and amplitude magnitude, providing a conservative test of significance. A p-value is then obtained by fitting a Gaussian function on the distribution of surrogate coupling values and calculating the area under the curve for values greater than the observed coupling value.

As pointed out in Jones (2016), it is important to note that observed oscillations can exhibit complexities such as non-sinusoidal features and brief transient events on single trials. Such nuances become aggregated when averaging signals, leading to the widely observed continuous rhythms. We therefore perform all analysis on single utterances.

For intra-layer interactions, a single population signal is used to extract both the low-frequency oscillation phase and the high-frequency oscillation amplitude. In a three-layered architecture, these interactions include nerve-nerve, first layer-first layer, second layer-second layer, and third layer-third layer couplings.

For inter-layer interactions, we consider couplings between the low-frequency oscillation phase in one layer and the high-frequency oscillation amplitude in all subsequent layers. These interactions include nerve-first layer, nerve-second layer, nerve-third layer, first layer-second layer, first layer-third layer, second layer-third layer couplings.

For all aforementioned intra- and inter-layer combinations, we use delta (0.5–4 Hz), theta (4–8 Hz), alpha (8–13 Hz), and beta (13–30 Hz) ranges as low-frequency modulating bands, and low-gamma (30–80 Hz) and high-gamma (80–150 Hz) ranges as high-frequency modulated bands. For a given model, we iterate through the 64 longest utterances in the TIMIT test set. For each utterance, we consider the 10 aforementioned intra- and inter-layer relations, as well as the eight possible combinations of low-frequency to high-frequency bands. We conduct PAC testing on each of the 5,120 resulting coupling scenarios, and only consider a coupling to be significant when both modulation index and mean vector length metrics yield p-values below 0.05.

3 Results

3.1 Architectural analysis

In order to draw a comparison with the human auditory pathway, we have introduced the physiologically inspired ASR pipeline illustrated in Figure 2. The proposed hybrid ANN-SNN architecture is trained in an end-to-end fashion on the TIMIT dataset (Garofolo et al., 1993) to predict phoneme sequences from speech waveforms. In the architectural design, we aimed to minimize the complexity of ANN components and favor the central SNN which will be the focus of the oscillation analysis. Here as a preliminary step, we examine how relevant hyperparameters affect the PER. On top of assessing the scalability of our approach to larger

TABLE 1 Hyperparameter tuning for the number of SNN layers and neurons per layer on the TIMIT dataset.

Number of layers and neurons per layer	Number of parameters	Test PER (%)	Validation PER (%)
1 × 512	740k/1.3M	23.3	21.8
2 × 512	1.1/1.7M	21.0	19.2
3 × 512	1.5/2.1M	20.5	18.2
4 × 512	1.9/2.5M	20.2	17.4
5 × 512	2.3/2.9M	20.0	17.6
6 × 512	2.7/3.3M	20.0	17.9
7 × 512	3.1/3.7M	20.5	18.0
3 × 64	91/394k	30.9	29.6
3 × 128	211/547k	25.5	24.1
3 × 256	537/938k	22.5	20.9
3 × 768	3.0/3.7M	19.6	17.4
3 × 1,024	4.9/5.7M	19.1	17.1
3 × 1,536	10.1/11.2M	19.0	17.3
3 × 2,048	17.1/18.5M	19.2	17.2
No nerve, no SNN	0/873k	34.2	32.0

The second column gives both the number of trainable parameters in the multi-layered SNN (left) and in the whole encoder (right). The PERs are reported after 50 training epochs using a 5 ms time step, 16 CNN channels, 50% of AdLIF neurons, 100% feedforward and 50% recurrent connectivity. The performance of the architecture when removing the SNN is also reported (bottom). Bold values indicate the lowest achieved PERs.

networks, we identify the importance of the interplay between recurrence and SFA.

3.1.1 Network scalability

As reported in Table 1, performance improves with added layers, peaking at 4–6 layers before declining, which suggests a significant contribution to the final representation from all layers within this range. Compared to conventional non-spiking RNN encoders used in ASR, our results support the scalability of surrogate gradient SNNs to relatively deep architectures. Additionally, augmenting the number of neurons until about 1,000 per layer consistently yields lower PERs, beyond which performance saturates.

3.1.2 Recurrent connections and spike-frequency adaptation

The impact of adding SFA in the neuron model as well as using recurrent connections are reported in Table 2. Interestingly, we find that without SFA, optimal performance is achieved by limiting the recurrent connectivity to 80%. When additionally using SFA, further limitation of the recurrent connectivity about 50% yields the lowest PER. This differs from conventional non-spiking RNNs, where employing FC recurrent matrices is favored. These results indicate that while requiring fewer parameters, an architecture with sparser recurrent connectivity and more selective parameter usage can achieve better task performance.

Overall, SFA and recurrent connections individually yield significant error rate reduction, although they, respectively, grow as $\mathcal{O}(N)$ and $\mathcal{O}(N^2)$ with the number of neurons N . In

line with previous studies on speech command recognition tasks (Perez-Nieves et al., 2021; Bittar and Garner, 2022a), our results emphasize the metabolic and computational efficiency gained by harnessing the heterogeneity of adaptive spiking neurons. Furthermore, effectively calibrating the interplay between unit-wise feedback from SFA and layer-wise feedback from recurrent connections appears crucial for achieving optimal performance.

3.1.3 Enforcing Dale's law

To align with common ANN practice, the previous results were obtained without restricting neurons to be either strictly excitatory or strictly inhibitory. We now train more physiologically inspired models that satisfy Dale's law, with results presented in Table 3. Although ASR performance decreased (1–4% absolute PER increase), this may simply be due to suboptimal weight initialization, which is known to affect performance (Li et al., 2024). This could likely be mitigated in future work by using the approach of Rossbroich et al. (2022), who derived fluctuation-driven initialization schemes compatible with Dale's law.

3.1.4 Supplementary findings

We observe in Supplementary Table S1 that decreasing the simulation time step does not affect the performance. Although making the simulation of spiking dynamics more realistic, one might have anticipated that backpropagating through more time steps could hinder the training and worsen the performance as observed in standard RNNs often suffering from vanishing or

TABLE 2 Ablation experiments for the recurrent connectivity and proportion of neurons with SFA on the TIMIT dataset.

Model type	Recurrent connectivity	Proportion of AdLIF neurons	Number of parameters	Test PER (%)	Validation PER (%)
No recurrence no SFA	0	0	1.1/1.7M	26.9	24.8
Recurrence only	0.2	0	1.3/1.8M	22.0	20.1
	0.5	0	1.5/2.1M	21.0	18.9
	0.8	0	1.8/2.3M	20.8	18.7
	1	0	1.9/2.5M	21.8	19.3
SFA only	0	0.2	1.1/1.7M	24.2	21.7
	0	0.5	1.1/1.7M	23.7	21.6
	0	0.8	1.1/1.7M	23.3	21.0
	0	1	1.1/1.7M	22.9	21.1
Recurrence and SFA	0.2	0.2	1.3/1.8M	20.9	19.3
	0.5	0.5	1.5/2.1M	20.5	18.2
	0.8	0.8	1.8/2.3M	21.2	18.8
	1	1	1.9/2.5M	23.3	21.5

PERs are reported after 50 epochs using a 5 ms time step, 16 CNN channels, three layers, 512 neurons per layer and 100% feedforward connectivity. Bold values indicate the lowest achieved PERs.

TABLE 3 Ablation experiments for the recurrent connectivity and proportion of neurons with SFA on the TIMIT dataset when additionally using Dale's law.

Model type	Recurrent connectivity	Proportion of AdLIF neurons	Number of parameters	Test PER (%)	Validation PER (%)
No recurrence no SFA	0	0	1.1/1.7M	30.7	28.7
Recurrence only	0.5	0	1.5/2.1M	23.6	20.6
SFA only	0	0.5	1.1/1.7M	25.1	22.9
Recurrence and SFA	0.5	0.5	1.5/2.1M	21.2	19.2

PERs are reported after 50 epochs using a 2 ms time step, 16 CNN channels, three layers, 512 neurons per layer, 100% feedforward connectivity and an excitatory-inhibitory ratio of 1. Bold values indicate the lowest achieved PERs.

exploding gradients (Bengio et al., 1994). With inputs ranging from 1,000 to over 7,000 steps using 1 ms intervals on TIMIT, our results demonstrate a promising scalability of surrogate gradient SNNs for processing longer sequences.

Secondly, as reported in Supplementary Table S2, we did not observe substantial improvement when increasing the number of auditory nerve fibers past ~5,000, even though there are ~30,000 of them in the human auditory system. This could be due to both the absence of a proper model for cochlear and hair cell processing in our pipeline and to the relatively low number of neurons (<1,000) in the subsequent layer.

As detailed in Supplementary Table S3, reduced feedforward connectivity in the SNN led to poorer overall performance. This contrasts with our earlier findings on recurrent connectivity, highlighting the distinct functional roles of feedforward and feedback mechanisms in the network.

Additionally, we incorporated trainable delays by replacing the fully connected feedforward matrices with dilated convolutions over the temporal dimension. While including delays resulted in similar overall performance, using the groups parameter to control connectivity led to more parameter-efficient models, as shown in Supplementary Table S4. This method of reducing connectivity

proved more effective than our previous approach of randomly masking fully connected matrices.

Finally, Supplementary Table S5 shows our results when using the more popular moving threshold formulation of SFA instead of the AdLIF model. Consistent with our previous findings on speech command recognition (Bittar and Garner, 2022a), the AdLIF implementation of SFA outperforms its moving threshold alternative, with the same number of parameters.

3.2 Oscillation analysis

Based on our previous architectural results that achieved satisfactory speech recognition performance using a physiologically inspired model, we hypothesize that the spiking dynamics of a trained network should, to some extent, replicate those occurring throughout the auditory pathway. Our investigation aims to discern if synchronization phenomena resembling brain rhythms manifest within the implemented SNNs as they process speech utterances to recognize phonemes.

3.2.1 Synchronized gamma activity produces low-frequency rhythms

As illustrated in [Figure 3A](#), the spike trains produced by passing a test-set speech utterance through the trained architecture exhibit distinct low-frequency rhythmic features in all layers. By looking at the histogram of single neuron firing rates illustrated in [Figure 3B](#), we observe that the distribution predominantly peaks at gamma range, with little to no activity below beta. This reveals that the low-frequency oscillations visible in [Figure 3A](#) actually emerge from the synchronization of gamma-active neurons. The resulting low-frequency rhythms appear to follow to some degree the intonation and syllabic contours of the input filterbank features and to persist across layers. Compared to the three subsequent layers, higher activity in the auditory nerve comes from the absence of inhibitory SFA and recurrence mechanisms. These initial observations suggest that the representation of higher-order features in the last layer is temporally modulated by lower level features already encoded in the auditory nerve fibers, even though each layer is seen to exhibit distinct rhythmic patterns. In the next section, we focus on measuring this modulation more rigorously via PAC analysis.

3.2.2 Phase-amplitude couplings within and across layers

By aggregating over the relevant spike trains, we compute distinct EEG-like population signals for the auditory nerve fibers and each of the three SNN layers. These are then filtered in the different frequency bands, as illustrated in [Figure 4A](#), which allows us to perform CFC analyses. We measure PAC within-layer and across-layers between all possible combinations of frequency bands and find multiple significant forms of coupling for every utterance.

An example of significant theta low-gamma coupling between the auditory nerve fibers and the last layer is illustrated in [Figure 4B](#). Here input low-frequency modulation is observed to significantly modulate the output gamma activity. This indicates that the network integrates and propagates intonation and syllabic contours across layers through synchronized neural activity along these perceptual cues.

On the majority of utterances, we found significant CFCs between the input waveform and the population signal of the last layer. It is important to note that the synchronization of neural signals to the auditory envelope emerged without imposing any theta or gamma activity in our network. The optimization of the PER combined with sufficiently realistic spiking neuronal dynamics therefore represent sufficient conditions to reproduce some broad properties of neural oscillations observed in the brain, suggesting a general functional role of facilitating information processing.

The activity of the final layer of the SNN stands out as the most significantly modulated overall. By architectural construction, modulation in that final layer has the greatest impact on the ASR task, as the spike trains from this layer are converted to phoneme probabilities using a small ANN. A higher number of couplings in the final layer correlates with a decrease in the PER. This suggests that CFCs may be associated with more selective integration of phonetic features, enhanced attentional processes, as well as improved assimilation of contextual information.

Alpha-band oscillations were the most frequently measured, consistent with biological evidence that the alpha rhythm is the most prominent oscillation in spontaneous EEG ([Berger, 1929](#)).

More generally, the patterns of coupling between different neural populations and frequency bands were found to differ from one utterance to another. These variations indicate that the neural processing of our network is highly dynamic and depends on the acoustic properties and linguistic content of the input. The observed rich range of intra- and inter-layer couplings suggests that the propagation of low-level input features such as intonation and syllabic rhythm is only one aspect of these synchronized neural dynamics.

3.2.3 Impact of Dale's law on oscillations

As illustrated in [Table 4](#), SNNs that satisfy Dale's law display significantly higher numbers of CFCs. In biological neural networks, this principle contributes to a more structured and organized form of network dynamics, as each neuron is either excitatory or inhibitory but not both. When applying Dale's principle to SNNs, it constrains the network with more defined roles for neurons, leading to more coherent oscillatory patterns and more CFCs.

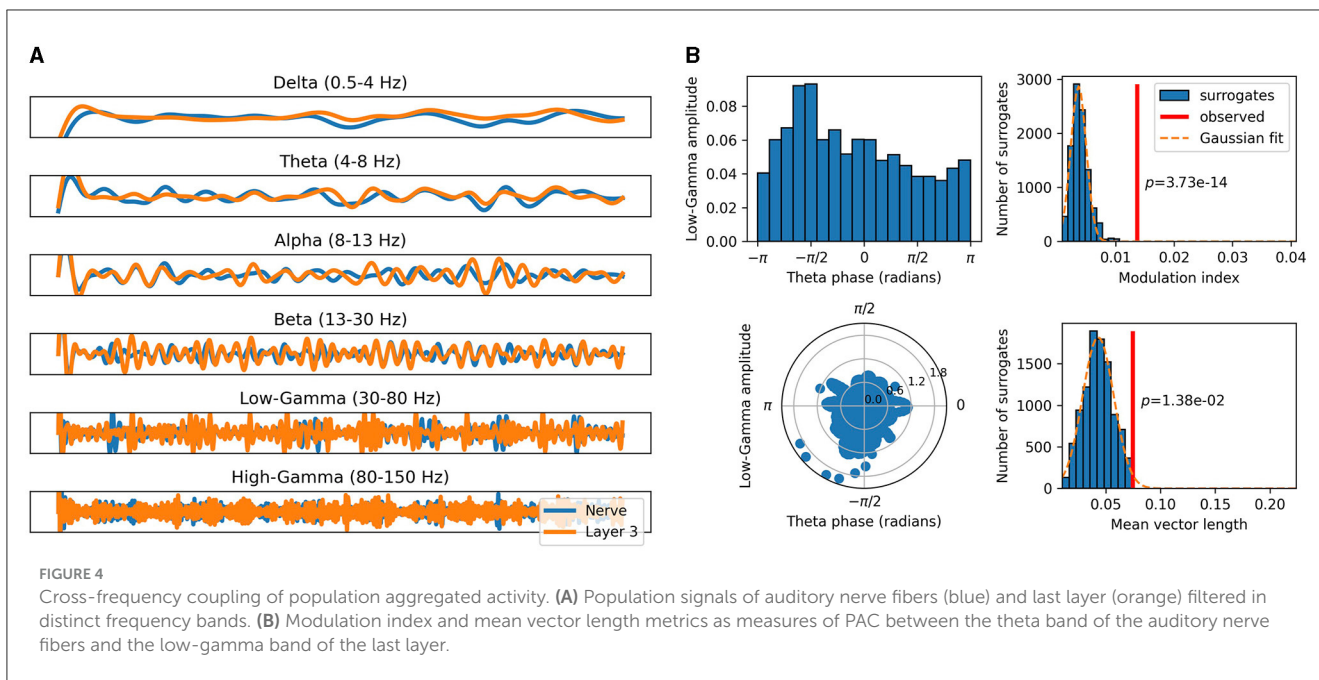
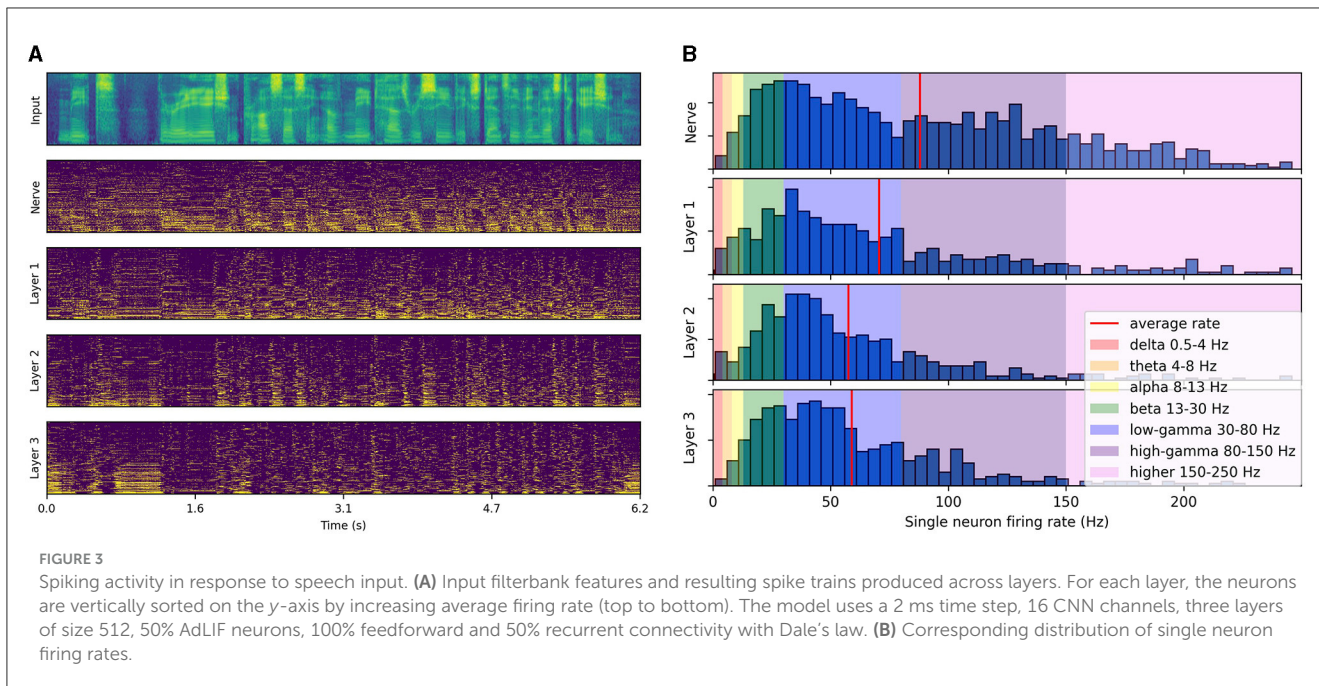
In this study, we use equivalent models for both excitatory and inhibitory neurons, and simply allow them to adapt their parameters within the same fixed ranges during training. We observe that trained values of both membrane and adaptation time constants τ_u and τ_w converge to lower average values across the inhibitory populations compared to the excitatory ones. This suggests that, even with initially equivalent models, the network naturally differentiates the dynamics of excitatory and inhibitory neurons to fulfill their distinct functional roles. Future work could focus on better defining these two types of neurons from the outset, incorporating more biologically plausible initial parameter ranges.

When using Dale's law, the fixed excitatory-inhibitory (E/I) ratio plays a crucial role in shaping the neuronal dynamics inside the SNN. In our study, we observed that increasing the proportion of inhibitory neurons ($E/I = 0.33$) resulted in a similar ASR performance ($\sim 1\%$ absolute PER difference) compared to the standard ratio ($E/I = 1$), but with a reduced average firing rate of 56 Hz instead of 68 Hz for equivalent models with 50% of neurons with SFA and 50% of recurrent connectivity (see [Table 4](#)). This finding suggests that a higher ratio of inhibitory neurons can achieve comparable task performance while maintaining a lower overall level of network activity.

3.2.4 Impact of spike-frequency adaptation and recurrent connections on oscillations

Across all model types (AdLIF or moving threshold SFA, with or without delays and Dale's law), both SFA and recurrent connections had an overall inhibitory effect, typically reducing the average network firing rate from around 100 Hz to roughly 60 Hz (see [Table 4](#)).

This regularization of the network activity appears to enable more effective parsing and encoding of speech information, as it reliably led to improved ASR performance. While both forms of feedback exhibit an overall inhibitory effect, SFA operates at the



individual neuron level whereas recurrent connections act at the layer level.

SFA is known to encourage and stabilize the synchronization of cortical networks (Crook et al., 1998) and to promote periodic signal propagation (Augustin et al., 2013). In our results, this effect is pronounced in SNNs without Dale's law, where the inclusion of SFA was consistently associated with a significant increase in the number of measured CFCs. However, when Dale's law is applied, the overall number of CFCs is significantly higher with no noticeable impact of SFA on CFCs. This suggests that the stricter constraints imposed by Dale's law may lead to more uniform behavior in CFCs, thereby reducing the observed influence of SFA.

In models with and without Dale's law, incorporating recurrent connections was consistently associated with a decrease in the number of inter-layer couplings, indicating more localized synchronization.

Finally, using the moving threshold formulation of SFA produces a lower firing rate and fewer significant PACs compared to our AdLIF baseline. The lower number of PACs might mean that the moving threshold formulation is less effective than the AdLIF at coordinating these interactions, which could explain the inferior task performance. Nevertheless, the discrepancy between the two approaches could potentially be narrowed with further hyperparameter optimization.

TABLE 4 Effect of Dale's law, SFA, recurrence and delays on oscillatory activity.

Model type	Model configuration	Test PER (%)	Firing rate (Hz)	Number of intra-layer PACs	Number of inter-layer PACs
AdLIF baseline	No recurrence no SFA	26.9	98	133	192
	SFA only	23.7	72	213	390
	Recurrence only	21.0	76	180	132
	Recurrence and SFA	20.5	71	265	177
AdLIF with Dale	No recurrence no SFA	30.7	100	513	643
	SFA only	25.1	77	564	603
	Recurrence only	23.6	67	432	278
	Recurrence and SFA	21.2	68	329	290
	Recurrence and SFA, E/I=0.33	22.3	56	526	496
AdLIF with delays	No recurrence no SFA	26.2	92	116	54
	SFA only	23.3	77	91	53
	Recurrence only	21.6	68	205	66
	Recurrence and SFA	20.8	68	253	60
Moving threshold	SFA only	26.3	56	136	131
	Recurrence and SFA	22.1	62	182	135

Comparison of the oscillatory activity resulting from passing the 64 longest TIMIT test-set utterances through different types of trained networks. The last two columns show the total number of significant PACs summed across all 64 utterances and frequency bands, for intra- and inter-layer relations, respectively. All networks use a 2 ms time step, 16 CNN channels, three layers, 512 neurons per layer and 100% feedforward connectivity (i.e., groups = 1 for delays). The E/I ratio is 1 for models with Dale's law, except the last one where it is 0.33.

3.2.5 Impact of delays on oscillations

As illustrated in Table 4, SNNs with delays produced significantly fewer PACs, especially for inter-layer couplings, compared to the baseline with fully connected feedforward matrices. Introducing trainable delays through dilated convolutions allows for temporal dispersion of signals, which may desynchronize neural activity and explain the observed reduction in CFCs.

3.2.6 Effects of training and input type on neuronal activity

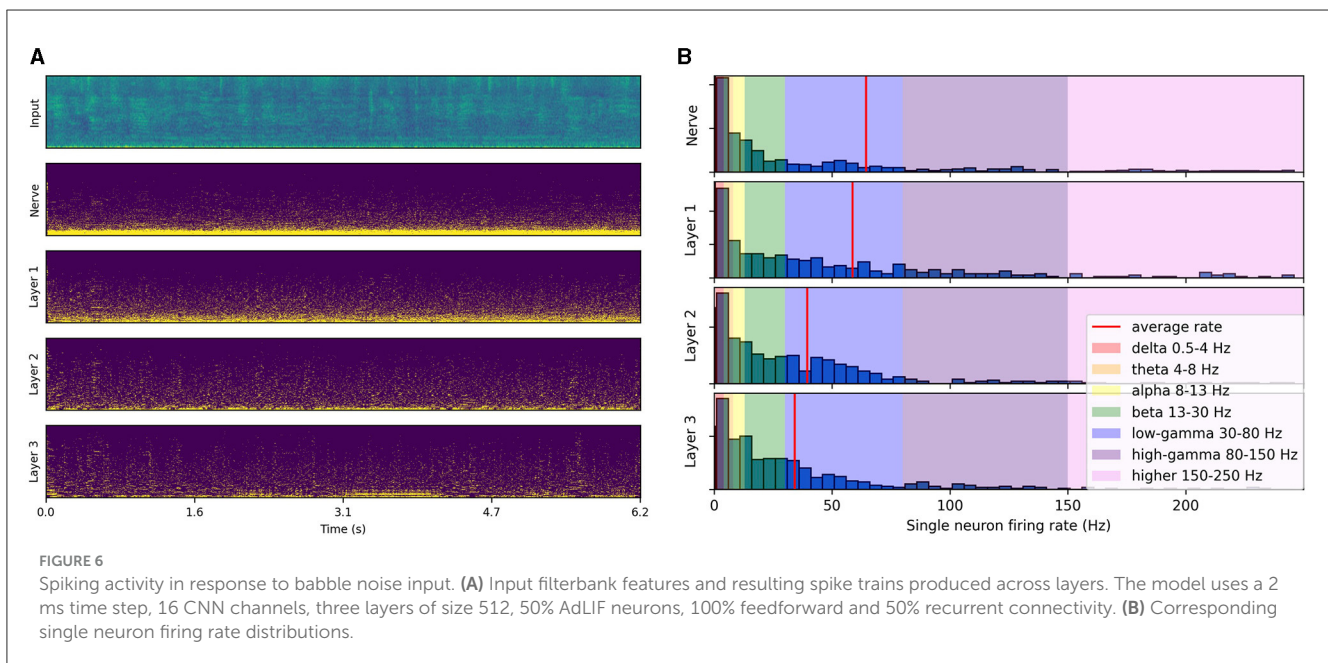
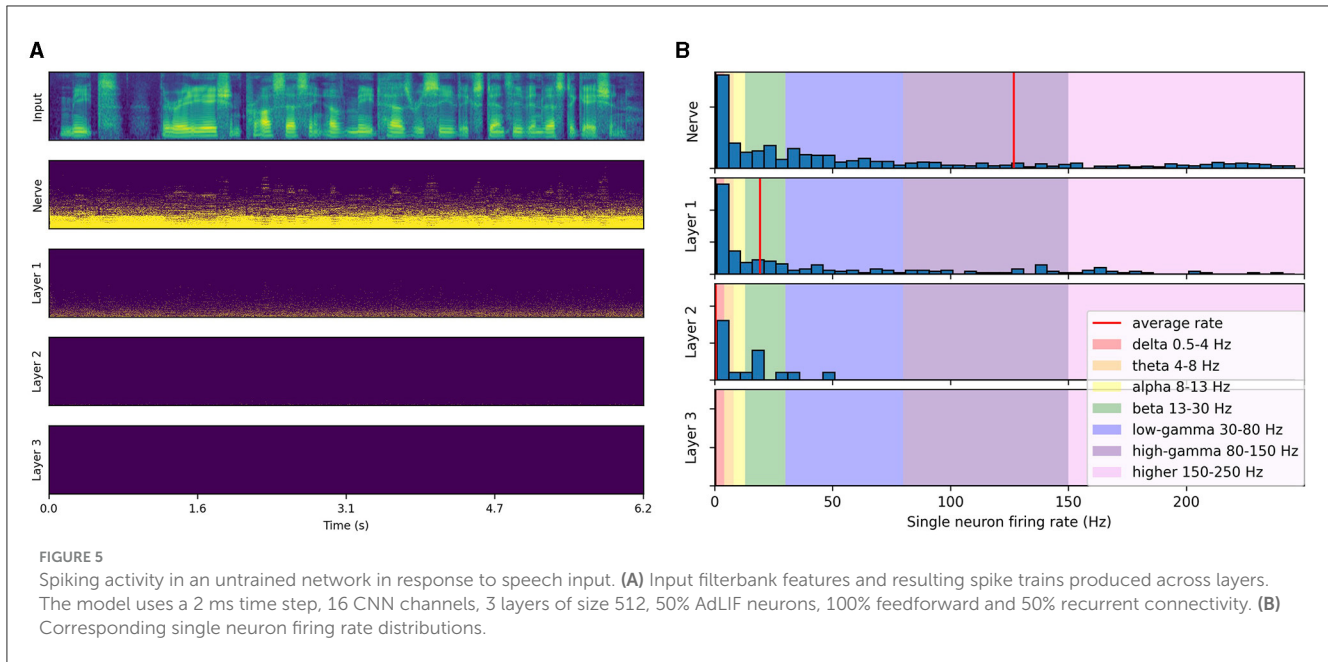
In order to further understand the emergence of coupled signals, we consider passing speech through an untrained network, as well as passing different types of noise inputs through a trained network.

Trained architectures exhibit persisting neuronal activity across layers compared to untrained ones, where the activity almost completely decays after the first layer, as illustrated in Figure 5. This decay across layers persists even when increasing the input magnitude up to saturating auditory nerve fibers. This phenomenon can be attributed to the random weights in untrained architectures, which transform structured input patterns into uncorrelated noise, leading to vanishing neuronal activity in deeper layers. Our CFC analysis shows no significant coupling, even in layers with sufficient spiking activity, i.e., within the auditory nerve population and the first layer.

In trained networks, noise inputs lead to single neuron firing rate distributions peaking at very low rates and where the activity gradually decreases across layers, as illustrated in Figure 6B. This contrasts with the response to speech inputs seen in Figure 3B where the activity was sustained across layers with most of the distribution in the gamma range. We tested uniform noise as well as different noise sources (air conditioner, babble, copy machine and typing) from the MS-SNSD dataset (Reddy et al., 2019). Compared to a speech input, all noise types yielded reduced average firing rates (from 60 Hz to about 40 Hz) with most of the neurons remaining silent. This highly dynamic processing of information is naturally efficient at attenuating its activity when processing noise or any input that does not induce sufficient synchronization. Interestingly, babble noises were found in certain cases to induce some significant PAC patterns, whereas other noise types resulted in no coupling at all. Even though babble noises resemble speech and produced some form of synchronization, they only triggered a few neurons per layer (see Figure 6). Overall, we showed that synchronicity of neural oscillations in the form of PAC results from training and is only triggered when passing an appropriate speech input.

3.2.7 Scaling to a larger dataset

Our approach was extended to the Librispeech dataset (Panayotov et al., 2015) with 960 h of training data. After 8 epochs, we reached 9.5% word error rate on the *test-clean* data split. As



observed on TIMIT, trained models demonstrated similar CFCs in their spiking activity.

3.2.8 Training on speech command recognition task

With our experimental setup, the encoder is directly trained to recognize phonemes on TIMIT and subwords on Librispeech. One could therefore assume that the coupled gamma activity emerges from that constraint. In order to test this hypothesis, we run additional experiments on a speech command recognition task where no phoneme or subword recognition is imposed by the training. Instead the model is directly trained to recognize a set of

short words. We use the same architecture as on TIMIT, except the average pooling layer is replaced by a readout layer as defined in Bittar and Garner (2022a) which reduces the temporal dimension altogether, as required by the speech command recognition task. Interestingly, using speech command classes as ground truths still produces significant PAC patterns, especially in the last layer. These results indicate that the emergence of the studied rhythms does not require phoneme-based training and may be naturally emerging from speech processing. Using the second version of the Google Speech Commands data set (Warden, 2018) with 35 classes, we achieve a test set accuracy of 97.05%, which, to the best of our knowledge, improves upon the current state-of-the-art performance using SNNs of 95.35% (Hammouamri et al., 2024).

4 Discussion

In this study, we introduced a physiologically inspired speech recognition architecture, centered around an SNN, and designed to be compatible with modern deep learning frameworks. As set out in the introduction, we first explored the capabilities and scalability of the proposed speech recognition architecture before analyzing neural oscillations.

Our preliminary architectural analysis demonstrated a satisfactory level of scalability to deeper and wider networks, as well as to longer sequences and larger datasets. This scalability was achieved through our approach of utilizing the surrogate gradient method to incorporate an SNN into an end-to-end trainable speech recognition pipeline. In addition, our ablation experiments emphasized the importance of including SFA within the neuron model, along with layer-wise recurrent connections, to attain optimal recognition performance. Notably, our implementation of SFA using the AdLIF model outperformed the more popular moving threshold formulation, which corroborates our previous results on speech command recognition (Bittar and Garner, 2022a).

The subsequent analysis of the spiking activity across our trained networks in response to speech stimuli revealed that neural oscillations, commonly associated with various cognitive processes in the brain, did emerge from training an architecture to recognize words or phonemes. Through CFC analyses, we measured similar synchronization phenomena to those observed throughout the human auditory pathway. During speech processing, trained networks exhibited several forms of PAC, including delta-gamma, theta-gamma, alpha-gamma, and beta-gamma, while no such coupling occurred when processing pure background noise. Our networks' ability to synchronize oscillatory activity in the last layer was also associated with improved speech recognition performance, which points to a functional role for neural oscillations in auditory processing. Even though we employ gradient descent training, which does not represent a biologically plausible learning algorithm, our approach was capable of replicating natural phenomena of macro-scale neural coordination. By leveraging the scalability offered by deep learning frameworks, our approach can therefore serve as a valuable tool for studying the emergence and role of brain rhythms.

Building upon the main outcome of replicating neural oscillations, our analysis on SFA and recurrent connections emphasized their key role in actively shaping neural responses and driving synchronization via inhibition in support of efficient auditory information processing. Our results point toward further work on investigating more realistic feedback mechanisms including efferent pathways across layers. More accurate neuron populations could also be obtained using clustering algorithms.

Further analysis incorporated Dale's law which constrains neurons to be exclusively excitatory or inhibitory. This physiologically inspired principle proved to be a crucial consideration as it significantly increased the number of measured oscillations.

Aside from the fundamental aspect of developing the understanding of biological processes, our research on SNNs also holds significance for the fields of neuromorphic computing

and energy efficient technology. Our exploration of the spiking mechanisms that drive dynamic and efficient information processing in the brain is particularly relevant for low-power audio and speech processing applications, such as on-device keyword spotting. In particular, the absence of synchronization in our architecture when handling background noise results in fewer computations, making our approach well-suited for always-on models.

Data availability statement

All datasets used to conduct our study are publicly available. The TIMIT dataset, along with relevant access information, can be found on the Linguistic Data Consortium website at <https://catalog.ldc.upenn.edu/LDC93S1>. The Librispeech and Google Speech Commands datasets are directly available at <https://www.openslr.org/12> and https://www.tensorflow.org/datasets/catalog/speech_commands respectively. To facilitate the advancement of spiking neural networks, we have made our code available open source at <https://github.com/idiap/sparse>.

Author contributions

AB: Conceptualization, Formal analysis, Investigation, Methodology, Software, Visualization, Writing – original draft. PG: Conceptualization, Funding acquisition, Project administration, Resources, Supervision, Writing – review & editing.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article. This project received funding under NAST: Neural Architectures for Speech Technology, Swiss National Science Foundation grant 185010 (<https://data.snf.ch/grants/grant/185010>).

Acknowledgments

We would like to thank the reviewers for their insightful suggestions, notably regarding the incorporation of Dale's law. All recommendations led to significant improvements of the manuscript.

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.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of

their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnins.2024.1449181/full#supplementary-material>

References

- Abubaker, M., Al Qasem, W., and Kvašňák, E. (2021). Working memory and cross-frequency coupling of neuronal oscillations. *Front. Psychol.* 12:756661. doi: 10.3389/fpsyg.2021.756661
- Attaheri, A., Choidealbha, Á. N., Di Liberto, G. M., Rocha, S., Brusini, P., Mead, N., et al. (2022). Delta-and theta-band cortical tracking and phase-amplitude coupling to sung speech by infants. *Neuroimage* 247:118698. doi: 10.1016/j.neuroimage.2021.118698
- Augustin, M., Ladenbauer, J., and Obermayer, K. (2013). How adaptation shapes spike rate oscillations in recurrent neuronal networks. *Front. Comput. Neurosci.* 7:9. doi: 10.3389/fncom.2013.00009
- Axmacher, N., Henseler, M. M., Jensen, O., Weinreich, I., Elger, C. E., and Fell, J. (2010). Cross-frequency coupling supports multi-item working memory in the human hippocampus. *Proc. Nat. Acad. Sci. U. S. A.* 107, 3228–3233. doi: 10.1073/pnas.0911531107
- Backus, A. R., Schoffelen, J.-M., Szebényi, S., Hanslmayr, S., and Doeller, C. F. (2016). Hippocampal-prefrontal theta oscillations support memory integration. *Curr. Biol.* 26, 450–457. doi: 10.1016/j.cub.2015.12.048
- Badel, L., Lefort, S., Berger, T. K., Petersen, C. C., Gerstner, W., and Richardson, M. J. (2008). Extracting non-linear integrate-and-fire models from experimental data using dynamic I-V curves. *Biol. Cybern.* 99, 361–370. doi: 10.1007/s00422-008-0259-4
- Baevski, A., Zhou, Y., Mohamed, A., and Auli, M. (2020). wav2vec 2.0: a framework for self-supervised learning of speech representations. *Adv. Neural Inf. Process. Syst.* 33, 12449–12460. doi: 10.48550/arXiv.2006.11477
- Banerjee, S., Snyder, A. C., Molholm, S., and Foxe, J. J. (2011). Oscillatory alpha-band mechanisms and the deployment of spatial attention to anticipated auditory and visual target locations: supramodal or sensory-specific control mechanisms? *J. Neurosci.* 31, 9923–9932. doi: 10.1523/JNEUROSCI.4660-10.2011
- Başar, E., Başar-Eroğlu, C., Karakaş, S., and Schürmann, M. (2000). Brain oscillations in perception and memory. *Int. J. Psychophysiol.* 35, 95–124. doi: 10.1016/S0167-8760(99)00047-1
- 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, Vol. 31*, eds. S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (Montréal, QC: Curran Associates, Inc.), 1412–1421.
- Bengio, Y., Simard, P., and Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Transact. Neural Netw.* 5, 157–166. doi: 10.1109/72.279181
- Berger, H. (1929). Über das Elektroencephalogramm des Menschen. *Arch. Psychiatr. Nervenkrankheiten* 87, 527–570. doi: 10.1007/BF01797193
- Bittar, A., and Garner, P. N. (2022a). A surrogate gradient spiking baseline for speech command recognition. *Front. Neurosci.* 16:865897. doi: 10.3389/fnins.2022.865897
- Bittar, A., and Garner, P. N. (2022b). Surrogate gradient spiking neural networks as encoders for large vocabulary continuous speech recognition. *arXiv* [preprint]. doi: 10.48550/arXiv.2212.01187
- Bonhage, C. E., Meyer, L., Gruber, T., Friederici, A. D., and Mueller, J. L. (2017). Oscillatory EEG dynamics underlying automatic chunking during sentence processing. *Neuroimage* 152, 647–657. doi: 10.1016/j.neuroimage.2017.03.018
- Brette, R., and Gerstner, W. (2005). Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.* 94, 3637–3642. doi: 10.1152/jn.00686.2005
- Brodbeck, C., Hannagan, T., and Magnuson, J. S. (2024). Recurrent neural networks as neuro-computational models of human speech recognition. *bioRxiv*. doi: 10.1101/2024.02.20.580731
- Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford: Oxford University Press.
- Buzsáki, G., and Moser, E. I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat. Neurosci.* 16, 130–138. doi: 10.1038/nn.3304
- Canolty, R. T., Edwards, E., Dalal, S. S., Soltani, M., Nagarajan, S. S., Kirsch, H. E., et al. (2006). High gamma power is phase-locked to theta oscillations in human neocortex. *Science* 313, 1626–1628. doi: 10.1126/science.1128115
- Colgin, L. L. (2013). Mechanisms and functions of theta rhythms. *Annu. Rev. Neurosci.* 36, 295–312. doi: 10.1146/annurev-neuro-062012-170330
- Cornford, J., Kalajdziewski, D., Leite, M., Lamarquette, A., Kullmann, D. M., and Richards, B. A. (2021). “Learning to live with dale’s principle: {ANN}s with separate excitatory and inhibitory units,” in *International Conference on Learning Representations* (Vienna).
- Crook, S. M., Ermentrout, G. B., and Bower, J. M. (1998). Spike frequency adaptation affects the synchronization properties of networks of cortical oscillators. *Neural Comput.* 10, 837–854. doi: 10.1162/089976698300017511
- Deckers, L., Van Damme, L., Van Leekwijck, W., Tsang, I. J., and Latré, S. (2024). Co-learning synaptic delays, weights and adaptation in spiking neural networks. *Front. Neurosci.* 18:1360300. doi: 10.3389/fnins.2024.1360300
- Devi, N., Sreeraj, K., Anuroopa, S., Ankitha, S., and Namitha, V. (2022). Q10 and tip frequencies in individuals with normal-hearing sensitivity and sensorineural hearing loss. *Ind. J. Otol.* 28, 126–129. doi: 10.4103/indianjotol.indianjotol_5_22
- Engel, A. K., Fries, P., and Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. *Nat. Rev. Neurosci.* 2, 704–716. doi: 10.1038/35094565
- FitzHugh, R. (1961). Impulses and physiological states in theoretical models of nerve membrane. *Biophys. J.* 1:445. doi: 10.1016/S0006-3495(61)86902-6
- Foxe, J. J., and Snyder, A. C. (2011). The role of alpha-band brain oscillations as a sensory suppression mechanism during selective attention. *Front. Psychol.* 2:10747. doi: 10.3389/fpsyg.2011.00154
- Fries, P., Reynolds, J. H., Rorie, A. E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291, 1560–1563. doi: 10.1126/science.1055465
- Ganguly, C., Bezugam, S. S., Abs, E., Payvand, M., Dey, S., and Suri, M. (2024). Spike frequency adaptation: bridging neural models and neuromorphic applications. *Commun. Eng.* 3:22. doi: 10.1038/s44172-024-00165-9
- Garofolo, J. S., Lamel, L. F., Fisher, W. M., Fiscus, J. G., and Pallett, D. S. (1993). *DARPA TIMIT Acoustic-Phonetic Continuous Speech Corpus CD-ROM. NIST Speech Disc 1-1.1*. NASA STI/Recon Technical Report, 93:27403. Philadelphia: Linguistic Data Consortium.
- Gerstner, W., and Kistler, W. M. (2002). *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge: Cambridge University Press.
- Ghitza, O. (2011). Linking speech perception and neurophysiology: speech decoding guided by cascaded oscillators locked to the input rhythm. *Front. Psychol.* 2:130. doi: 10.3389/fpsyg.2011.00130
- Giraud, A.-L., and Poeppel, D. (2012). Cortical oscillations and speech processing: emerging computational principles and operations. *Nat. Neurosci.* 15, 511–517. doi: 10.1038/nn.3063
- Graves, A., Fernández, S., Gomez, F., and Schmidhuber, J. (2006). “Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks,” in *Proceedings of the 23rd International Conference on Machine Learning* (Pittsburgh, PA), 369–376.
- Gulati, A., Qin, J., Chiu, C.-C., Parmar, N., Zhang, Y., Yu, J., et al. (2020). “Conformer: convolution-augmented transformer for speech recognition,” in *Interspeech* (Shanghai), 5036–5040.
- Gundersen, T., Skarstein, Ø., and Sikkeland, T. (1978). A study of the vibration of the basilar membrane in human temporal bone preparations by the use of the mossbauer effect. *Acta Otolaryngol.* 86, 225–232. doi: 10.3109/00016487809124740
- Hammouamri, I., Khalfouji-Hassani, I., and Masquelier, T. (2024). “Learning delays in spiking neural networks using dilated convolutions with learnable spacings,” in *The Twelfth International Conference on Learning Representations* (Vienna).

- Henningsen-Schomers, M. R., and Pulvermüller, F. (2022). Modelling concrete and abstract concepts using brain-constrained deep neural networks. *Psychol. Res.* 86, 2533–2559. doi: 10.1007/s00426-021-01591-6
- Hodgkin, A. L., and Huxley, A. F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *J. Physiol.* 117:500. doi: 10.1113/jphysiol.1952.sp004764
- Hovsepian, S., Olasagasti, I., and Giraud, A.-L. (2020). Combining predictive coding and neural oscillations enables online syllable recognition in natural speech. *Nat. Commun.* 11:3117. doi: 10.1038/s41467-020-16956-5
- Hülsemann, M. J., Naumann, E., and Rasch, B. (2019). Quantification of phase-amplitude coupling in neuronal oscillations: comparison of phase-locking value, mean vector length, modulation index, and generalized-linear-modeling-cross-frequency-coupling. *Front. Neurosci.* 13:573. doi: 10.3389/fnins.2019.00573
- Hummos, A., and Nair, S. S. (2017). An integrative model of the intrinsic hippocampal theta rhythm. *PLoS ONE* 12:e0182648. doi: 10.1371/journal.pone.0182648
- Hyafil, A., Fontolan, L., Kabdebon, C., Gutkin, B., and Giraud, A.-L. (2015). Speech encoding by coupled cortical theta and gamma oscillations. *Elife* 4:e06213. doi: 10.7554/eLife.06213
- Itskov, V., Pastalkova, E., Mizuseki, K., Buzsáki, G., and Harris, K. D. (2008). Theta-mediated dynamics of spatial information in hippocampus. *J. Neurosci.* 28, 5959–5964. doi: 10.1523/JNEUROSCI.5262-07.2008
- Izhikevich, E. M. (2003). Simple model of spiking neurons. *IEEE Transact. Neural Netw.* 14, 1569–1572. doi: 10.1109/TNN.2003.820440
- Jensen, O., and Colgin, L. L. (2007). Cross-frequency coupling between neuronal oscillations. *Trends Cogn. Sci.* 11, 267–269. doi: 10.1016/j.tics.2007.05.003
- Jensen, O., and Mazaheri, A. (2010). Shaping functional architecture by oscillatory alpha activity: gating by inhibition. *Front. Hum. Neurosci.* 4:186. doi: 10.3389/fnhum.2010.00186
- Jirsa, V., and Müller, V. (2013). Cross-frequency coupling in real and virtual brain networks. *Front. Comput. Neurosci.* 7:78. doi: 10.3389/fncom.2013.00078
- Jones, S. R. (2016). When brain rhythms aren't 'rhythmic': implication for their mechanisms and meaning. *Curr. Opin. Neurobiol.* 40, 72–80. doi: 10.1016/j.conb.2016.06.010
- Kaiser, J., Mostafa, H., and Neftci, E. (2020). Synaptic plasticity dynamics for deep continuous local learning (DECOLLE). *Front. Neurosci.* 14:424. doi: 10.3389/fnins.2020.00424
- Klimesch, W. (2012). Alpha-band oscillations, attention, and controlled access to stored information. *Trends Cogn. Sci.* 16, 606–617. doi: 10.1016/j.tics.2012.10.007
- Kucewicz, M. T., Berry, B. M., Kremen, V., Brinkmann, B. H., Sperling, M. R., Jobst, B. C., et al. (2017). Dissecting gamma frequency activity during human memory processing. *Brain* 140, 1337–1350. doi: 10.1093/brain/awx043
- Li, B., Pang, R., Sainath, T. N., Gulati, A., Zhang, Y., Qin, J., et al. (2021). "Scaling end-to-end models for large-scale multilingual ASR," in *Automatic Speech Recognition and Understanding Workshop (ASRU)* (Cartagena: IEEE), 1011–1018.
- Li, P., Cornford, J., Ghosh, A., and Richards, B. (2024). Learning better with Dale's law: a spectral perspective. *Adv. Neural Inf. Process. Syst.* 36:546924. doi: 10.1101/2023.06.28.546924
- MacKay, W. A. (1997). Synchronized neuronal oscillations and their role in motor processes. *Trends Cogn. Sci.* 1, 176–183. doi: 10.1016/S1364-6613(97)01059-0
- Magnuson, J. S., You, H., Luthra, S., Li, M., Nam, H., Escabi, M., et al. (2020). Earshot: a minimal neural network model of incremental human speech recognition. *Cogn. Sci.* 44:e12823. doi: 10.1111/cogs.12823
- Mesgarani, N., Cheung, C., Johnson, K., and Chang, E. F. (2014). Phonetic feature encoding in human superior temporal gyrus. *Science* 343, 1006–1010. doi: 10.1126/science.1245994
- Millet, J., Caucheteux, C., Boubenec, Y., Gramfort, A., Dunbar, E., Pallier, C., et al. (2022). Toward a realistic model of speech processing in the brain with self-supervised learning. *Adv. Neural Inf. Process. Syst.* 35, 33428–33443. doi: 10.48550/arXiv.2206.01685
- Millet, J., and King, J.-R. (2021). Inductive biases, pretraining and fine-tuning jointly account for brain responses to speech. *arXiv [preprint]*. doi: 10.31219/osf.io/fq6gd
- Mizuseki, K., Sirota, A., Pastalkova, E., and Buzsáki, G. (2009). Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* 64, 267–280. doi: 10.1016/j.neuron.2009.08.037
- Morris, C., and Lecar, H. (1981). Voltage oscillations in the barnacle giant muscle fiber. *Biophys. J.* 35, 193–213. doi: 10.1016/S0006-3495(81)84782-0
- Neftci, E. O., Mostafa, H., and Zenke, F. (2019). Surrogate gradient learning in spiking neural networks: bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Process. Mag.* 36, 51–63. doi: 10.1109/MSP.2019.2931595
- Obleser, J., and Weisz, N. (2012). Suppressed alpha oscillations predict intelligibility of speech and its acoustic details. *Cereb. Cortex* 22, 2466–2477. doi: 10.1093/cercor/bhr325
- Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. (2015). "Librispeech: an ASR corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (Brisbane, QLD: IEEE), 5206–5210.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., et al. (2017). "Automatic differentiation in pytorch," in *NIPS Workshops* (Long Beach, CA).
- Perez-Nieves, N., Leung, V. C., Dragotti, P. L., and Goodman, D. F. (2021). Neural heterogeneity promotes robust learning. *Nat. Commun.* 12:5791. doi: 10.1038/s41467-021-26022-3
- Pulvermüller, F. (2023). Neurobiological mechanisms for language, symbols and concepts: clues from brain-constrained deep neural networks. *Progr. Neurobiol.* 230:102511. doi: 10.1016/j.pneurobio.2023.102511
- Pulvermüller, F., Tomasello, R., Henningsen-Schomers, M. R., and Wennekers, T. (2021). Biological constraints on neural network models of cognitive function. *Nat. Rev. Neurosci.* 22, 488–502. doi: 10.1038/s41583-021-00473-5
- Radford, A., Kim, J. W., Xu, T., Brockman, G., McLeavey, C., and Sutskever, I. (2023). "Robust speech recognition via large-scale weak supervision," in *International Conference on Machine Learning* (Hawaii: PMLR), 28492–28518.
- Ramos-Murguialday, A., and Birbaumer, N. (2015). Brain oscillatory signatures of motor tasks. *J. Neurophysiol.* 113, 3663–3682. doi: 10.1152/jn.00467.2013
- Ravanelli, M., Parcollet, T., Plantinga, P., Rouhe, A., Cornell, S., Lugosch, L., et al. (2021). SpeechBrain: a general-purpose speech toolkit. *arXiv [preprint]*. doi: 10.48550/arXiv.2106.04624
- Reddy, C. K., Beyrami, E., Pool, J., Cutler, R., Srinivasan, S., and Gehrke, J. (2019). "A scalable noisy speech dataset and online subjective test framework," in *Interspeech* (Graz), 1816–1820.
- Reddy, L., Self, M. W., Zoefel, B., Poncet, M., Possel, J. K., Peters, J. C., et al. (2021). Theta-phase dependent neuronal coding during sequence learning in human single neurons. *Nat. Commun.* 12:4839. doi: 10.1038/s41467-021-25150-0
- Rossbroich, J., Gygax, J., and Zenke, F. (2022). Fluctuation-driven initialization for spiking neural network training. *Neuromorp. Comp. Eng.* 2:044016. doi: 10.1088/2634-4386/ac97bb
- Saenz, M., and Langers, D. R. (2014). Tonotopic mapping of human auditory cortex. *Hear. Res.* 307, 42–52. doi: 10.1016/j.heares.2013.07.016
- Salaj, D., Subramoney, A., Krašniković, C., Bellec, G., Legenstein, R., and Maass, W. (2021). Spike frequency adaptation supports network computations on temporally dispersed information. *Elife* 10:e65459. doi: 10.7554/eLife.65459
- Senkowski, D., Talsma, D., Grigutsch, M., Herrmann, C. S., and Woldorff, M. G. (2007). Good times for multisensory integration: effects of the precision of temporal synchrony as revealed by gamma-band oscillations. *Neuropsychologia* 45, 561–571. doi: 10.1016/j.neuropsychologia.2006.01.013
- Shaban, A., Bezugam, S. S., and Suri, M. (2021). An adaptive threshold neuron for recurrent spiking neural networks with nanodevice hardware implementation. *Nat. Commun.* 12, 1–11. doi: 10.1038/s41467-021-24427-8
- Strauß, A., Kotz, S. A., Scharinger, M., and Obleser, J. (2014a). Alpha and theta brain oscillations index dissociable processes in spoken word recognition. *Neuroimage* 97, 387–395. doi: 10.1016/j.neuroimage.2014.04.005
- Strauß, A., Wöstmann, M., and Obleser, J. (2014b). Cortical alpha oscillations as a tool for auditory selective inhibition. *Front. Hum. Neurosci.* 8:350. doi: 10.3389/fnhum.2014.00350
- Sun, P., Chua, Y., Devos, P., and Botteldoorn, D. (2023). Learnable axonal delay in spiking neural networks improves spoken word recognition. *Front. Neurosci.* 17:1275944. doi: 10.3389/fnins.2023.1275944
- Tort, A. B., Komorowski, R. W., Manns, J. R., Kopell, N. J., and Eichenbaum, H. (2009). Theta-gamma coupling increases during the learning of item-context associations. *Proc. Nat. Acad. Sci. U. S. A.* 106, 20942–20947. doi: 10.1073/pnas.0911331106
- Tort, A. B., Kramer, M. A., Thorn, C., Gibson, D. J., Kubota, Y., Graybiel, A. M., et al. (2008). Dynamic cross-frequency couplings of local field potential oscillations in rat striatum and hippocampus during performance of a T-maze task. *Proc. Nat. Acad. Sci. U. S. A.* 105, 20517–20522. doi: 10.1073/pnas.0810524105
- Vinck, M., Womelsdorf, T., Buffalo, E. A., Desimone, R., and Fries, P. (2013). Attentional modulation of cell-class-specific gamma-band synchronization in awake monkey area v4. *Neuron* 80, 1077–1089. doi: 10.1016/j.neuron.2013.08.019
- Warden, P. (2018). Speech commands: a dataset for limited-vocabulary speech recognition. *arXiv [preprint]*. doi: 10.48550/arXiv.1804.03209

Womelsdorf, T., and Fries, P. (2007). The role of neuronal synchronization in selective attention. *Curr. Opin. Neurobiol.* 17, 154–160. doi: 10.1016/j.conb.2007.02.002

Wöstmann, M., Lim, S.-J., and Obleser, J. (2017). The human neural alpha response to speech is a proxy of attentional control. *Cereb. Cortex* 27, 3307–3317. doi: 10.1093/cercor/bhx074

Yin, B., Corradi, F., and Bohté, S. M. (2020). “Effective and efficient computation with multiple-timescale spiking recurrent neural networks,” in *International Conference on Neuromorphic Systems* (Oak Ridge, TN), 1–8.

Yin, B., Corradi, F., and Bohté, S. M. (2021). Accurate and efficient time-domain classification with adaptive spiking recurrent neural networks. *Nat. Mach. Intell.* 3, 905–913. doi: 10.1038/s42256-021-00397-w