

From Neurons to Computation: Biological Reservoir Computing for Pattern Recognition

Ludovico Iannello^{1*}[0009-0003-1392-7307], Luca Ciampi^{1*}[0000-0002-6985-0439], Gabriele Lagani^{1*}[0000-0003-4739-5778], Fabrizio Tonelli²[0000-0003-4167-555X], Eleonora Crocco², Lucio Maria Calcagnile³[0000-0003-3572-6154], Angelo Di Garbo³, Federico Cremisi²[0000-0003-4925-2703], and Giuseppe Amato¹[0000-0003-0171-4315]

¹ Institute of information science and technologies (CNR-ISTI), Pisa, Italy

² Biology Laboratory of Scuola Normale Superiore (Bio@SNS), Pisa, Italy

³ Institute of biophysics (CNR-IBF), Pisa, Italy

Abstract. In this paper, we introduce a paradigm for reservoir computing (RC) that leverages a pool of cultured biological neurons as the reservoir substrate, creating a *biological reservoir computing* (BRC). This system operates similarly to an echo state network (ESN), with the key distinction that the neural activity is generated by a network of cultured neurons, rather than being modeled by traditional artificial computational units. The neuronal activity is recorded using a multi-electrode array (MEA), which enables high-throughput recording of neural signals. In our approach, inputs are introduced into the network through a subset of the MEA electrodes, while the remaining electrodes capture the resulting neural activity. This generates a nonlinear mapping of the input data to a high-dimensional biological feature space, where distinguishing between data becomes more efficient and straightforward, allowing a simple linear classifier to perform pattern recognition tasks effectively. To evaluate the performance of our proposed system, we present an experimental study that includes various input patterns, such as positional codes, bars with different orientations, and a digit recognition task. The results demonstrate the feasibility of using biological neural networks to perform tasks traditionally handled by artificial neural networks, paving the way for further exploration of biologically-inspired computing systems, with potential applications in neuromorphic engineering and bio-hybrid computing.

Keywords: Reservoir Computing · Bio-Inspired Neural Networks · Multi-Electrode Array · Neural Computing.

1 Introduction

Reservoir Computing (RC) [40] is a powerful machine learning paradigm that leverages a highly nonlinear dynamical system with numerous degrees of freedom to project input data into a high-dimensional latent space. In this space, data samples are often more easily linearly separable, which facilitates efficient downstream tasks such as classification or regression. Several models have been proposed to implement the reservoir, among which the Liquid State Machine (LSM) [41] and the Echo State Network (ESN) [25, 17] are the most prominent. The LSM employs a pool of nonlinear spiking neural units to generate a dynamic representation of input signals, while the

* Corresponding authors, they contributed equally to this work: ludovico.iannello@isti.cnr.it, luca.ciampi@isti.cnr.it, gabriele.lagani@isti.cnr.it

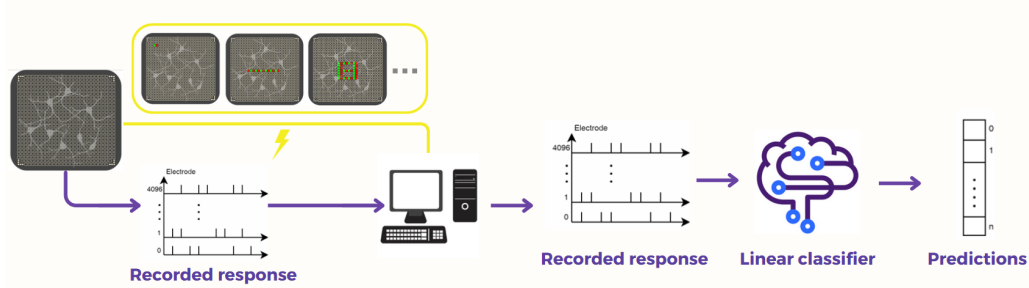


Fig. 1: **Overview of our biological reservoir computing (BRC) approach.** A multi-electrode array (MEA) serves as an interface to both stimulate and record activity from a cultured biological neural network. Each input is mapped to a specific subset of MEA electrodes, which deliver the corresponding stimulus to the network. The resulting spiking activity is recorded by a different set of electrodes and assembled into a vector that represents the input in a high-dimensional latent space. The complex dynamics of the BRC system ensure that this mapping is highly nonlinear. Finally, a linear classifier is trained to recognize the category of the original input based on its latent representation.

ESN uses continuous-valued neurons and does not aim to replicate biologically realistic spiking behavior. These models have enabled successful applications in various domains, including time-series prediction [3] and speech recognition [57].

In this paper, we introduce an extension to the RC paradigm: the use of a cultured network of real biological neurons as the reservoir substrate. This approach, which we refer to as the *biological reservoir computing* (BRC), employs a *biological* dynamical system to map input samples into a high-dimensional space of biological features. Unlike traditional RC systems, where the reservoir is typically composed of artificial neurons or units, our biological system leverages real neurons grown *in vitro*, providing a natural bridge between computational and biological systems. This paradigm has the potential to improve energy efficiency – a significant limitation of current deep learning technologies [2, 26] – while also advancing our understanding of biological neuron behavior. Indeed, neuroscientific observations suggest that transient dynamics in the brain carry a significant portion of the input information, indicating that some form of RC may play a role in high-level biological cognition [14]. In more detail, we derive our neurons from stem cells using well-established neurobiological techniques [7, 8]. These cultured networks are then interfaced with a high-density multi-electrode array (MEA) device [4], which contains 4096 electrodes for both electrical stimulation and neural response recording (see also Fig. 1 for a visual overview of our approach).

To evaluate the ability of our BRC system to generate discriminative feature representations across varying input complexities, we conduct an experimental study spanning tasks of several difficulties. Specifically, we feed features derived from biological culture responses into a single-layer perceptron and assess classification performance of the combined system (i.e., our BRC system followed by the perceptron). We begin with simple pointwise stimuli as basic input patterns, then progress to more complex geometric stimuli such as bars with varying orientations. Finally, we evaluate digit recognition to test the utility of the BRC system in real-world classification scenarios. Results are promising, indicating that biologically cultured networks can function as effective

substrates for reservoir computing, producing high-dimensional features that improve pattern recognition accuracy.

In summary, the contributions of this paper are as follows:

- We introduce a reservoir computing paradigm that leverages biologically cultured networks as substrates for generating high-dimensional feature representations, which we term biological reservoir computing (BRC). Our approach interfaces with these cultured networks using a high-density multi-electrode array (MEA) device, enabling both electrical stimulation and the recording of neural responses.
- We define and calibrate a stimulation protocol for interacting with neurons via the MEA device, ensuring robust and consistent neuronal responses across different stimulation parameters.
- We present an experimental investigation to validate the proposed BRC system in several pattern recognition tasks, assessing the efficacy of the biological feature representations in distinguishing various categories of input patterns.

The remainder of this manuscript is structured as follows: Sec. 2 provides an overview of related work in the field of reservoir computing and biological computing; Sec. 3 details the proposed methodology, including the neurobiological techniques and MEA interfacing; Sec. 4 discusses the experimental setup, results, and analysis; and finally, Sec. 5 presents our conclusions and outlines future directions.

2 Related Work

A central theme in deep learning (DL) research is the biological plausibility of current methods, with several studies questioning whether existing architectures reflect the complexity and functionality of biological neural systems [43, 20, 46, 55]. In response, there has been a growing interest in exploring biologically inspired alternatives to push the boundaries of machine learning and cognitive computation [35, 34]. A wide range of models has been proposed to bridge this gap, either by adopting biologically realistic neural computations [12, 16, 38, 54] or by refining synaptic dynamics and learning mechanisms [30, 23, 32, 33, 31, 37, 27, 36, 9, 10].

Reservoir Computing (RC) [40, 50] represents one of the bio-inspired paradigms that has gained considerable attention for its potential to model complex neural dynamics. Specifically, Echo State Networks (ESNs) [17, 25] and Liquid State Machines (LSMs) [41, 58] are two prominent models within the RC framework. While LSMs leverage biologically inspired spiking neuron models [18, 1] to implement highly nonlinear neural dynamics, ESNs rely on continuous-valued neurons and random recurrent connectivity to generate rich internal dynamics. In both cases, the reservoir transforms input signals into a high-dimensional feature space, facilitating pattern separation for downstream tasks. Although the connectivity is typically random, specific algorithms are employed to ensure that the network dynamics remain stable and computationally effective [17, 49]. Several extensions and related models have been proposed to enhance the biological plausibility of reservoir computing systems. For instance, some variants of ESNs [56] incorporate biologically inspired mechanisms such as spike-timing-dependent plasticity (STDP) [52, 18], which adjusts synaptic strengths based on the timing of pre- and post-synaptic activity. Another biologically motivated model, the Self-Organizing Recurrent Network (SORN) [39], combines STDP with homeostatic plasticity [6] to maintain network stability and support unsupervised learning in recurrent architectures. Additionally, a gating mechanism proposed in [49] enhances ESN performance by improving the recall of

long-term dependencies, which are often degraded in highly nonlinear recurrent systems. A related technique, the Extreme Learning Machine (ELM) [22, 21], employs a large pool of feedforward nonlinear neurons with randomly initialized weights to project inputs into a high-dimensional feature space. Although ELMs lack recurrent connections and differ architecturally from RC models, they share the core principle of using random nonlinear mappings to facilitate efficient learning. These biologically inspired and computationally efficient models have demonstrated success in various applications, including speech processing [58] and continual learning [11], among others.

In contrast to previous works, this contribution enhances the biological plausibility of RC approaches by leveraging a real biological neural network as the computational reservoir. Specifically, we introduce the concept of biological reservoir computing (BRC), in which cultured biological neurons serve as the reservoir substrate. Prior studies have explored interfacing with biological neurons via microelectrode array (MEA) devices [51, 47, 15, 24, 19, 45, 28], but our work specifically investigates the use of cultured neuronal networks as a potential RC substrate. A few previous efforts have also employed biological neurons in RC frameworks; however, they either focused on input separation through distinct stimuli [13] or relied on alternative interfacing methods such as optogenetics and calcium imaging instead of MEA-based systems [53]. Our approach offers a promising direction for advancing the field by integrating biological computation with machine learning methodologies, potentially yielding benefits such as improved energy efficiency and deeper insights into the dynamics of biological neural systems.

3 Reservoir of Biological Neurons

A network of biological neurons can be conceptualized as a collection of randomly interconnected computing units that exhibit complex, nonlinear dynamics. When such a network is cultured on a multi-electrode array (MEA) device – a grid of electrodes capable of both electrical stimulation and activity recording (see Fig. 2) – it becomes possible to deliver controlled input signals and monitor the corresponding network responses.

To utilize the biological network as a reservoir, we establish a direct mapping between each input sample and a predefined subset of MEA electrodes. For example, when processing an image, its $H \times W$ grid of pixels is mapped onto an $H \times W$ subset of the electrode array. Each pixel is associated with a specific electrode, and its intensity determines the stimulation parameters used to elicit a neural response. Once the electrical stimulus is delivered, we record the evoked spiking activity from the remaining electrodes on the array. The resulting spike trains are aggregated into a single high-dimensional vector that represents the high-dimensional latent space of the input sample. Thus, in this configuration, the biological reservoir functions as a biologically grounded feature extractor, or *bioware*, that transforms structured input signals into nonlinear, high-dimensional representations [8, 47].

To assess the quality of these representations, we feed the extracted features into a single-layer perceptron and evaluate the classification performance of the entire pipeline. A schematic overview of the system is provided in Fig. 1. The following sections provide details on the stimulation protocol, as well as the training and testing procedures concerning the single-layer perceptron classifier used in our experimental study.

3.1 Stimulation Protocol

Interfacing with the biological network presents unique challenges, particularly the need to calibrate multiple stimulation parameters – such as pulse amplitude, frequency, and duration – which must be




Fig. 2: **Schematic illustration of a cultured biological network interfaced with the MEA device.** Each square represents an individual MEA electrode. Input patterns are mapped onto the MEA by assigning elements of the patterns to specific electrodes. Electrical pulses are delivered based on the corresponding input intensities, and the evoked activity of the network is recorded via the remaining electrodes. The resulting spiking responses are used to construct a high-dimensional representation of the input in the feature space of the biological reservoir. Neuron scale in the figure is not to scale and is adjusted for visibility.

carefully optimized to ensure robust and consistent neuronal responses. To address these challenges, we investigate and define optimal stimulation protocols aimed at facilitating efficient neuronal activity and enabling reliable data recording.

Specifically, to interface with the cultured neural network, specific MEA electrodes are selected for delivering electrical stimulation. For bipolar stimulation, designated electrodes are assigned as positive and negative poles. Each electrical pulse delivered to the network is configured as a rectangular waveform – either monophasic or biphasic – with a specified amplitude A (measured in μA) and pulse width δ_+ or δ_- (measured in μs) for the positive and negative phases, respectively. Because the MEA distributes the total stimulation current across all selected electrode pairs, the amplitude A is interpreted on a per-pair basis. Stimulation is applied as a sequence of N identical pulses, spaced by a fixed inter-pulse interval T (in seconds). When configuring the stimulation parameters, we carefully balanced signal strength and hardware longevity. The stimulus must be strong enough to reliably elicit measurable neuronal responses, yet not so intense or repetitive as to risk damaging the electrodes – an issue observed during preliminary trials involving high-intensity or high-frequency stimulation.

Throughout the stimulation protocol, network activity was recorded continuously. To evaluate the efficacy of the stimulation parameters, we developed a visualization tool that allows for the assessment of stimulation quality. The tool represents the 64×64 MEA electrode grid as an image, where each pixel corresponds to a specific electrode and the recorded signals are processed using the precise spike time detection (PTSD) algorithm [42], a widely-used method for accurately extracting

spiking activity from extracellular recordings. Specifically, for the stimulus given at time t_s , for each electrode (i, j) , the tool calculates the activity a during a temporal chunk C of the recorded activity, defined as the number of spikes observed during this chunk:

$$a_{ij}^C(t_s) = \sum_{t=t_s}^{t_s+C} s_{ij}(t). \quad (1)$$

Here, $s_{ij}(t)$ is a binary variable indicating whether a spike was recorded at time t , and the sum runs over discrete time steps, determined by the recording device’s frame rate, which in our case was the maximum available rate of $f_s = 20,000$ Hz. To quantify the response of the network to each stimulus, we define the response r on a given electrode as the difference between the activity in a temporal chunk C immediately following the stimulus and the activity in a chunk $-C$ immediately preceding it:

$$r_{ij}^C(t_s) = a_{ij}^C(t_s) - a_{ij}^{-C}(t_s). \quad (2)$$

The tool then visualizes the response of each electrode to a given stimulus by encoding it into the color intensity of the corresponding pixel. Specifically, since we have N recordings (one for each stimulus), we take the average of r_{ij}^C over the N repetitions. In addition to the average, we compute the standard deviation of the observations and estimate confidence intervals for the response using a Student’s t-estimator. Red pixels indicate a positive response (increased activity) at a confidence level above 99%, green indicates a positive response at a 95% confidence level, cyan denotes a significant negative response (decreased activity) at a 95% confidence level, and yellow marks a negative response at a 99% confidence level. Figure 3 shows a representative example of the network response visualized using our tool.

3.2 Classifier Training Procedure

The training phase of the single-layer perceptron classifier begins by selecting a collection of input patterns to be used for network stimulation through the MEA interface. Each pattern is mapped to a predefined subset of MEA electrodes designated for stimulation. For every input, the corresponding electrodes are activated via electrical pulses repeated $N = 25$ times to obtain a distribution of responses suitable for statistical analysis. To mitigate potential carry-over effects and ensure reliable neuronal responses, a $T = 10$ s interval is enforced between successive stimuli, allowing the cultured network to return to its resting state [47]. We set a temporal chunk $C = 10$ ms – corresponding to approximately two synaptic transmission cycles [48] – and the number of spikes detected on each electrode within this window is counted. This results in a 4096-dimensional feature vector representing the input pattern in the latent space defined by the biological reservoir. To test its effectiveness in transforming the input and propagating information from the stimulated neurons to the others, we excluded a square region around the stimulated area.

Once latent representations for all input patterns have been collected, a linear classifier is trained to associate each feature vector with its corresponding input category. For each class, the recordings are randomly shuffled and split into two independent sample sets for training and testing – we refer to $N_{training}$ and $N_{testing}$ to express their cardinalities, where $N = N_{training} + N_{testing}$. Specifically, we set $N_{training} = 20$ and $N_{testing} = 5$. Classification is performed using a single-layer perceptron trained via stochastic gradient descent (SGD) [5], minimizing the cross-entropy loss function [29]. No separate validation set is used, as no early stopping is performed. This decision maximizes the




Fig. 3: **Visualization of the average network response to N stimuli over a temporal chunk C .** Each pixel represents an electrode in the 64×64 MEA grid, with the color intensity encoding the average level of response following the stimuli. We set $C = 10$ ms and $N = 25$.

amount of data available for both training and evaluation. The classifier is trained for 20 epochs, a value determined empirically to ensure convergence of the learning process. Finally, the trained model is used in the evaluation phase to assess the performance of the biological reservoir as a high-dimensional feature extractor.

3.3 Classifier Testing Procedure

To evaluate the performance of the trained model, we use the remaining $N_{testing} = 5$ network activity responses for each input pattern that were withheld during training. The linear classifier is kept fixed in the state achieved at the final training epoch, without further updates or fine-tuning.

Each test sample is acquired following the same protocol used during training: the input pattern is mapped onto the MEA electrodes, electrical stimulation is delivered accordingly, and the resulting network activity response is recorded. This activity is then processed using the same spike detection algorithm [42] to extract precise spiking events. From the extracted spike trains, we compute the number of spikes recorded on each electrode within a temporal window of duration $C = 10$ ms, immediately following the stimulus. This spike count vector serves as the latent representation of the test sample.

Because the testing procedure mirrors the training methodology, the distribution of test features remains consistent with that of the training data, supporting a fair and meaningful evaluation. Each test representation is passed through the trained classifier, and the resulting predictions are compared against the ground truth. We compute classification accuracy – defined as the ratio of correct predictions to the total number of samples – as the primary metric for assessing the effectiveness of our biologically-grounded reservoir computing system.

4 Experimental Evaluation

In this section, we present the experiments conducted to assess our BRC system and discuss the obtained results. Specifically, we evaluate the classification performance of the biological feature representations using three sets of experiments, each based on a different input pattern: (i) pointwise stimuli, (ii) oriented bars, and (iii) digits. Figure 4 illustrates the various input patterns on the MEA device. Table 1, on the other hand, summarizes the classification accuracy results. To enable a direct comparison between our biological reservoir and an artificial one, we implemented an ESN with 4,096 recurrently connected rate-based units, matching the dimensionality of the feature vectors obtained from the MEA, which consists of 4,096 electrodes. The recurrent connection matrix was initialized with a sparsity of 10%, consistent with commonly used ESN configurations and intended to reflect the sparse connectivity observed in cortical circuits. To ensure stable and rich internal dynamics, we set the spectral radius to 0.9 – a common practice in ESN design that balances dynamic richness with fading memory. Although our reservoir state does not evolve over many time steps, and long-term memory stability was not required, this setting ensures consistent and meaningful dynamics. To mirror our experimental setup, the ESN was initialized to a resting state with all units set to zero and driven by the same external stimulus patterns used in the biological experiments. To emulate the variability of spontaneous activity observed in the biological reservoir, we introduced input noise modeled from empirical data: we analyzed $N = 25$ randomly selected 10 ms windows of spontaneous neural activity and computed the average spike count. This estimate was used to generate realistic noise, which was superimposed on the input before being fed into the ESN. After stimulus presentation, the ESN dynamics evolved for one additional time step, and the resulting state was fed to a single-layer perceptron trained to classify the input patterns. While the ESN serves as a useful benchmark, it should be considered an upper bound, as we expect it to perform better by design. Nevertheless, our BRC system retains the distinctive advantages of a biological substrate.

In the following, we detail the specific setups and obtained results for each set of experiments.

4.1 Pointwise Stimuli Experiments

In this experimental setup, we applied pointwise input patterns by electrically stimulating single pairs of adjacent electrodes, designating one as positive and the other as negative. A simple monophasic electrical signal was employed for the stimulation in this scenario. To select the region of interest for pointwise stimulation, we first analyzed the spontaneous network activity recorded through the MEA. Specifically, we identified four distinct electrodes in different locations that exhibited high spontaneous spike counts and selected them as positive stimulation endpoints. For each positive endpoint, the immediately adjacent electrodes were designated as negative counterparts (see also the first row of Fig. 4). Electrical stimuli were delivered to each electrode pair with an amplitude of $A = 10 \mu\text{A}$ and a pulse width of $\delta = 20 \mu\text{s}$.

Following the training and testing procedure, the recorded network activity was fed into a linear classifier for pattern recognition, and the classifier was trained to differentiate between the four distinct stimuli based on the spiking network activity. The results, shown in Tab. 1, indicate a classification accuracy of $98\% \pm 2\%$, consistent with expectations given the simplicity of the task. The stimuli consisted of spatial patterns that were already linearly separable, likely contributing to the high performance. Notably, our BRC system outperforms the artificial ESN competitor, which struggles despite the simplicity of the task. We argue that this is due to the way noise

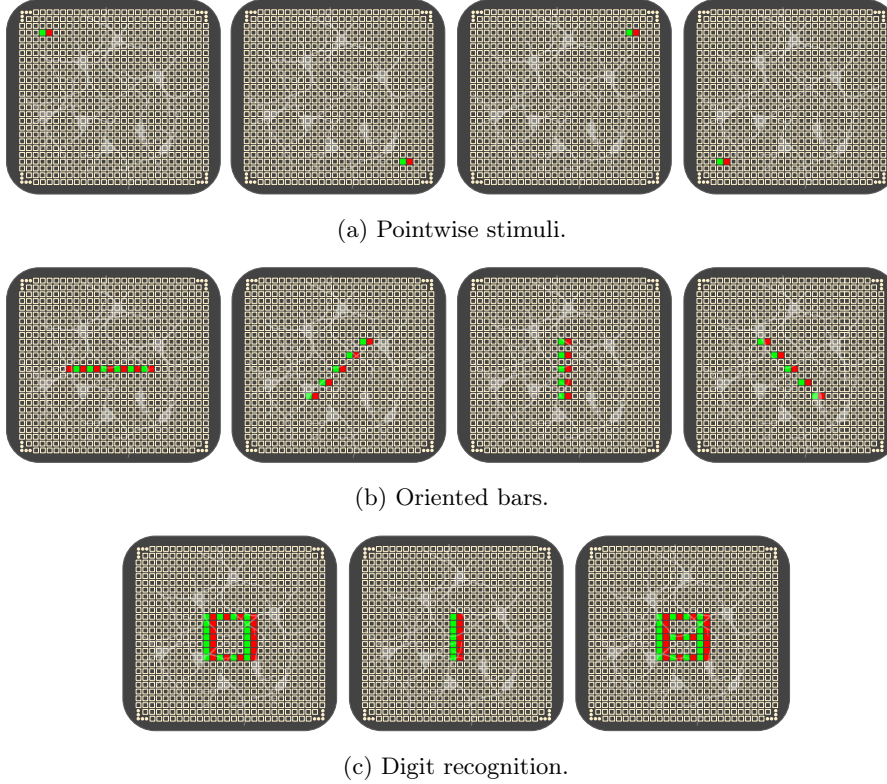


Fig. 4: **Visual representation of the input patterns used across the three experimental scenarios.** The first row displays the pointwise stimuli experiments, where single pairs of adjacent electrodes were stimulated. The second row shows the oriented bar experiments with four categories of bars at 0, 45, 90, and 135 degrees. The third row illustrates the digit recognition experiments with three patterns representing the digits 0, 1, and 8. Red and green indicated the positive and negative electrodes of the MAE, respectively.

is introduced in the ESN: when applied to pointwise stimuli, the noise can disrupt the signal by affecting specific input points, thereby degrading the informative content. This issue is specific to this set of experiments, as the subsequent scenarios rely on spatially distributed input patterns rather than isolated pointwise stimulation. While the simplicity of this task facilitated high accuracy, it nonetheless provided valuable confirmation that the biological reservoir responds effectively under controlled, linearly separable conditions.

4.2 Oriented Bars Experiments

The previous experiments with spatially distinct pointwise stimuli served as an effective preliminary stage for calibrating the stimulation protocol and validating the feasibility of linear classification as a downstream task for our BRC system. However, the simplicity of that setup, in which the

Table 1: **Experimental Results.** Classification accuracy (light-blue column) achieved by our BRC system across the three experimental settings, corresponding to different input patterns: (i) pointwise stimuli, (ii) oriented bars, and (iii) digit recognition. For each set, we show the mean \pm standard deviation across individual experiments, along with the overall average. For comparison, we also report the performance of an artificial ESN system that mirrors the same experimental setup.

Scenario	BRC Acc. (%)	Artificial ESN Acc. (%)
Pointwise Stimuli		
Point 1	100% \pm 0%	89% \pm 13%
Point 2	100% \pm 0%	71% \pm 19%
Point 3	100% \pm 0%	85% \pm 18%
Point 4	95% \pm 9%	79% \pm 24%
Average	98% \pm 2%	82% \pm 6%
Oriented Bars		
Bar 1 (0 degrees)	95% \pm 9%	100% \pm 0%
Bar 2 (45 degrees)	91% \pm 11%	100% \pm 0%
Bar 3 (90 degrees)	87% \pm 27%	98% \pm 3%
Bar 4 (135 degrees)	95% \pm 9%	100% \pm 0%
Average	92% \pm 6%	98% \pm 3%
Digit Recognition		
Digit 0	87% \pm 19%	100% \pm 0%
Digit 1	100% \pm 0%	100% \pm 0%
Digit 8	100% \pm 0%	92% \pm 10%
Average	95% \pm 6%	97% \pm 3%

patterns were linearly separable and spatially isolated, provided limited insight into the capacity of the reservoir to handle more complex and overlapping inputs.

To probe the ability of the system to discriminate among more challenging spatiotemporal patterns, we designed a second experimental scenario using stimuli composed of bar-shaped activation patterns. Specifically, we defined four stimulus categories corresponding to bars oriented at 0, 45, 90, and 135 degrees. Each pattern consists of a sequence of five pairs of adjacent electrodes (each comprising one positive and one negative endpoint) aligned along the designated orientation, with each pair spaced apart by one dilation step (see also the second row of Fig. 4). As in the previous experiment, we employed monophasic electrical pulses. Importantly, all four bar stimuli were presented at the same location on the MEA grid – centered on the region exhibiting the highest spontaneous activity. This design introduces a significant challenge: different stimulus classes activate spatially overlapping regions of the neural culture, thereby eliciting partially overlapping patterns of network activity. As a result, the downstream classifier can no longer rely on simple spatial separability; instead, the BR must perform meaningful spatiotemporal transformations to map these inputs into distinct and linearly separable latent representations. The remaining aspects of the stimulation protocol and classification pipeline were kept consistent with those described previously.

The results are shown in Tab. 1. The introduction of spatially overlapping bar-shaped stimuli increased the task complexity for the downstream linear classifier. The classifier achieved an average accuracy of $92\% \pm 6\%$, a slight decrease compared to the accuracy observed in the simpler pointwise experiment. These findings suggest that, despite the increased overlap in evoked activity across stimulus classes, the BRC system remains capable of producing distinct and linearly separable representations. This indicates the intrinsic ability of the BRC system to encode and transform input stimuli into meaningful feature embeddings, even under conditions of high spatial interference. As expected, performance was slightly lower than that of the artificial ESN, yet the BRC system maintained competitive results.

4.3 Digit Recognition Experiments

Finally, we evaluated the performance of the BRC system in a more complex pattern recognition task: digit recognition. In this scenario, patterns shaped like digits were defined using subsets of MEA electrodes. Specifically, we designed three distinct patterns representing the digits 0, 1, and 8, each mapped onto the MEA as *electronic clock* patterns. This approach mimics the display of a digital clock, where each electrode subset corresponds to an LED segment activated to form the digit. Specifically, for digit 1, a vertical sequence of seven electrode pairs (each consisting of a positive and a negative electrode) forms a vertical bar. The digit 0 is represented by two vertical bars of five electrode pairs spaced four electrodes apart, along with two horizontal bars at the top and bottom. Digit 8 is structured similarly to 0, but with an additional horizontal bar in the center. These input patterns are illustrated in the third row of Fig. 4. They were selected for their high degree of overlap and because they are not linearly separable, making the classification task more challenging than in the previous scenarios.

To support this task, we adopted a more complex biphasic stimulation protocol. The stimulus amplitude was reduced to $4\ \mu\text{A}$ per electrode pair, maintaining consistent neural activation despite the lower intensity. This was feasible due to the broader distribution of the stimulation across more electrodes. Additionally, the approach improved the balance of current flow between the positive and negative endpoints, helping to preserve electrode integrity over time. As in the previous scenario, all digit patterns were presented within the same region of the MEA grid, corresponding to the area of maximal spiking activity.

Results are presented in Tab. 1. This third experimental scenario ranks between the two previous ones, with the linear classifier achieving an accuracy of $95\% \pm 6\%$. Although the digit recognition task is, in principle, more challenging than the oriented bar experiments, we hypothesize that the higher performance is due to the greater number of stimulated electrodes. In the oriented bar scenario, fewer electrodes are activated, likely resulting in more localized and weaker network responses. The artificial ESN competitor performs on average slightly better, reaching an accuracy of 97%, as expected. Nevertheless, our BRC system delivers competitive results while retaining the inherent advantages of a biological reservoir.

5 Conclusion and Future Works

In this work, we introduced an innovative approach to RC by employing a network of cultured biological neurons as the reservoir substrate. This bio-inspired paradigm presents promising prospects for enhancing energy efficiency, given the inherently lower energy consumption of biological neurons compared to artificial models. Additionally, studying biological neural networks offers valuable

insights into the core mechanisms driving cognition, advancing our understanding of brain function and computation. We successfully developed an interfacing protocol via multi-electrode array (MEA) devices and conducted an experimental investigation to evaluate the ability of the biological reservoir to recognize several input patterns. The results were promising, suggesting that biological neurons can indeed serve as a viable substrate for RC tasks, generating high-dimensional feature representations for pattern recognition. These findings lay the groundwork for further exploration of biological networks in the context of RC.

Looking ahead, we intend to expand this investigation by incorporating more complex pattern recognition tasks, allowing us to assess the scalability and robustness of the biological reservoir in a broader range of applications. Additionally, we plan to explore the integration of optogenetic stimulation protocols [44], complementing traditional electrical stimulation, as a strategy to mitigate electrode degradation and enhance the precision and reliability of neural stimulation. This approach could lead to more sustainable and flexible neural interfacing solutions. Lastly, investigating stimulation protocols that induce synaptic plasticity will help us understand how biological learning mechanisms could be leveraged to optimize network behavior, steering it toward more task-specific and efficient responses. These future endeavors hold significant potential for advancing both the understanding and practical applications of biological reservoir computing, bringing us closer to more energy-efficient and biologically grounded cognitive computing systems.

Acknowledgments. The research was in part supported by the Matteo Caleo Foundation, by Scuola Normale Superiore (FC), by the PRIN AICult grant #2022M95RC7 from the Italian Ministry of University and Research (MUR) (FC) and by the Tuscany Health Ecosystem - THE grant from MUR (FC, GA, ADG).

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Abbott, L.F., van Vreeswijk, C.: Asynchronous states in networks of pulse-coupled oscillators. *Phys. Rev. E* **48**, 1483–1490 (Aug 1993). <https://doi.org/10.1103/PhysRevE.48.1483>
2. Badar, A., Varma, A., Staniec, A., Gamal, M., Magdy, O., Iqbal, H., Arani, E., Zonooz, B.: Highlighting the importance of reducing research bias and carbon emissions in cnns. In: *AIxIA 2021 - Advances in Artificial Intelligence - 20th International Conference of the Italian Association for Artificial Intelligence, Virtual Event, December 1-3, 2021, Revised Selected Papers. Lecture Notes in Computer Science*, vol. 13196, pp. 515–531. Springer (2021). https://doi.org/10.1007/978-3-031-08421-8_36
3. Bianchi, F.M., Scardapane, S., Løkse, S., Jenssen, R.: Reservoir computing approaches for representation and classification of multivariate time series. *IEEE Trans. Neural Networks Learn. Syst.* **32**(5), 2169–2179 (2021). <https://doi.org/10.1109/TNNLS.2020.3001377>
4. Bonifazi, P., Fromherz, P.: Silicon chip for electronic communication between nerve cells by non-invasive interfacing and analog–digital processing. *Advanced Materials* **14**(17), 1190–1193 (2002). [https://doi.org/https://doi.org/10.1002/1521-4095\(20020903\)14:17<1190::AID-ADMA1190>3.0.CO;2-#](https://doi.org/https://doi.org/10.1002/1521-4095(20020903)14:17<1190::AID-ADMA1190>3.0.CO;2-#)
5. Bottou, L.: *On-line learning and stochastic approximations*, p. 9–42. Cambridge University Press, USA (1999)
6. Carlson, K.D., Richert, M., Dutt, N.D., Krichmar, J.L.: Biologically plausible models of homeostasis and STDP: stability and learning in spiking neural networks. In: *The 2013 International Joint Conference on Neural Networks, IJCNN 2013, Dallas, TX, USA, August 4-9, 2013*. pp. 1–8. IEEE (2013). <https://doi.org/10.1109/IJCNN.2013.6706961>

7. Chambers, S.M., Fasano, C.A., Papapetrou, E.P., Tomishima, M., Sadelain, M., Studer, L.: Highly efficient neural conversion of human es and ips cells by dual inhibition of smad signaling. *Nature Biotechnology* **27**(3), 275–280 (Mar 2009). <https://doi.org/10.1038/nbt.1529>, <http://dx.doi.org/10.1038/nbt.1529>
8. Chiaradia, I., Lancaster, M.A.: Brain organoids for the study of human neurobiology at the interface of in vitro and in vivo. *Nature neuroscience* **23**(12), 1496–1508 (December 2020). <https://doi.org/10.1038/s41593-020-00730-3>
9. Ciampi, L., Lagani, G., Amato, G., Falchi, F.: A biologically-inspired approach to biomedical image segmentation. In: *Computer Vision - ECCV 2024 Workshops - Milan, Italy, September 29-October 4, 2024, Proceedings, Part XIV. Lecture Notes in Computer Science*, vol. 15636, pp. 158–171. Springer (2024). https://doi.org/10.1007/978-3-031-91578-9_10, https://doi.org/10.1007/978-3-031-91578-9_10
10. Ciampi, L., Lagani, G., Amato, G., Falchi, F.: Biologically-inspired semi-supervised semantic segmentation for biomedical imaging. *CoRR* **abs/2412.03192** (2024). <https://doi.org/10.48550/ARXIV.2412.03192>
11. Cossu, A., Bacciu, D., Carta, A., Gallicchio, C., Lomonaco, V.: Continual learning with echo state networks. In: *29th European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning, ESANN 2021, Online event (Bruges, Belgium), October 6-8, 2021* (2021). <https://doi.org/10.14428/ESANN/2021.ES2021-80>
12. Diehl, P.U., Cook, M.: Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Frontiers Comput. Neurosci.* **9**, 99 (2015). <https://doi.org/10.3389/FNCOM.2015.00099>
13. Dockendorf, K.P., Park, I., He, P., Príncipe, J.C., DeMarse, T.B.: Liquid state machines and cultured cortical networks: The separation property. *Biosystems* **95**(2), 90–97 (2009). <https://doi.org/https://doi.org/10.1016/j.biosystems.2008.08.001>
14. Durstewitz, D., Deco, G.: Computational significance of transient dynamics in cortical networks. *European Journal of Neuroscience* **27**(1), 217–227 (2008). <https://doi.org/https://doi.org/10.1111/j.1460-9568.2007.05976.x>
15. Ferrández, J.M., Lorente, V., de la Paz, F., Fernández, E.: Training biological neural cultures: Towards hebbian learning. *Neurocomputing* **114**, 3–8 (2013). <https://doi.org/10.1016/J.NEUCOM.2012.09.031>
16. Ferré, P., Mamalet, F., Thorpe, S.J.: Unsupervised feature learning with winner-takes-all based STDP. *Frontiers Comput. Neurosci.* **12**, 24 (2018). <https://doi.org/10.3389/FNCOM.2018.00024>
17. Gallicchio, C., Micheli, A.: Deep echo state network (deepesn): A brief survey. *CoRR* **abs/1712.04323** (2017)
18. Gerstner, W., Kistler, W.M.: *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press (2002). <https://doi.org/10.1017/CBO9780511815706>
19. Goel, A., Buonomano, D.V.: Temporal interval learning in cortical cultures is encoded in intrinsic network dynamics. *Neuron* **91**(2), 320–327 (2016). <https://doi.org/https://doi.org/10.1016/j.neuron.2016.05.042>, <https://www.sciencedirect.com/science/article/pii/S0896627316302537>
20. Hassabis, D., Kumaran, D., Summerfield, C., Botvinick, M.: Neuroscience-inspired artificial intelligence. *Neuron* **95**(2), 245–258 (Jul 2017). <https://doi.org/10.1016/j.neuron.2017.06.011>, <http://dx.doi.org/10.1016/j.neuron.2017.06.011>
21. Huang, G.: An insight into extreme learning machines: Random neurons, random features and kernels. *Cogn. Comput.* **6**(3), 376–390 (2014). <https://doi.org/10.1007/S12559-014-9255-2>
22. Huang, G., Zhu, Q., Siew, C.K.: Extreme learning machine: Theory and applications. *Neurocomputing* **70**(1-3), 489–501 (2006). <https://doi.org/10.1016/J.NEUCOM.2005.12.126>
23. Illing, B., Gerstner, W., Brea, J.: Biologically plausible deep learning - but how far can we go with shallow networks? *Neural Networks* **118**, 90–101 (2019). <https://doi.org/10.1016/J.NEUNET.2019.06.001>
24. Isomura, T., Kotani, K., Jimbo, Y.: Cultured cortical neurons can perform blind source separation according to the free-energy principle. *PLoS Comput. Biol.* **11**(12) (2015). <https://doi.org/10.1371/JOURNAL.PCBI.1004643>
25. Jaeger, H.: The “echo state” approach to analysing and training recurrent neural networks-with an erratum note. Bonn, Germany: German National Research Center for Information Technology GMD Technical Report **148**(34), 13 (2001)

26. Javed, F., He, Q., Davidson, L.E., Thornton, J.C., Albu, J., Boxt, L., Krasnow, N., Elia, M., Kang, P., Heshka, S., Gallagher, D.: Brain and high metabolic rate organ mass: contributions to resting energy expenditure beyond fat-free mass¹²³⁴. *The American Journal of Clinical Nutrition* **91**(4), 907–912 (2010). <https://doi.org/https://doi.org/10.3945/ajcn.2009.28512>
27. Journé, A., Rodriguez, H.G., Guo, Q., Moraitis, T.: Hebbian deep learning without feedback. In: *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net (2023)
28. Kagan, B.J., Kitchen, A.C., Tran, N.T., Habibollahi, F., Khajehnejad, M., Parker, B.J., Bhat, A., Rollo, B., Razi, A., Friston, K.J.: In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron* **110**(23), 3952–3969.e8 (2022). <https://doi.org/https://doi.org/10.1016/j.neuron.2022.09.001>
29. Kline, D., Berardi, V.L.: Revisiting squared-error and cross-entropy functions for training neural network classifiers. *Neural Comput. Appl.* **14**(4), 310–318 (2005). <https://doi.org/10.1007/S00521-005-0467-Y>
30. Krotov, D., Hopfield, J.J.: Unsupervised learning by competing hidden units. *Proc. Natl. Acad. Sci. USA* **116**(16), 7723–7731 (2019). <https://doi.org/10.1073/PNAS.1820458116>
31. Lagani, G., Bacciu, D., Gallicchio, C., Falchi, F., Gennaro, C., Amato, G.: Deep features for CBIR with scarce data using hebbian learning. In: *CBMI 2022: International Conference on Content-based Multimedia Indexing, Graz, Austria, September 14 - 16, 2022*. pp. 136–141. ACM (2022). <https://doi.org/10.1145/3549555.3549587>
32. Lagani, G., Falchi, F., Gennaro, C., Amato, G.: Training convolutional neural networks with competitive hebbian learning approaches. In: *Machine Learning, Optimization, and Data Science - 7th International Conference, LOD 2021, Grasmere, UK, October 4-8, 2021, Revised Selected Papers, Part I. Lecture Notes in Computer Science, vol. 13163*, pp. 25–40. Springer (2021). https://doi.org/10.1007/978-3-030-95467-3_2
33. Lagani, G., Falchi, F., Gennaro, C., Amato, G.: Comparing the performance of hebbian against back-propagation learning using convolutional neural networks. *Neural Comput. Appl.* **34**(8), 6503–6519 (2022). <https://doi.org/10.1007/S00521-021-06701-4>, <https://doi.org/10.1007/s00521-021-06701-4>
34. Lagani, G., Falchi, F., Gennaro, C., Amato, G.: Spiking neural networks and bio-inspired supervised deep learning: A survey. *CoRR* **abs/2307.16235** (2023). <https://doi.org/10.48550/ARXIV.2307.16235>
35. Lagani, G., Falchi, F., Gennaro, C., Amato, G.: Synaptic plasticity models and bio-inspired unsupervised deep learning: A survey. *CoRR* **abs/2307.16236** (2023). <https://doi.org/10.48550/ARXIV.2307.16236>
36. Lagani, G., Falchi, F., Gennaro, C., Fassold, H., Amato, G.: Scalable bio-inspired training of deep neural networks with fasthebb. *Neurocomputing* **595**, 127867 (2024). <https://doi.org/10.1016/J.NEUCOM.2024.127867>
37. Lagani, G., Gennaro, C., Fassold, H., Amato, G.: Fasthebb: Scaling hebbian training of deep neural networks to imagenet level. In: *Similarity Search and Applications - 15th International Conference, SISAP 2022, Bologna, Italy, October 5-7, 2022, Proceedings. Lecture Notes in Computer Science, vol. 13590*, pp. 251–264. Springer (2022). https://doi.org/10.1007/978-3-031-17849-8_20
38. Lagani, G., Mazziotti, R., Falchi, F., Gennaro, C., Cicchini, G.M., Pizzorusso, T., Cremisi, F., Amato, G.: Assessing pattern recognition performance of neuronal cultures through accurate simulation. In: *10th International IEEE/EMBS Conference on Neural Engineering, NER 2021, Virtual Event, Italy, May 4-6, 2021*. pp. 726–729. IEEE (2021). <https://doi.org/10.1109/NER49283.2021.9441166>
39. Lazar, A., Pipa, G., Triesch, J.: SORN: a self-organizing recurrent neural network. *Frontiers Comput. Neurosci.* **3**, 23 (2009). <https://doi.org/10.3389/NEURO.10.023.2009>
40. Lukosevicius, M., Jaeger, H.: Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* **3**(3), 127–149 (2009). <https://doi.org/10.1016/J.COSREV.2009.03.005>
41. Maass, W., Natschläger, T., Markram, H.: Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Comput.* **14**(11), 2531–2560 (2002). <https://doi.org/10.1162/089976602760407955>

42. Maccione, A., Gandolfo, M., Massobrio, P., Novellino, A., Martinoia, S., Chiappalone, M.: A novel algorithm for precise identification of spikes in extracellularly recorded neuronal signals. *Journal of Neuroscience Methods* **177**(1), 241–249 (2009). <https://doi.org/https://doi.org/10.1016/j.jneumeth.2008.09.026>
43. Marblestone, A.H., Wayne, G., Körding, K.P.: Toward an integration of deep learning and neuroscience. *Frontiers Comput. Neurosci.* **10**, 94 (2016). <https://doi.org/10.3389/FNCOM.2016.00094>
44. Meloni, I., Sachidanandan, D., Thum, A.S., Kittel, R.J., Murawski, C.: Controlling the behaviour of *Drosophila melanogaster* via smartphone optogenetics. *Scientific Reports* **10**(1), 1–11 (2020)
45. Pastore, V.P., Massobrio, P., Godjowski, A., Martinoia, S.: Identification of excitatory-inhibitory links and network topology in large-scale neuronal assemblies from multi-electrode recordings. *PLoS Comput. Biol.* **14**(8) (2018). <https://doi.org/10.1371/JOURNAL.PCBI.1006381>
46. Richards, B., Lillicrap, T., Beaudoin, P., Bengio, Y., Bogacz, R., Christensen, A., Clopath, C., Costa, R., de Berker, A., Ganguli, S., Gillon, C., Hafner, D., Kepecs, A., Kriegeskorte, N., Latham, P., Lindsay, G., Miller, K., Naud, R., Pack, C., Poirazi, P., Roelfsema, P., Sacramento, J., Saxe, A., Scellier, B., Schapiro, A., Senn, W., Wayne, G., Yamins, D., Zenke, F., Zylberberg, J., Therien, D., Kording, K.: A deep learning framework for neuroscience. *Nature Neuroscience* **22**(11), 1761–1770 (2019). <https://doi.org/10.1038/s41593-019-0520-2>
47. Ruaro, M.E., Bonifazi, P., Torre, V.: Toward the neurocomputer: image processing and pattern recognition with neuronal cultures. *IEEE Trans. Biomed. Eng.* **52**(3), 371–383 (2005). <https://doi.org/10.1109/TBME.2004.842975>
48. Sabatini, B.L., Regehr, W.G.: Timing of synaptic transmission. *Annual Review of Physiology* **61**(Volume 61, 1999), 521–542 (1999). <https://doi.org/https://doi.org/10.1146/annurev.physiol.61.1.521>
49. Sarli, D.D., Gallicchio, C., Micheli, A.: Gated echo state networks: a preliminary study. In: *International Conference on INnovations in Intelligent SysTems and Applications, INISTA 2020, Novi Sad, Serbia, August 24-26, 2020*. pp. 1–5. IEEE (2020). <https://doi.org/10.1109/INISTA49547.2020.9194681>
50. Scardapane, S., Butcher, J.B., Bianchi, F.M., Malik, Z.K.: Advances in biologically inspired reservoir computing. *Cogn. Comput.* **9**(3), 295–296 (2017). <https://doi.org/10.1007/S12559-017-9469-1>
51. Shahaf, G., Marom, S.: Learning in networks of cortical neurons. *Journal of Neuroscience* **21**(22), 8782–8788 (2001). <https://doi.org/10.1523/JNEUROSCI.21-22-08782.2001>
52. Song, S., Miller, K.D., Abbott, L.F.: Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience* **3**(9), 919–926 (Sep 2000). <https://doi.org/10.1038/78829>
53. Sumi, T., Yamamoto, H., Katori, Y., Ito, K., Moriya, S., Konno, T., Sato, S., Hirano-Iwata, A.: Biological neurons act as generalization filters in reservoir computing. *Proceedings of the National Academy of Sciences* **120**(25) (2023). <https://doi.org/10.1073/pnas.2217008120>
54. Sun, Z., Cutsuridis, V., Caiafa, C.F., Solé-Casals, J.: Brain simulation and spiking neural networks. *Cogn. Comput.* **15**(4), 1103–1105 (2023). <https://doi.org/10.1007/S12559-023-10156-1>
55. Tenenbaum, J.: Building machines that learn and think like people. In: André, E., Koenig, S., Dastani, M., Sukthankar, G. (eds.) *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*. p. 5. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM (2018)
56. Wang, Q., Li, P.: D-LSM: deep liquid state machine with unsupervised recurrent reservoir tuning. In: *23rd International Conference on Pattern Recognition, ICPR 2016, Cancún, Mexico, December 4-8, 2016*. pp. 2652–2657. IEEE (2016). <https://doi.org/10.1109/ICPR.2016.7900035>
57. Yonemura, Y., Katori, Y.: Dynamical predictive coding with reservoir computing performs noise-robust multi-sensory speech recognition. *Frontiers Comput. Neurosci.* **18** (2024). <https://doi.org/10.3389/FNCOM.2024.1464603>
58. Zhang, Y., Li, P., Jin, Y., Choe, Y.: A digital liquid state machine with biologically inspired learning and its application to speech recognition. *IEEE Trans. Neural Networks Learn. Syst.* **26**(11), 2635–2649 (2015). <https://doi.org/10.1109/TNNLS.2015.2388544>