

Brain-inspired, interpretable, resonant recurrent neural networks

Mark A. Kramer 

Department of Mathematics and Statistics and Center for Systems Neuroscience, [Boston University](#), Boston, Massachusetts 02215, USA
 and Department of Applied Mathematics and Statistics, [Johns Hopkins University](#), Baltimore, Maryland 21218, USA



(Received 23 June 2025; accepted 9 May 2026; published 1 June 2026)

Traditional artificial neural networks consist of nodes with nonoscillatory dynamics. Biological neural networks, on the other hand, consist of oscillatory components embedded in an oscillatory environment. Motivated by this feature of biological neurons, we describe a neural network framework with explicit damped, oscillatory node dynamics. We express the oscillatory dynamics using two history-dependent terms to connect these dynamics with standard recurrent neural network formulations, and apply physical constraints from observed brain dynamics to choose the oscillatory frequencies and stationary constraints to reduce the number of free parameters. We then optimize and illustrate network performance by classifying hand-written digits and simulated neuronal spike train activity and show that these oscillatory network elements support accurate classification with few trainable parameters. Choosing oscillator frequencies according to a proposed theory for brain rhythms improves classification accuracy compared to alternative frequency configurations and compared to standard recurrent neural network frameworks with comparable numbers of parameters. Compared to existing approaches, the proposed *resonant* recurrent network (RRN) utilizes oscillatory dynamics expressed as a straightforward extension of standard recurrent neural networks, produces interpretable features for classification, and performs well with few parameters when oscillator frequencies follow a configuration observed *in vivo*. We propose that RRNs may serve as efficient, biologically inspired building blocks to achieve complex goals in biological and artificial neural networks.

DOI: [10.1103/phlp-rhny](https://doi.org/10.1103/phlp-rhny)

I. INTRODUCTION

Biological systems rely on stable dynamics, which typically oscillate when perturbed as compensatory mechanisms prevent divergent activity. For example, a tree perturbed by the wind or blood glucose perturbed by sugar oscillates in its return to equilibrium. These intrinsically oscillatory systems exist in the Earth's environment, which imposes oscillations across temporal scales, from ocean waves, to diurnal cycles, to lunar cycles, to yearly cycles. Biological systems on Earth communicate with oscillations (e.g., acoustic disturbances like speech, bird calls, or dolphin clicks) and interact with the world to produce and control oscillations (e.g., breathing [1], movement [2], and digestion [3]). Importantly, the building blocks of animal intelligence—neurons—often exhibit oscillatory dynamics alone and in circuits with other neurons [4] to support functions like memory consolidation [5], executive control [6], and navigation [7].

Features of the brain's biological processes have motivated the development of artificial neural networks. Like biological neurons, artificial neurons respond when sufficient input drives stereotyped output, which subsequently excites or inhibits other neurons and impacts the connections between

neurons [8–10]. To process spatially structured input images, biological neurons in early visual cortex respond preferentially to localized regions and features of the visual field [11], which motivated the development of local receptive fields in convolutional neural networks [12]. To process temporally structured input signals, biological brains may exploit the temporal characteristics of individual neurons or neural population dynamics (e.g., rhythms [4,13], resonance [14,15], coherence [16,17], and cross-frequency coupling [18,19]). More recently, biological brain rhythms have motivated the development of artificial neural networks with imposed [20] or intrinsic [21] rhythmic structure. These recent works have shown that artificial neural networks with rhythmic features support multiplexing of complex inputs [20], and enhance learning speed, noise tolerance, and parameter efficiency [22].

Motivated by the utility of brain rhythms shown *in vivo* and *in silico* [4,16,23], we propose a modified recurrent neural network with rhythmic structure. Traditional recurrent neural networks are commonly written either as discrete-time difference equations in which the hidden state at time index t depends on the state at $t - 1$ [24] or as first-order continuous-time dynamical systems [25,26]; while such formulations can exhibit oscillatory trajectories, these formulations do not explicitly parametrize intrinsic, frequency-selective modes. To include rhythmic activity, neural networks with explicit bandpass filters have been proposed [27,28], and recent efforts have implemented networks as complex-valued equations [29] or second-order differential equations modeling networks of oscillators [21,22,30]. While these oscillatory networks benefit from physical interpretability (e.g., as damped, driven oscillators) and direct mathematical

*Contact author: mkramer@jhu.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](#) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

analysis via classical coupled-oscillator theory [31], these sophisticated formulations differ from traditional recurrent neural network expressions; a simplified framework may improve interpretability between, and insights from, different artificial neural network strategies.

To that end, we propose a resonant recurrent network (RRN) consisting of coupled, damped, driven oscillators that respond to time series input at different frequencies. Unlike previous approaches, we explicitly formulate the oscillator dynamics as a discrete-time recurrent neural network with two history-dependent terms. In this way, we link the physical oscillator dynamics to traditional recurrent neural network and statistical formulations, and we apply biological and stability constraints to limit the number of hyperparameters. To illustrate RRN performance, we analyze data from a well-established task: classification accuracy of handwritten digits. We show that, consistent with alternative coupled-oscillator network approaches [21,22,30], the RRN performs well in this classification task compared to existing approaches [a traditional recurrent neural network and long-short term memory (LSTM) network] with a comparable number of parameters. In addition, we show that choosing RRN oscillatory frequencies to mimic one proposed theory for the brain's rhythmic organization results in networks that outperform alternative frequency configurations. We find consistent results for classification of synthetic modulated or unmodulated spike train data. The proposed conceptual framework links techniques from statistics (spectral estimation theory) and physics (coupled-oscillator dynamics) to a traditional recurrent neural network formulation and provides a simple, biologically motivated computing strategy with few parameters and immediately interpretable dynamics.

II. RESULTS

A. A resonant recurrent network with stability constraints

The RRN consists of K coupled nodes. We begin by expressing the dynamics of the k th individual (i.e., uncoupled) node in the RRN as a damped, driven oscillator:

$$\ddot{x}^k = -2\beta_k \dot{x}^k - \omega_k^2 x^k + \xi^k(t) + U^k(t), \quad (1)$$

where x^k is the position, β_k is the damping constant, and ω_k is the natural frequency of the k th oscillator. We include an optional noise term $\xi^k(t)$ [27] to the k th oscillator evaluated at time t to create spontaneous activity consistent with biological neural networks [13]. Finally, we include an external input $U^k(t)$ to the k th oscillator at time t .

Replacing the derivatives with discrete approximations, Eq. (1) becomes

$$x_t^k = \alpha_1^k x_{t-1}^k + \alpha_2^k x_{t-2}^k + \epsilon_t^k + u_t^k, \quad (2)$$

where α_1^k and α_2^k are functions of the physical parameters β_k and ω_k , and ϵ_t^k (normally distributed with zero mean and standard deviation σ) and u_t^k are the noise and driving terms from Eq. (1), respectively, scaled by the time interval of the discrete approximation; see Appendix A of Ref. [32]. We note that, in this way, the damped, driven oscillator is equivalent to an autoregressive model of order 2 [i.e., an AR(2)] driven by external input u_t^k .

The complete RRN model consists of K coupled nodes with dynamics

$$x_t^k = \alpha_1^k x_{t-1}^k + \alpha_2^k x_{t-2}^k + \epsilon_t^k + u_t^k + \sum_{j \neq k} W^{kj} f(u_t^j x_{t-1}^j), \quad (3)$$

where we update Eq. (2) to include interactions between nodes. The matrix W indicates the connectivity, where W^{kj} represents the connectivity from node j to node k , and we set $W^{kk} = 0$ (i.e., no self-connections). We couple node dynamics using the product of the external input to node k (u_t^k) and the position of node j with nonlinearity $f(x) = \tanh(x)$. With this choice of coupling, the external input directly drives node k and modulates the impact of node j on node k (consistent with the phenomenon of gain modulation [33,34]).

The RRN model in Eq. (3) supports multiple interpretations. First, we motivated the RRN in Eq. (3) using a physical system: a coupled network of damped, driven harmonic oscillators [32]. Expressed as a discrete difference, the dynamical equation for the physical system becomes more like typical recurrent neural networks, but not quite. Typical recurrent neural networks (e.g., simple recurrent neural networks [24]) or networks with advanced gates (e.g., LSTM [35] or GRU [36] networks) include dependence only on the previous state (i.e., one history-dependent term). Here, motivated by the physical model of a damped oscillator, we extend this original formulation to include two history-dependent terms; the inclusion of the new term ($\alpha_2^k x_{t-2}^k$) allows intrinsic oscillatory dynamics at each node with features (i.e., natural frequency and damping) we control through parameters α_1^k and α_2^k .

We leverage these alternative formulations in different ways to support simplicity, enhance interpretability, and apply biological and mathematical constraints. We choose the formulation in Eq. (3)—rather than formulating the systems as coupled oscillators—for its simplicity and consistency with typical recurrent neural networks with one history-dependent term (i.e., with $\alpha_2^k = 0$). To establish damped oscillatory dynamics at center or resonant frequencies f_0 , we choose

$$\begin{aligned} \alpha_1^k &= 2r \cos\left(2\pi \frac{f_0}{F_s}\right), \\ \alpha_2^k &= -r^2, \end{aligned} \quad (4)$$

where r is the damping parameter and F_s is the sampling frequency of our discrete-time system; in what follows, we set $r = 0.9999$ for narrow-band oscillations and light damping, and we fix $F_s = 1000$ Hz. To enforce stationarity, we require the roots of the polynomial

$$1 - \alpha_1^k z^{-1} - \alpha_2^k z^{-2} = 0 \quad (5)$$

lie inside of the unit circle in the complex plane [37]. We additionally constrain $\alpha_1^k > 0$ to avoid sign-alternating dynamics and restrict frequencies to $f_0 < \frac{F_s}{2}$ to prevent aliasing. We set $f_0 \geq 2$ Hz to focus on rhythms in the traditional frequency bands studied *in vivo* [38,39]. In what follows, we choose the frequencies f_0 based on different proposals for the organization of mammalian brain rhythms *in vivo*.

As an initial illustration of the RRN dynamics, we simulate a network with 11 nodes, choosing frequencies spaced by factors of the golden ratio $f_0 \in \{2, 3.2, 5.2, 8.5, 13.7, 22.2, 35.9, 58.1, 93.9, 152.0, 245.9\}$ Hz, motivated




FIG. 1. The resonant recurrent network consists of damped, coupled oscillators. (a) Theoretical spectra computed for 11 nodes with center frequencies separated by factors of the golden ratio. (b) Example dynamics of this RRN driven by noise; colors correspond to the spectra in panel (a).

by observed distributions of brain rhythms *in vivo* [32,39,40]. We set the input to 0 [$u_t^k = 0$ in Eq. (3)] and remove all connections between nodes [$W = 0$ in Eq. (3)]. The dynamics of each node then consist of damped oscillations (i.e., fading memory) with narrow spectral peaks [example theoretical spectra [37] with these parameters shown in Fig. 1(a)]. The noise [ϵ_t^k in Eq. (3) with $\sigma = 1$] results in noisy, rhythmic fluctuations in the activity of each oscillator [Fig. 1(b)]. This example illustrates the interpretability of the RRN network dynamics as damped, driven oscillators.

In what follows, we train RRNs to classify input signals (i.e., both u_t^k and W are no longer restricted to 0) and show how performance depends on node frequencies. We simplify by setting the noise to zero [$\epsilon_t^k = 0$ for all k, t in Eq. (3)] and applying the same input to all nodes, so that u_t^k becomes u_t in Eq. (3). We note that, in doing so, we fix the input weight to 1 for all nodes; i.e., the same input signal drives all network elements. In the examples that follow, we apply only minimal preprocessing to each input signal (e.g., simple scaling). We compute per-node features as the time-averaged squared amplitude and pass these features through batch normalization and a linear readout to produce features for classification. Intuitively, we expect nodes will respond to the frequency content of the input signals and learn cross-frequency interactions to support classification. We show that RRNs outperform comparable recurrent neural networks in two classification tasks when the RRNs implement a biologically motivated choice of node frequencies.

B. The resonant recurrent network—with golden ratio spacing between frequencies—outperforms alternative configurations and comparable architectures in a common classification task

To characterize RRN performance, we analyzed the well-established problem of classifying the hand-written digits in the MNIST dataset [see Sec. IV A and examples in Fig. 2(a)].

We first selected three configurations of node frequencies, motivated by observations from *in vivo* brain dynamics (Table I). These configurations separate the brain's discrete rhythms by factors of Euler's number ($e \approx 2.718$) [41], two [42], or the golden ratio ($\phi \approx 1.618$) [32]. We trained and tested each configuration in the same way, with two hyperparameters (weight decay and label-smoothing coefficient) selected for each architecture to maximize accuracy on held-out validation data from the training set (see Secs. IV B and IV C). Training estimated the connectivity between nodes [W in Eq. (3)] and parameters for the linear classifier and batch normalization. We trained each architecture with ten realizations of the MNIST data, using the same randomized train/validation data splits for each architecture, and computed the classification accuracy of the test data (not used in training or validation). We found that the RRN with frequencies separated by factors of the golden ratio—the *golden RRN*—performed with higher accuracy than the other two biologically motivated configurations [Table I, $p < 1e^{-14}$, Fig. 2(b)]. Moreover, the golden RRN outperformed a network with the same number of parameters but linearly spaced frequencies ($p = 4e^{-14}$, Table I). Finally, compared to existing recurrent neural networks with a comparable number of parameters (an LSTM and standard Elman RNN), the golden RRN again performed with higher classification accuracy (LSTM: $p = 0.0098$; standard RNN: $p = 1.1e^{-11}$; see Table I). Confusion matrices demonstrated comparable or superior performance of the golden RRN compared to the other architectures [per-class accuracy of the golden RRN exceeded or matched all other architectures except for digits 0 (LSTM and standard RNN) and 8 (standard RNN); Fig. 2(c)]. We conclude that the RRN with frequencies separated by factors of the golden ratio outperformed other biologically motivated RRNs and existing recurrent neural network approaches with comparable numbers of parameters.




FIG. 2. The resonant recurrent network with golden ratio spacing classifies hand-written digits with highest accuracy. (a) Example MNIST images and corresponding scanline time series. (b) Average test accuracies for each network architecture; black vertical lines indicate ± 2 standard errors of the mean (SEM). (c) Average confusion matrices for the golden RRN, LSTM, and RRN with factor of 2 spacing.

TABLE I. Characteristics and performance of different network architectures in two classification tasks. RRNs consist of node frequencies f_0 spaced by factors of the golden ratio (“Golden”, $n = 11$ nodes), two ($n = 7$ nodes), or Euler’s number ($n = 5$ nodes), or linear spacing from 2 to 245.9 Hz ($n = 11$ nodes). A standard recurrent neural network (RNN, $n = 11$ nodes) and a long short-term memory (LSTM, $n = 6$ nodes) network were also compared. Characteristics include the number of estimable parameters and values for two hyperparameters (weight decay and label-smoothing coefficient). Performance includes accuracy classifying hand-written digits (MNIST) and spike train modulation (spike train).

	Node frequencies (Hz)	Parameters	Hyperparameters (decay, smoothing)	MNIST accuracy Mean (SEM)	Spike train accuracy Mean (SEM)
Golden ($n = 11$)	$f_0 \in \{2, 3.2, 5.2, 8.5, 13.7, 22.2, 35.9, 58.1, 93.9, 152.0, 245.9\}$	252	0, 0	0.759 (0.007)	0.635 (0.001)
Two ($n = 7$)	$f_0 \in \{2, 4, 8, 16, 32, 62, 128\}$	136	0.01, 0	0.583 (0.007)	0.622 (0.001)
Euler’s ($n = 5$)	$f_0 \in \{2, 5.4, 14.8, 40.2, 109.2\}$	90	0.01, 0	0.504 (0.004)	0.615 (0.001)
Linear ($n = 11$)	$f_0 \in \{2, 26.4, 50.8, 75.2, 99.6, 124.0, 148.3, 172.7, 197.1, 221.5, 245.9\}$	252	0, 0	0.592 (0.004)	0.604 (0.001)
RNN ($n = 11$)	—	296	0, 0.05	0.617 (0.018)	0.575 (0.002)
LSTM ($n = 6$)	—	298	0.001, 0	0.715 (0.019)	0.546 (0.002)




FIG. 3. RRN outputs depend on input digit and oscillator frequency. The average (solid curves) and 95% confidence intervals (shaded regions) of the golden RRN output for four digits (see legend) from the testing dataset. The RRN response is frequency and digit dependent. Distinct responses (e.g., digits 1 and 4) are easier to classify than overlapping responses (e.g., digits 2 and 3).

C. Examination of RRN dynamics provides insights into classification performance

Inspection of the RRN outputs provides interpretable insights into classification performance. To illustrate this claim, we investigated how the golden RRN responded to different hand-written digits. To do so, we selected a trained golden RRN (uniform randomly from the ten trained golden RRN architectures) and computed the response (the accumulated amplitude squared of each oscillator) to a randomly chosen digit from the test set (not used in training). For a fixed digit value (e.g., “1”), we repeated this process 100 times and determined the response mean and SEM across the 100 repetitions. Visual inspection of the example RRN output (the accumulated amplitude squared of each oscillator) revealed distinct responses to digits with high correct predictions (e.g., digits 1 and 4 in Fig. 3). Digit 1 maintains a relative spectral peak at ≈ 36 Hz, corresponding to the nearly periodic spectral trace of this digit [see example in Fig. 2(a)], with approximate period of 28 samples sampled at 1000 Hz (28 ms). Digits 2 and 3 have lower mean correct prediction ratios (0.66 and 0.69, respectively), and these two digits are more often mislabeled [see off-diagonal elements in the confusion matrix of Fig. 2(c)]. The response to digit 2 (orange curve in Fig. 3) overlaps with digit 3 (green curve in Fig. 3). In this way, the RRN produces interpretable features, which provide insight into classification performance.

D. Estimated recurrent connections indicate cross-frequency interactions

Having found the RRN with golden ratio frequency spacing (i.e., the “golden RRN”) performs best, we examined the resulting estimated connectivity between nodes. To do

so, we trained the golden RRN with 100 realizations of the MNIST data, each using a different random seed. While the resulting connectivity matrix (W) varied across realizations [examples in Fig. 4(a)], visual inspection of the mean connectivity matrix suggests more positive connection strengths from lower to higher frequency nodes [2–8.5 to 22.2–93.9 Hz, solid black square in Fig. 4(b)] compared to connection strengths from higher to lower frequency nodes [dashed black square in Fig. 4(b)]. Computing the average connectivity value for low-to-high frequency connections and high-to-low frequency connections, we find different weight distributions [Fig. 4(c)]. The mean low-to-high frequency connection strength exceeds the mean high-to-low frequency connection strength (t -test, $t = 12.1$, $p \approx 1e^{-25}$).

E. Golden ratio spacing in the resonant recurrent network classifies different spike train patterns

To illustrate RRN performance in a task more closely related to neuronal time series data, we considered the challenge of separating two different spike train patterns. The patterns consisted of brief (100 ms, sampling rate 1000 Hz) spike train data simulated in two conditions: (1) random spiking with an average firing rate of 100 Hz or (2) random spiking with an average firing rate of 100 Hz modulated by a sinusoid with frequency chosen uniformly between 10 and 50 Hz; see Sec. IV D. Visual inspection of example traces demonstrates the difficulty of distinguishing between the constant and modulated scenarios [Fig. 5(a)]. To classify these spike train data, we trained and tested each network architecture using the fixed hyperparameters (weight decay and label-smoothing coefficient) selected for the MNIST data (see Secs. IV B and IV C). Training estimated the connectivity between nodes [W in Eq. (3)] and parameters for the linear classifier and batch normalization. To match comparisons across architectures, we trained ($N = 1000$ examples), validated ($N = 200$ examples), and tested ($N = 200$ examples) each architecture using the same simulated spike train data. We repeated the entire training and testing process for 1000 realizations of the simulated spike train data. Consistent with the results for the MNIST dataset, we found that the RRN with frequencies separated by factors of the golden ratio—the golden RRN—performed with higher accuracy than the other architectures [Table I, $p < 5e^{-11}$; Fig. 5(b)]. Confusion matrices demonstrated comparable or superior performance of the golden RRN compared to the other architectures [Fig. 5(c)]. Inspection of example RRN output (the accumulated amplitude squared of each oscillator) revealed an increased response of oscillators in the 10–50 Hz frequency range for the modulated data, as expected for these simulated spike trains [Fig. 5(d)]. Finally, the trained connectivity matrices [W , Fig. 5(e)] demonstrated increased low-to-high frequency connection weights compared to high-to-low frequency connection weights [Fig. 5(f); t -test, $t = -4.9$, $p = 1.2e^{-6}$], again consistent with the MNIST classification results. We conclude that the golden RRN performed consistently for the spike train data and the MNIST data. In both cases, the golden RRN outperformed the other architectures, produced interpretable features, and learned recurrent




FIG. 4. Network connectivities learned during training exhibit cross-frequency interactions. (a) Example connectivity matrices W estimated during different training runs. (b) Average connectivity matrix over all training runs. Black squares indicate low-to-high frequency connectivity (solid) and high-to-low frequency connectivity (dashed). (c) Distributions of average cross-frequency coupling strengths from low-to-high (orange) and high-to-low (blue) frequencies.

connectivity with similar patterns of cross-frequency coupling.

III. DISCUSSION

We propose a resonant recurrent network (RRN) consisting of damped oscillatory elements, which we express as a traditional recurrent neural network with an additional history-dependent term. We show that the RRN outperforms standard recurrent neural network approaches to classify hand-written digits and simulated spike train data. Compared to existing recurrent neural network approaches, nodes in the RRN intrinsically oscillate and the system is simple to express [see Eq. (3)] and parametrize; the main parameter choice is the center frequency of each oscillator. Moreover, the structure and output of the RRN are directly interpretable; the RRN consists of damped, coupled oscillators and the outputs correspond to the summed amplitudes squared. Well-established constraints enforce stable node dynamics with fading (or short-term) memory and the input signals require no elaborate preprocessing. Finally, the RRN benefits from biological knowledge—The node frequencies can be chosen to follow the observed distributions of brain rhythms *in vivo*.

We propose two interpretations of the RRN: either as a network of discrete-time recurrent neural network elements with

two history-dependent terms or as a network of continuous-time, coupled, damped, driven oscillators. Each interpretation offers insights into network performance and design. Typical recurrent neural networks implement first-order discrete dynamical systems [43–47] consistent with autoregressive models of order 1 [AR(1)]. The spectrum of an AR(1) consists of a low-frequency peak and each node acts to low-pass filter input signals [37]. In this way, the node dynamics in typical recurrent neural networks are limited and cannot resonate with narrow-band higher-frequency oscillations. The two history-dependent terms in the RRN extend these typical systems to an autoregressive model of order 2 [AR(2)], thus extending the capabilities of each node to narrow-band filtering. Alternatively, conceptualizing the RRN as a network of damped oscillators provides physical intuition and offers opportunities to connect the neural network dynamics to existing well-developed theories of coupled oscillators [48–51]. Application of these existing theories may facilitate a deeper understanding of the RRN dynamics, providing insights difficult to obtain in typical deep neural network formulations.

We showed that frequencies separated by factors of the golden ratio outperformed alternative proposed frequency spacing scenarios (e.g., logarithmic [41] or factor of 2 [42]) in the RRN. These results are consistent with existing theoretical [32,52–54] and experimental [39,55] works, which




FIG. 5. The RRN with golden frequency spacing classifies spike train patterns. (a) Example constant and modulated spike train patterns. Each row indicates one spike train realization (duration 100 ms). Fifty examples (rows) are shown. (b) Average test accuracy for each network architecture; black vertical lines indicate ± 2 SEM. (c) Average confusion matrices ($\pm$ SEM) for each network architecture. (d) Average (solid curves) and 95% confidence intervals (shaded regions) of the golden RRN output for each spike train condition (constant: blue; modulated: orange). (e) Average connectivity matrix over all training runs. Black squares indicate low-to-high frequency connectivity (solid) and high-to-low frequency connectivity (dashed). (f) Distributions of average cross-frequency coupling strengths from low-to-high (orange) and high-to-low (blue) frequencies.

suggest that golden ratio spacing supports both independent communication channels and coupled cross-frequency interactions. Training the node-to-node connectivity matrices using a standard artificial neural network approach (backpropagation through time), we find evidence of cross-frequency interactions, with low-to-high frequency oscillators more positively coupled than high-to-low frequency oscillators. These cross-frequency interactions may be especially important for characterizing the sharp, abrupt changes characteristic of

the scanline MNIST or simulated spike train time series data. How these interactions between simulated damped oscillators relate to the biological mechanisms of interacting neurons or neural populations remains a topic for future investigation.

The proposed RRN differs from existing oscillatory recurrent neural networks [21,22] in the following ways. First, the RRN formulation provides a direct link from the physical system—damped oscillators—to a simple extension of

existing recurrent neural network models. The fundamental insight is to augment the traditional recurrent neural network model ($x_t = \alpha x_{t-1}$) to include a second-order term ($x_t = \alpha_1 x_{t-1} + \alpha_2 x_{t-2}$). Second, we choose oscillator frequencies in the RRN to match the spacing between the brain's frequency bands observed *in vivo* [32,41,42]. Doing so, we find that golden ratio spacing outperforms alternative scenarios and standard recurrent neural networks (LSTM and Elman RNN) with a comparable number of parameters. Third, the number of estimated parameters in the RRN is small; determining the recurrent connections between N nodes requires estimation of $N^2 - N$ parameters (no self-connections), with $N \leq 11$. Fourth, the RRN produces interpretable outputs, the accumulated amplitude squared (or response) of each oscillator, and investigation of these responses provides insight into classification performance.

Here, we used training data to estimate connectivity between RRN nodes. An alternative approach is to constrain or impose network connectivity motivated by physical or biological insight. For example, replacing the inferred network architecture with a finite difference approximation of the Laplacian, the activation function with the node activity, and limiting the frequency differences between nodes would simulate diffusive coupling between approximately homogeneous damped, driven oscillators. In this way, the *in silico* RRN connectivity and dynamics may be altered to approximate existing physical reservoir computing strategies (e.g., liquid media [56,57], memristors [58], nanoscale magnetic oscillators [59], or human brain organoids [60]) and, for example, support waves of neural activity [61], which may implement computational principles similar to the transformer networks of large language models [62]. Updating the network architecture to use *in vivo* brain connectivity [63,64] would incorporate additional biological features, and the nonlocal, distributed interactions may better support computations than the local interactions of simpler physical systems [56].

We note that an appropriate choice of network architecture could combine two nodes, each with one history-dependent term (i.e., $t - 1$), to implement second-order dynamics. This connectivity could be imposed or, perhaps, result from training. However, in this case, interpretation of the node activity becomes less direct. For example, we might interpret some nodes as representing the position and other nodes the velocity of a dynamical system. While mathematically equivalent, this first-order system is conceptually more complex. Here, we instead impose second-order dynamics at each node, or equivalently dynamics that depend on two history-dependent terms at $t - 1$ and $t - 2$. An advantage of this framework is a direct physical interpretation and connection to a clear physical mechanism (i.e., a damped oscillator) or to narrow-band filtering at each node.

Traditional artificial neural networks draw inspiration from the brain's neuronal dynamics. Here, we continue this approach by mimicking an additional, well-established feature of neuronal dynamics: rhythms [4]. Most artificial recurrent neural networks implement first-order dynamics and therefore do not intrinsically oscillate. The brain, however, oscillates in both health and disease [23,65,66]. Why biological neurons and brain circuits oscillate remains unknown. Oscillations have been proposed to support brain processes, including

TABLE II. For the MNIST dataset, the number of examples for each digit used in training and testing.

Digit	Training count	Testing count
0	5923	980
1	6742	1135
2	5958	1032
3	6131	1010
4	5842	982
5	5421	892
6	5918	958
7	6265	1028
8	5851	974
9	5949	1009
Total	60 000	10 000

temporal filtering via resonance [14,15], communication through coherence [16,17], and cross-frequency interactions [18,19]. More generally, biological brains exist in an oscillatory world and therefore may utilize oscillatory components to exploit these oscillatory inputs. Because many arbitrary signals can be represented as a summation of sinusoids, developing a biological system that decomposes signals as sinusoids supports a rich feature space. Alternatively, biological neural networks may have no other choice as the only available components (biological neurons) necessarily oscillate [67,68]. In that case, artificial brains, which are not restricted to an oscillatory environment or oscillatory components, may have an advantage over biological brains. However, recent work directly comparing oscillatory and nonoscillatory neural networks suggests that oscillations may support performance [22]. Beyond oscillations, our biological neural networks may possess intrinsic advantages not immediately accessible to artificial neural networks. If so, protecting these advantages may serve strategic outcomes if biological and artificial brains must ultimately compete [69].

IV. METHODS

A. MNIST dataset

The MNIST dataset [70] consisted of ten hand-written digits from 0 to 9. The training set contained 60 000 images and the testing set 10 000 images (Table II). Each digit consisted of a 28×28 matrix of intensity values, which we transformed (scanline order) to a (1×784) time series with assumed sampling frequency of 1000 Hz and rescaled to values between 0 and 1 [see examples in Fig. 2(a)].

B. Model training

All models were trained using a common PyTorch [71] pipeline with supervised minibatch optimization. For the MNIST data, for a chosen random seed, a fixed train/validation split was generated by shuffling the 60 000 training examples and holding out 5000 for validation; the standard 10 000-example MNIST test set was used only for final evaluation. For the synthetic spike train data, for a chosen random seed, synthetic data were generated with constant or modulated firing rate, with 1000 training examples

(500 per class), 200 validation examples (100 per class), and 200 testing examples (100 per class). Models were trained for 30 epochs (batch size 128) with AdamW, cross-entropy (optional label smoothing), global gradient clipping (norm 1.0), and OneCycle LR scheduling (cosine annealing, stepped each minibatch). After each epoch, validation accuracy was computed; the best-validation checkpoint was saved, reloaded after training, and evaluated on test accuracy. For the RRN configurations, the diagonal of the recurrent weight matrix W was constrained to zero (preventing self-connections) by applying a postupdate projection after each optimizer step.

C. Hyperparameter selection

We selected hyperparameters via a grid search performed separately for each architecture variant (each RRN type, LSTM, and the recurrent neural network). For each candidate setting, models were trained for 30 epochs with the same training loop and objective (multiclass cross-entropy with optional label smoothing), varying only the optimizer/loss hyperparameters: learning rate $\in 3 \times 10^{-4}, \{10^{-3}, 3 \times 10^{-3}\}$, weight decay $\in 0, 10^{-3}, 10^{-2}$, and label-smoothing coefficient $\in 0, 0.05$. Each grid point was evaluated across three independent random seeds; for each seed, the model was reinitialized, and 5000 MNIST examples were held out for validation. During training, validation accuracy was computed for each epoch and the checkpoint achieving the maximum validation accuracy over the full training trajectory was retained. After completing the grid, hyperparameters were selected per architecture by maximizing the mean best-validation accuracy across the three seeds, requiring that all three seeds were present for a candidate configuration. The

test set was not used for hyperparameter selection; the chosen hyperparameters were subsequently fixed for subsequent evaluations. We note that, for all network architectures, a learning rate of 3×10^{-3} optimized performance.

D. Simulated spike train data

Spike train data were simulated as realizations of a Poisson process with either constant intensity of 100 Hz or modulated intensity. For the modulated intensity,

$$\lambda_t = f_c + f_m \sin(2\pi f_m t),$$

where $f_c = 100$ Hz and f_m was selected (uniformly) between 10 and 50 Hz for each spike train realization.

E. Statistical tests

To test for differences in accuracy between models, we estimated linear models with output test accuracy and categorical predictors corresponding to each model type (reference category set to the golden RRN) and a constant predictor.

F. Code availability

All methods required to reproduce the results in this paper are available in Ref. [72] for reuse and further development.

ACKNOWLEDGMENTS

This work was supported in part by NIH R01NS119483 and NIH R01NS110669.

DATA AVAILABILITY

The data that support the findings of this article are openly available in Ref. [72].

-
- [1] A. B. L. Tort, J. Brankač, and A. Draguhn, Respiration-entrained brain rhythms are global but often overlooked, *Trends Neurosci.* **41**, 186 (2018).
 - [2] E. Marder and D. Bucher, Central pattern generators and the control of rhythmic movements, *Curr. Biol.* **11**, R986 (2001).
 - [3] E. Marder and D. Bucher, Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs, *Annu. Rev. Physiol.* **69**, 291 (2007).
 - [4] G. Buzsáki, *Rhythms of the Brain* (Oxford University Press, New York, NY, USA, 2011).
 - [5] B. P. Staresina, Coupled sleep rhythms for memory consolidation, *Trends Cogn. Sci.* **28**, 339 (2024).
 - [6] M. Lundqvist, E. K. Miller, J. Nordmark, J. Liljefors, and P. Herman, Beta: Bursts of cognition, *Trends Cogn. Sci.* **28**, 662 (2024).
 - [7] G. Buzsáki and E. I. Moser, Memory, navigation and theta rhythm in the hippocampal-entorhinal system, *Nat. Neurosci.* **16**, 130 (2013).
 - [8] W. McCulloch and W. Pitts, A logical calculus of the ideas immanent in nervous activity, *Bull. Math. Biol.* **5**, 115 (1943).
 - [9] J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities, *Proc. Natl. Acad. Sci. USA* **79**, 2554 (1982).
 - [10] F. Rosenblatt, The perceptron: A probabilistic model for information storage and organization in the brain, *Psychol. Rev.* **65**, 386 (1958).
 - [11] D. H. Hubel and T. N. Wiesel, Receptive fields of single neurons in the cat's striate cortex, *J. Physiol.* **148**, 574 (1959).
 - [12] Y. LeCun and Y. Bengio, Convolutional networks for images, speech, and time series, in *The Handbook of Brain Theory and Neural Networks*, 2nd ed., edited by M. A. Arbib (MIT Press, Cambridge, USA, 1998), p. 255.
 - [13] W. Singer, The cerebral cortex: A delay-coupled recurrent oscillator network?, in *Reservoir Computing: Theory, Physical Implementations, and Applications*, edited by K. Nakajima and I. Fischer, (Springer Nature Singapore, Singapore, 2021), pp. 3–28.
 - [14] G. Hahn, A. Ponce-Alvarez, G. Deco, A. Aertsen, and A. Kumar, Portraits of communication in neuronal networks, *Nat. Rev. Neurosci.* **20**, 117 (2019).
 - [15] R. F. Helfrich, A. Breska, and R. T. Knight, Neural entrainment and network resonance in support of top-down guided attention, *Curr. Opin. Psychol.* **29**, 82 (2019).
 - [16] P. Fries, A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence, *Trends Cogn. Sci.* **9**, 474 (2005).

- [17] P. Fries, Rhythms for cognition: Communication through coherence, *Neuron* **88**, 220 (2015).
- [18] R. T. Canolty and R. T. Knight, The functional role of cross-frequency coupling, *Trends Cogn. Sci.* **14**, 506 (2010).
- [19] A. Hyafil, A.-L. Giraud, L. Fontolan, and B. Gutkin, Neural cross-frequency coupling: Connecting architectures, mechanisms, and functions, *Trends Neurosci.* **38**, 725 (2015).
- [20] K. Duecker, M. Idiart, M. van Gerven, and O. Jensen, Oscillations in an artificial neural network convert competing inputs into a temporal code, *PLoS Comput. Biol.* **20**, e1012429 (2024).
- [21] T. K. Rusch and S. Mishra, Coupled Oscillatory Recurrent Neural Network (coRNN): An accurate and (gradient) stable architecture for learning long time dependencies, *International Conference on Learning Representations* (2021).
- [22] F. Effenberger, P. Carvalho, I. Dubinin, and W. Singer, The functional role of oscillatory dynamics in neocortical circuits: A computational perspective, *Proc. Natl. Acad. Sci. USA* **122**, e2412830122 (2025).
- [23] X.-J. Wang, Neurophysiological and computational principles of cortical rhythms in cognition, *Physiol. Rev.* **90**, 1195 (2010).
- [24] J. L. Elman, Finding structure in time, *Cogn. Sci.* **14**, 179 (1990).
- [25] K. Funahashi and Y. Nakamura, Approximation of dynamical systems by continuous time recurrent neural networks, *Neural Networks* **6**, 801 (1993).
- [26] K. Rajan, L. F. Abbott, and H. Sompolinsky, Stimulus-dependent suppression of chaos in recurrent neural networks, *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* **82**, 011903 (2010).
- [27] G. Holzman and H. Hauser, Echo state networks with filter neurons and a delay&sum readout, *Neural Networks* **23**, 244 (2010).
- [28] F. Wyffels, B. Schrauwen, D. Verstraeten, and D. Stroobandt, Band-pass reservoir computing, in *2008 IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, Hong Kong, China (IEEE, Piscataway, NJ, USA, 2008), pp. 3204–3209.
- [29] L. H. B. Liboni *et al.*, Image segmentation with traveling waves in an exactly solvable recurrent neural network, *Proc. Natl. Acad. Sci. USA* **122**, e2321319121 (2025).
- [30] A. Ceni *et al.*, Random oscillators network for time series processing, in *Proceedings of the 27th International Conference on Artificial Intelligence and Statistics* (PMLR, 2024), pp. 4807–4815.
- [31] S. H. Strogatz, From Kuramoto to Crawford: Exploring the onset of synchronization in populations of coupled oscillators, *Physica D* **143**, 1 (2000).
- [32] M. A. Kramer, Golden rhythms as a theoretical framework for cross-frequency organization, *Neuron Behav. Data Anal. Theory* **1** (2022).
- [33] T. Akam and D. M. Kullmann, Oscillatory multiplexing of population codes for selective communication in the mammalian brain, *Nat. Rev. Neurosci.* **15**, 111 (2014).
- [34] K. A. Ferguson and J. A. Cardin, Mechanisms underlying gain modulation in the cortex, *Nat. Rev. Neurosci.* **21**, 80 (2020).
- [35] S. Hochreiter and J. Schmidhuber, Long short-term memory, *Neural Comput.* **9**, 1735 (1997).
- [36] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, Learning phrase representations using RNN encoder-decoder for statistical machine translation, *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, edited by A. Moschitti, B. Pang, and W. Daelemans (Association for Computational Linguistics, Doha, Qatar, 2014).
- [37] D. Percival and A. Walden, *Spectral analysis for physical applications* (Cambridge University Press, Cambridge, UK, 1993).
- [38] G. Buzsáki and A. Draguhn, Neuronal oscillations in cortical networks, *Science* **304**, 1926 (2004).
- [39] A. K. Roopun *et al.*, Temporal interactions between cortical rhythms, *Front. Neurosci.* **2**, 145 (2008).
- [40] M. A. Kramer, A. K. Roopun, L. M. Carracedo, R. D. Traub, M. A. Whittington, and N. J. Kopell, Rhythm generation through period concatenation in rat somatosensory cortex, *PLoS Comput. Biol.* **4**, e1000169 (2008).
- [41] M. Penttonen and G. Buzsáki, Natural logarithmic relationship between brain oscillators, *Thalamus Relat. Syst.* **2**, 145 (2003).
- [42] W. Klimesch, An algorithm for the EEG frequency architecture of consciousness and brain body coupling, *Front. Hum. Neurosci.* **7**, 1 (2013).
- [43] H. Jaeger, The ‘echo state’ approach to analysing and training recurrent neural networks—With an Erratum note, *GMD Rep.* **148**, 1 (2001).
- [44] W. Maass, T. Natschläger, and H. Markram, Real-time computing without stable states: A new framework for neural computation based on perturbations, *Neural Comput.* **14**, 2531 (2002).
- [45] H. Jaeger, M. Lukoševičius, D. Popovici, and U. Siewert, Optimization and applications of echo state networks with leaky-integrator neurons, *Neural Networks* **20**, 335 (2007).
- [46] G. Tanaka, T. Matsumori, H. Yoshida, and K. Aihara, Reservoir computing with diverse timescales for prediction of multiscale dynamics, *Phys. Rev. Res.* **4**, L032014 (2022).
- [47] C. Gallicchio, A. Micheli, and L. Pedrelli, Design of deep echo state networks, *Neural Networks* **108**, 33 (2018).
- [48] E. M. Izhikevich and F. C. Hoppensteadt, *Weakly Connected Neural Networks*, 1st ed. (Springer, New York, NY, USA, 1997).
- [49] T. Aonishi, K. Kurata, and M. Okada, Statistical mechanics of an oscillator associative memory with scattered natural frequencies, *Phys. Rev. Lett.* **82**, 2800 (1999).
- [50] A. T. Winfree, Biological rhythms and the behavior of populations of coupled oscillators, *J. Theor. Biol.* **16**, 15 (1967).
- [51] Y. Kuramoto, Collective synchronization of pulse-coupled oscillators and excitable units, *Physica D* **50**, 15 (1991).
- [52] E. M. Izhikevich, Weakly connected quasi-periodic oscillators, fm interactions, and multiplexing in the brain, *SIAM J. Appl. Math.* **59**, 2193 (1999).
- [53] E. M. Izhikevich, Weakly-pulse-coupled oscillators, FM interactions, synchronization, and oscillatory associative memory, *IEEE Trans. Neural Networks* **10**, 508 (1999).
- [54] F. C. Hoppensteadt and E. M. Izhikevich, Thalamo-cortical interactions modeled by weakly connected oscillators: Could the brain use FM radio principles? *Biosystems* **48**, 85 (1998).
- [55] A. K. Roopun *et al.*, Period concatenation underlies interactions between gamma and beta rhythms in neocortex, *Front. Cell. Neurosci.* **2**, (2008).
- [56] T. Natschläger, W. Maass, and H. Markram, The ‘liquid computer’: A novel strategy for real-time computing on time series, *Telematik* **8**, 39 (2002).

- [57] C. Fernando and S. Sojakka, Pattern recognition in a bucket, in *Advances in Artificial Life*, edited by W. Banzhaf, J. Ziegler, T. Christaller, P. Dittrich, and J. T. Kim (Springer, Berlin, 2003), pp. 588–597.
- [58] Y. Zhong *et al.*, A memristor-based analogue reservoir computing system for real-time and power-efficient signal processing, *Nat. Electron.* **5**, 672 (2022).
- [59] J. Torrejon *et al.*, Neuromorphic computing with nanoscale spintronic oscillators, *Nature (London)* **547**, 428 (2017).
- [60] H. Cai *et al.*, Brain organoid reservoir computing for artificial intelligence, *Nat. Electron.* **6**, 1032 (2023).
- [61] T. A. Keller, L. Muller, T. Sejnowski, and M. Welling, Traveling waves encode the recent past and enhance sequence learning, *The Twelfth International Conference on Learning Representations* (2024).
- [62] L. Muller, P. S. Churchland, and T. J. Sejnowski, Transformers and cortical waves: Encoders for pulling in context across time, *Trends Neurosci.* **47**, 788 (2024).
- [63] F. Damicelli, C. C. Hilgetag, and A. Goulas, Brain connectivity meets reservoir computing, *PLoS Comput. Biol.* **18**, e1010639 (2022).
- [64] L. E. Suárez *et al.*, Connectome-based reservoir computing with the conn2res toolbox, *Nat. Commun.* **15**, 656 (2024).
- [65] G. Buzsáki and M. Vöröslakos, Brain rhythms have come of age, *Neuron* **111**, 922 (2023).
- [66] N. J. Kopell, H. J. Gritton, M. A. Whittington, and M. A. Kramer, Beyond the connectome: The dynamo, *Neuron* **83**, 1319 (2014).
- [67] K. M. Stiefel and G. B. Ermentrout, Neurons as oscillators, *J. Neurophysiol.* **116**, 2950 (2016).
- [68] M. Whittington, R. Traub, N. Kopell, B. Ermentrout, and E. Buhl, Inhibition-based rhythms: Experimental and mathematical observations on network dynamics, *Int. J. Psychophysiol.* **38**, 315 (2000).
- [69] D. Kototajlo, S. Alexander, T. Larsen, E. Lifland, and R. Dean, AI 2027, AI Futures Project, 2025.
- [70] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, Gradient-based learning applied to document recognition, *Proc. IEEE* **86**, 2278 (1998).
- [71] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Köpf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, PyTorch: An imperative style, high-performance deep learning library, *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), Vol. 32.
- [72] M. Kramer, Mark-Kramer/Resonant-Recurrent-Network, Jupyter Notebook, <https://github.com/Mark-Kramer/Resonant-Recurrent-Network>, accessed 11 May 2026.