

From Neural Activity to Computation: Biological Reservoirs for Pattern Recognition in Digit Classification

Ludovico Iannello^{*1}, Luca Ciampi^{*1}, Fabrizio Tonelli³, Gabriele Lagani¹, Lucio Maria Calcagnile², Federico Cremisi³, Angelo Di Garbo², Giuseppe Amato¹
¹ISTI-CNR, Italy ²IBF-CNR, Italy ³Bio@SNS, Italy

^{*}ludovico.iannello@isti.cnr.it

^{*}luca.ciampi@isti.cnr.it

Abstract

In this paper, we present a biologically grounded approach to reservoir computing (RC), in which a network of cultured biological neurons serves as the reservoir substrate. This system, referred to as biological reservoir computing (BRC), replaces artificial recurrent units with the spontaneous and evoked activity of living neurons. A multi-electrode array (MEA) enables simultaneous stimulation and readout across multiple sites: inputs are delivered through a subset of electrodes, while the remaining ones capture the resulting neural responses, mapping input patterns into a high-dimensional biological feature space. We evaluate the system through a case study on digit classification using a custom dataset. Input images are encoded and delivered to the biological reservoir via electrical stimulation, and the corresponding neural activity is used to train a simple linear classifier. To contextualize the performance of the biological system, we also include a comparison with a standard artificial reservoir trained on the same task. The results indicate that the biological reservoir can effectively support classification, highlighting its potential as a viable and interpretable computational substrate. We believe this work contributes to the broader effort of integrating biological principles into machine learning and aligns with the goals of human-inspired vision by exploring how living neural systems can inform the design of efficient and biologically plausible models.

1. Introduction

Reservoir computing (RC) [37] is a machine learning framework that transforms input data through the dynamics of a high-dimensional system. This transformation often results in representations that are more easily separable by simple classifiers, making RC a practical and efficient approach for tasks such as classification and regression (e.g.,

time-series analysis [3] and speech recognition [51]). Its effectiveness has been demonstrated across a range of implementations. One of the most prominent RC models is the echo state network (ESN) [15, 23], in which a large pool of randomly connected recurrent units transforms input into a high-dimensional nonlinear representation that facilitates downstream learning. A related paradigm is the liquid state machine (LSM) [38, 52], which employs spiking neurons to capture rich temporal dynamics.

In this work, we propose an instantiation of the reservoir computing (RC) paradigm where the reservoir is not simulated but physically realized through a living network of cultured neurons. We refer to this approach as *biological reservoir computing* (BRC). Unlike conventional RC systems that rely on artificial units, our method leverages the intrinsic dynamics of a biological neural culture to project input stimuli into a high-dimensional feature space. This setup offers a unique opportunity to explore the computational potential of real neural tissue, while also contributing to the broader goal of integrating biological principles into machine learning. Indeed, this biologically grounded approach not only provides a natural substrate for computation, but also aligns with neuroscientific evidence suggesting that transient neural dynamics play a key role in information processing [12]. Moreover, by offloading part of the computation to a physical system, BRC may offer advantages in terms of energy efficiency—an increasingly relevant concern in modern AI systems [2, 24].

Specifically, the BRC system is built upon a high-density multi-electrode array (MEA) [4], which enables both electrical stimulation and high-resolution recording of neural activity. Neurons are derived from stem cells using established differentiation protocols [6, 7], and once matured, they form spontaneously active networks. Input patterns are encoded as spatially distributed stimulation sequences and delivered to the culture via the MEA. The resulting neural responses are recorded and used to construct feature vectors for downstream classification. To evaluate the ability of our BRC system to generate discriminative feature repre-

^{*}Equal contribution.

sentations, we conducted an experimental study focused on a single yet challenging pattern recognition task: the classification of ten digit-like spatial input patterns (from 0 to 9). Each pattern was defined by a distinct spatial configuration of electrical stimulation sites on a multi-electrode array (MEA). The evoked neural responses were recorded and used as features for a simple readout layer to classify the stimulus category.

Despite the inherent variability in biological responses due to noise, spontaneous activity, and differences across stimulation sessions and biological replicates, our system achieved promising levels of accuracy. These results demonstrate that cultured neuronal networks, even without considering plasticity or learning mechanisms, can serve as effective reservoirs that transform static spatial inputs into rich, high-dimensional representations suitable for downstream classification tasks. This manuscript is structured as follows. Sec. 2 surveys existing literature on reservoir computing and biological computation. Sec. 3 details the experimental methodology, including MEA interfacing and stimulation protocols. In Sec. 4, we present and analyze our experimental setup and results. Finally, Sec 5 shapes the conclusion and outlines directions for future research.

2. Related Works

A central topic in deep learning (DL) research concerns the biological plausibility of existing computational frameworks. Numerous studies have critically examined the extent to which contemporary neural architectures reflect the structural and functional complexity of biological neural systems [18, 39, 41, 48]. In light of this, there has been a growing movement toward the development of biologically inspired alternatives aimed at expanding the capabilities of machine learning and cognitive modeling [34, 35]. Efforts in this direction have led to a wide array of models that attempt to bridge the gap between artificial and biological intelligence, either by implementing more biologically grounded neural computations [11, 14, 30, 47] or by refining the underlying synaptic and learning mechanisms [8, 9, 21, 25, 28, 29, 31–33, 36].

Among these biologically inspired paradigms, reservoir computing (RC) [37, 44] has attracted significant attention for its capacity to emulate complex neural dynamics. Two prominent subclasses within RC are echo state networks (ESNs) [15, 23] and liquid state machines (LSMs) [38, 52], which differ in their computational mechanisms and biological plausibility. While LSMs typically employ spiking neuron models [1, 16] to generate rich temporal dynamics, ESNs rely on a large, recurrent network of randomly connected continuous-valued units to embed input signals into a high-dimensional space conducive to downstream classification tasks. Although the recurrent architecture is initialized stochastically, specific strategies are employed to

ensure that the network operates within a stable and computationally useful dynamic regime [15, 43]. Several enhancements to the ESN framework have been proposed to increase its biological plausibility. For instance, models incorporating synaptic plasticity mechanisms, such as spike-timing-dependent plasticity (STDP) [16, 46], aim to more closely emulate biological learning processes [50]. More recently, gating mechanisms have been introduced to improve the recall of long-term dependencies in nonlinear recurrent systems [43]. These RC-inspired techniques have demonstrated promising results across various application domains, including speech recognition [52] and continual learning [10].

Building upon this foundation, the present work introduces an extension to the RC paradigm by incorporating biological neural networks as computational reservoirs, i.e., a biological reservoir computing (BRC) wherein cultured neuronal populations serve as the dynamic substrate for computation. Although prior studies have explored the use of multi-electrode array (MEA) devices to interface with biological neurons [13, 17, 22, 26, 40, 42, 45], only a few works have investigated the potential of employing biological neural networks as computational reservoirs [5, 19]. Notably, Cai et al. [5] employed brain organoids as biological reservoirs, demonstrating their capacity for speech recognition by leveraging the rich temporal dynamics inherent in such 3D biological structures. Their approach capitalizes on time-dependent neural activity patterns to encode and process input sequences. In contrast, our study focuses on spatially distributed stimulation patterns delivered to 2D cultured biological neural networks, rather than 3D organoids. Specifically, we employ spatially patterned electrical stimulation across a microelectrode array (MEA) without relying on temporal sequencing, aiming instead at static pattern recognition. This represents a distinct paradigm, emphasizing spatial encoding of input patterns over their temporal evolution. Similarly, our preliminary work [19] explored the feasibility of stimulating 2D cultured biological neural networks using a limited set of spatial patterns and network configurations. However, [19] was primarily exploratory, involving a small number of stimuli and simpler classification tasks applied to a single biological network. By systematically investigating a broader range of spatially distributed stimulation patterns and evaluating their effects across three independent biological replicates (BRs), the present study constitutes one of the first comprehensive efforts to assess the feasibility of using cultured biological neural networks as functional reservoirs within a reservoir computing (RC) framework for static pattern recognition. This approach not only bridges biological neural substrates and machine learning architectures, but also opens new avenues for leveraging the intrinsic properties of biological networks – such as energy efficiency and complex nonlin-

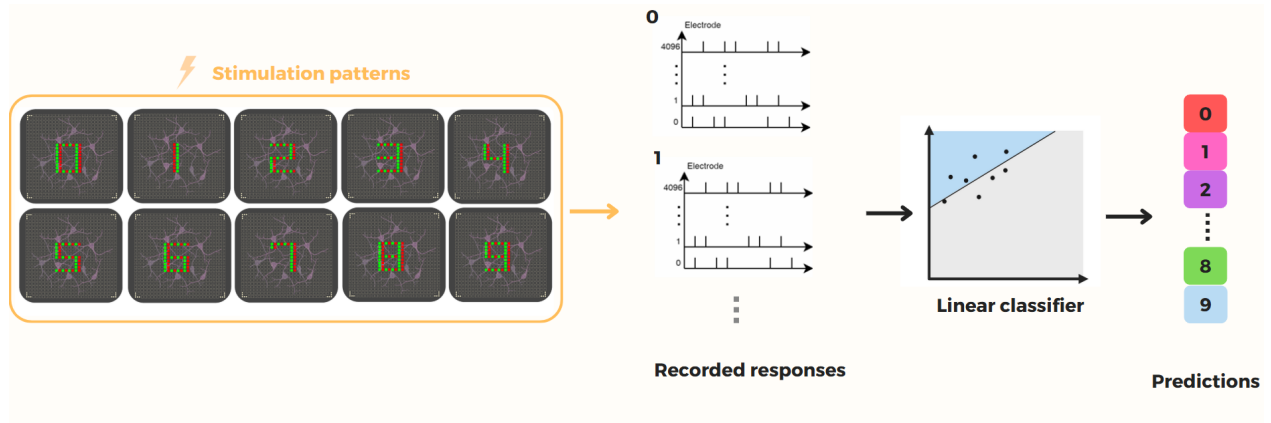


Figure 1. **Framework of our Biological Reservoir Computing (BRC) paradigm.** In this approach, a multi-electrode array (MEA) functions as a bidirectional interface, enabling both the stimulation of and the recording from a cultured biological neural network. Discrete inputs are encoded by selectively activating specific subsets of MEA electrodes, which deliver targeted stimuli to the network. The evoked spiking responses are captured via a separate set of electrodes and transformed into high-dimensional vectors that encode the input within a latent computational space. Due to the inherent complexity and rich dynamics of the biological reservoir, this transformation is intrinsically nonlinear. Subsequently, a linear classifier is trained to infer the category of the original input from its corresponding latent representation.




Figure 2. **Illustration of a cultured biological network plated on a MEA device.** Each square represents a single 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.

ear dynamics – to advance neuromorphic computation.

3. Biological Reservoir Computing

A cultured network of biological neurons may be regarded as a complex system of randomly interconnected processing elements capable of generating rich, nonlinear dynamical behavior. When interfaced with a multi-electrode array (MEA) – a bidirectional platform equipped for both electrical stimulation and electrophysiological recording (see

Fig. 2) – this neural substrate can be functionally harnessed to receive structured input stimuli and produce measurable responses.

In our reservoir computing paradigm, we define a direct correspondence between individual input samples and specific subsets of MEA electrodes. For instance, when an image is used as input, its pixel grid of dimensions $H \times W$ is projected onto an $H \times W$ subset of the electrode array. Each pixel is uniquely mapped to a designated electrode, and the pixel intensity modulates the stimulation parameters delivered to the underlying neurons. Following the stimulation phase, we collect the evoked spiking activity through the remaining electrodes on the MEA. These recorded spike trains are then aggregated into a high-dimensional feature vector, effectively embedding the input into the latent space defined by the biological reservoir.

In this configuration, the cultured neural network operates as a biologically instantiated feature extractor, transforming structured inputs into complex, nonlinear, high-dimensional representations [7, 42].

To quantitatively assess the quality and discriminative power of these biologically derived representations, we employ a single-layer perceptron as a downstream classifier. The full computational pipeline – including stimulation protocol and classification framework – is illustrated in Fig. 1. The following sections detail the procedures for stimulation, feature extraction, training, and performance evaluation of the perceptron within our experimental setup.

3.1. Stimulation Protocol

Interfacing with a biological neural network involves several technical challenges, particularly with respect to the

calibration of stimulation parameters such as pulse amplitude, frequency, and duration. These parameters must be finely tuned to ensure that stimulation elicits consistent and robust neural responses while maintaining the integrity of the hardware. In this work, we systematically investigate and define a set of optimized stimulation protocols designed to activate the cultured network effectively and support reliable electrophysiological recordings.

To deliver stimulation, specific electrodes within the MEA are designated as active stimulation sites. We used a bipolar stimulation, where each electrode pair is configured with opposing polarities, serving as the positive and negative poles. Each stimulus is delivered as a rectangular biphasic waveform with a defined amplitude A (expressed in μA) and phase durations δ_+ and δ_- (in μs) corresponding to the positive and negative components, respectively. This approach enhanced the current balance between the positive and negative poles, contributing to the long-term stability and integrity of the electrodes. Since the MEA distributes current across multiple electrode pairs, the amplitude is interpreted on a per-pair basis. When determining these stimulation parameters, particular care is taken to balance signal efficacy with hardware preservation: the stimuli must be sufficiently strong to reliably evoke spiking activity, yet not so intense or frequent as to cause electrode degradation—a risk observed during initial trials involving high-current or high-frequency configurations.

To mitigate any potential bias in the network’s response due to order effects or temporal dependencies, all stimulation patterns are presented to the neuronal culture in a randomized sequence. A fixed inter-stimulus interval of $T = 10\text{ s}$ is enforced between successive inputs. This delay allows the network to return to a baseline or resting state before the subsequent stimulation is delivered, thereby ensuring the independence of evoked responses and reducing carry-over effects [42].

To implement this protocol on the 3Brain high-density multi-electrode array (HD-MEA) platform, we developed a custom Python script using the official API provided by the manufacturer. This custom script enables precise control over the stimulation sequence, including the randomization of input order, the timing of pulse delivery, and the selection of specific electrode subsets for stimulation, ensuring reproducibility and flexibility in our experimental pipeline.

3.2. Biological Network Readout

For each stimulus presentation, neural activity is continuously recorded from the MEA for a time window extending from 2s before to 2s after the stimulus onset. To extract spiking activity from the raw extracellular signals in real time, we employed a spike detection algorithm based on a double-threshold strategy, which enables online identification of spike times.

The algorithm operates using three parameters: a sliding time window w_s , a low detection threshold thr_l , and a high detection threshold thr_h used for final spike identification. Initially, for each recording channel, the algorithm scans the signal within a local window of duration $w_s = 2\text{ ms}$ to identify candidate peaks exceeding $thr_l \times \sigma$, where σ denotes the standard deviation of the full recorded signal.

To refine the detection, all segments of the signal that cross the low threshold are temporarily excluded, and a new standard deviation σ_n is computed on the remaining data. Final spike times are then determined as the local maxima that exceed the updated high threshold, defined as $thr_h \times \sigma_n$.

Now, let t_s denote the onset time of a stimulus. For each electrode (i, j) , we compute the spiking activity $a_{ij}^C(t_s)$ within a temporal window W following the stimulus, defined as the number of spikes observed:

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

where $s_{ij}(t)$ is a binary variable that equals 1 if a spike is detected at time t , and 0 otherwise. Time is discretized according to the sampling rate of the acquisition system, which in our setup was the maximum available frequency $f_s = 20,000\text{ Hz}$. We use $a_{ij}^W(t_s)$ with a properly defined time window (W) as the readout of the stimulus. The result is a 4096-dimensional feature vector that encodes the input pattern in the latent space of the biological reservoir. To evaluate the real network’s ability to propagate information beyond the stimulation site, a square region surrounding the stimulated electrodes is excluded from this feature vector.

3.3. Classifier Training and Testing Phase

The training phase of the single-layer perceptron classifier begins with the selection of a set of input patterns, each of which is mapped onto a predefined subset of MEA electrodes for stimulation. For every input sample, electrical pulses are applied through the designated electrodes and repeated $N = 20$ times to generate a distribution of responses amenable to statistical analysis. This repetition is essential because the biological networks exhibit spontaneous activity [49] and are subject to intrinsic noise and varying dynamical states at the moment of stimulation. As a result, the same input pattern elicits a range of different responses across trials. Once all feature vectors have been collected, a linear classifier is trained to associate each high-dimensional representation with its corresponding input label. For every class, the recorded responses are randomly shuffled and split into disjoint sets for training and testing. Training is conducted using a single-layer perception (SLP) optimized via stochastic gradient descent (SGD), with the objective of minimizing the cross-entropy loss [27]. No validation set is used, as no early stopping criterion is applied;

this choice allows for maximal utilization of the available data for both training and evaluation. The model is trained for 1000 epochs.

To assess the generalization performance of the trained model, we employ a 5-fold cross-validation strategy. In each fold, the dataset is randomly partitioned into five equally sized subsets. Four subsets are used for training the linear classifier, while the remaining one is held out for testing. This process is repeated five times, such that each subset serves exactly once as the test set. The classifier remains fixed at the state reached at the end of training in each fold, with no additional updates or fine-tuning performed during testing.

Since the testing pipeline replicates the training setup, the statistical properties of the test features closely match those of the training features, ensuring a consistent and unbiased evaluation. For each test sample, its latent representation is fed into the trained classifier, and the resulting prediction is compared with the corresponding ground truth label. Performance is quantified using classification accuracy, defined as the ratio of correctly predicted samples to the total number of test samples. This metric provides a robust measure of the system’s ability to serve as a biologically grounded reservoir computing architecture.

4. Experimental Evaluation

This section presents the experimental setup used to evaluate our Biological Reservoir Computing (BRC) system and discusses the results obtained using a custom digit dataset. Specifically, we assess the classification performance of the latent representations generated by the biological reservoir through three biological replicates of the same neural culture. The stimulus patterns are constructed following an *electronic clock* layout. This approach emulates the segments of a digital display, where each active electrode pair represents a lit LED segment forming part of a digit. Fig. 3 provides a visual overview of the stimuli employed and Fig. 4 provides two examples of the recorded responses, computed over a time window of $W = 10$ ms.

4.1. Artificial Reservoir Comparison

To establish a reference point and enable direct comparison, we implemented an artificial reservoir (AR) based on a network consisting of 4,096 recurrent units with a connection sparsity of 10%—i.e., only 10% of all possible interconnections between units were active. The AR was initialized in a resting state, with all unit activations set to zero, and driven by the same input stimuli used in the biological experiments to ensure comparability of conditions.

To replicate the variability inherent in spontaneous neural activity, we incorporated biologically realistic input noise into the AR model. This noise was empirically estimated by analyzing $N = 100$ randomly selected time

windows of spontaneous activity, with a temporal duration of W , recorded from the biological network. The average spike count computed from these segments was used to generate synthetic noise, which was superimposed on the input stimuli before being passed to the AR. After the presentation of each stimulus, the AR dynamics were allowed to evolve for one additional time step—excluding a square region around the stimulation units, in line with the masking procedure applied in the biological experiments—was used as input to a single-layer perceptron for classification.

While the AR model serves as a valuable performance benchmark and is expected to yield superior results due to its engineered structure, it does not possess the physical or energetic characteristics of a real biological system. Consequently, the AR provides an upper bound for performance evaluation, whereas the BRC system offers insights into biologically grounded computation and energy-efficient information processing.

4.2. Quantitative Results

To evaluate the performance of the BRC system, it is first essential to determine the most informative time window for reading out the network’s response to stimulation. For this purpose, we analyzed the classification accuracy across different stimulation sessions while systematically varying the duration W of the post-stimulus time window used to extract neural activity.

Fig. 5 shows the results obtained by aggregating data from $n = 9$ independent stimulation sessions performed across multiple days and biological replicates. A clear decreasing trend emerges, with the highest classification accuracy consistently observed when $W = 5$ ms. This suggests that the most informative response — in terms of discriminability among different stimuli — occurs within the first few milliseconds following stimulation.

This result is expected, as the initial segment of the neural response primarily reflects the immediate, first-order activation of neurons directly connected to the stimulated electrodes. In biological terms, a 5 ms window roughly corresponds to the timescale of a single synaptic transmission, capturing the direct postsynaptic potentials triggered by the stimulus. At these early latencies, the signal is minimally affected by recurrent processing, spontaneous activity, or noise propagation, making it highly specific to the input pattern. As the time window W increases beyond this point, the recorded activity becomes increasingly influenced by indirect responses, recurrent network dynamics, and spontaneous background activity.

The stimulus amplitude was set to $4\ \mu\text{A}$ per electrode pair. All digit patterns were presented within the same MEA region, selected for its maximal spontaneous activity and signal quality. The average classification accuracy per category computed over the three stimulation sessions (Day 1,

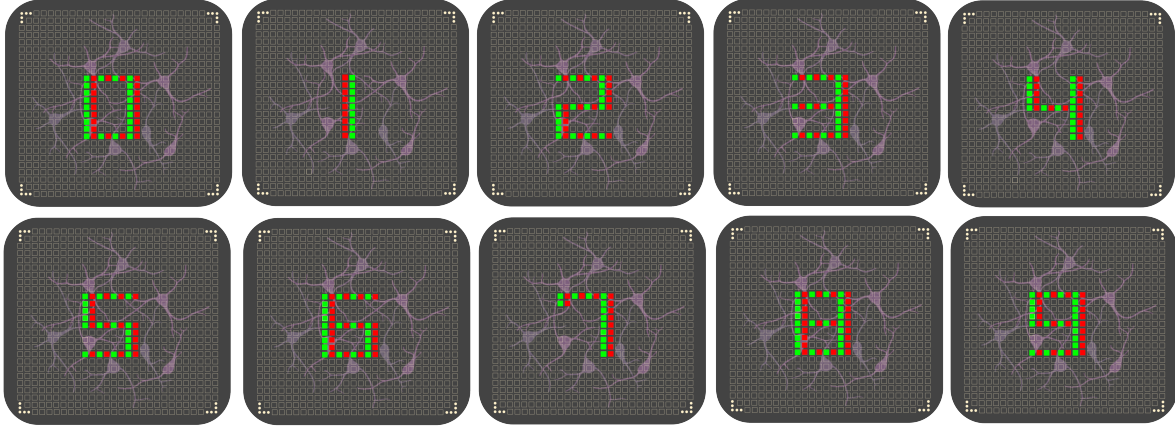


Figure 3. **Visual representation of the input patterns.** Digit recognition: input patterns represent the digits from 0 to 9.




Figure 4. **Heatmaps of neural activity following stimulation:** The left panel shows the recorded response within a 10 ms time window following stimulation with the input pattern "0", while the right panel depicts the corresponding response for the input pattern "1". The heatmaps illustrate the spatial distribution of spiking activity across the MEA electrodes.

Day 2, Day 3), and the comparison with the average classification accuracy of the artificial reservoir (AR) computed over different realizations of the noise are summarized in Fig. 6. The mean classification accuracy during each stimulation session across days of stimulation and biological replicates is reported in Tab. 1.

Interestingly, all biological reservoirs achieved performance levels comparable to those of the artificial reservoir, with an average classification accuracy of approximately 75%. This consistency was observed across multiple stimulation sessions and among the different biological replicates, underscoring the reliability and reproducibility of the system. Notably, our biological reservoir benchmark capitalizes on the inherent advantages of using a living neural substrate, such as rich, nonlinear dynamics and natural variability, which may offer valuable insights and potential benefits over conventional artificial models.




Figure 5. **Accuracy variation across different readout windows.** For each stimulation session ($n=9$), classification accuracy was computed using neural responses extracted from different post-stimulus time windows, ranging from 5 to 50 ms. The plot reports the mean accuracy \pm standard error of the mean (SEM). This analysis highlights how the temporal integration window influences the effectiveness of the reservoir readout.

4.3. Accuracy Variation Across Days

In order to assess the stability and generalizability of stimulus-evoked responses over time, we evaluated the classification performance of the classifier trained on data from the first stimulation session and subsequently tested it on data collected in the following days. This cross-day testing approach allows us to probe whether the neural dynamics elicited by a specific set of stimuli remain consistent across days, or if they are subject to spontaneous drift and reorganization.

It is well established that neural cultures evolve dynam-




Figure 6. **Mean classification accuracy per input category for each biological replicate.** Each box displays the average classification accuracy for a given stimulus category, computed over multiple stimulation sessions for each single biological replicate (BR 1, BR 2, BR 3). Accuracies are calculated using 5-fold cross-validation within each session and then averaged across the three stimulation sessions (Day 1, Day 2, Day 3). The last row in the heatmap displays the average classification accuracy per category computed over different realizations of the noise for the Artificial Reservoir (AR).




Figure 7. **Mean classification accuracy per input category across days.** Each box represents the average classification accuracy for a specific stimulus category, computed by training the classifier on data from Day 1 and testing on subsequent stimulation sessions (Day 2 and Day 3). Accuracies were averaged across three independent biological replicates, with each replicate evaluated using five cross-validation folds. Error bars denote the standard deviation across replicates, reflecting variability due to biological and experimental factors. The last row displays the average accuracy per category obtained for the classifier trained on Day 1 and tested on subsequent stimulation sessions by randomizing the neuronal responses for each stimulus, i.e., maintaining the absolute response but shuffling the spatial structure of the spiking activity recorded by the electrodes.

ically over time, spontaneously transitioning through various activity regimes—from uncoordinated, random spiking to highly synchronized bursting patterns [20, 49]. Such spontaneous reorganization is driven by intrinsic developmental processes and changes in connectivity, and it can significantly influence how the network responds to external stimulation. As a consequence, the evoked responses to identical stimulation patterns may vary substantially across days, reflecting underlying shifts in the network’s functional state.

Then, we used all the neuronal responses to the different stimulus patterns recorded during the first stimulation session (Day 1) to train a linear classifier. The trained model was then evaluated on the data acquired during the subsequent independent stimulation sessions (Day 2 and Day 3), allowing us to assess the temporal stability and generalizability of the biological reservoir’s encoding capabilities. The average classification accuracy per category, computed across the three biological replicates, is summarized in Fig. 7. As an additional control, we assessed the clas-

Stim. day	Biological Replicates(%)			Mean±SD(%)
	BR 1	BR 2	BR 3	
Day 1	77 ± 5	67 ± 9	78 ± 6	74 ± 6
Day 2	78 ± 5	66 ± 9	73 ± 4	72 ± 6
Day 3	72 ± 6	68 ± 9	67 ± 4	69 ± 3

Table 1. **Classification accuracy across stimulation day and biological replicates.** Results are reported for each stimulation session (Day 1, Day 2, Day 3) and for the three biological replicates (BR 1, BR 2, BR 3). Each value indicates the mean classification accuracy and standard deviation across 5-fold validation. The final column shows the average performance across replicates per day.

sifier’s performance on a surrogate dataset in which each neuronal response from the test days was randomized across recording channels. More in detail, we randomized the neuronal responses for each stimulus, maintaining the absolute response but shuffling the spatial structure of the spiking activity recorded by the electrodes. As expected, the classifier failed to discriminate among the stimulus classes under these conditions, yielding a mean accuracy of $12 \pm 5\%$, which is consistent with the chance level expected for a 10-class problem. The results revealed a progressive decrease in classification performance across days, with the accuracy dropping to $46 \pm 15\%$ on Day 2 and further to $37 \pm 19\%$ on Day 3. Although these values remain above chance level (10%), they indicate a substantial decline in the network’s ability to preserve stimulus-specific encoding over time. This performance degradation suggests that the underlying neural dynamics and connectivity patterns within the biological reservoir evolve significantly across days, likely due to spontaneous activity and plasticity phenomena intrinsic to neuronal cultures.

Interestingly, two stimulus categories—specifically, patterns "1" and "7"—consistently maintained higher classification accuracy across all sessions. While the precise reason for this robustness remains unclear, one possible explanation is that these patterns involve a lower number of stimulation electrodes and exhibit minimal spatial overlap with other stimuli. Such properties may lead to more distinct and less noisy evoked responses, facilitating their identification. However, further investigation is required to rigorously assess the relationship between stimulus complexity, spatial interference, and classification stability.

5. Conclusion

In this work, we proposed a novel approach to reservoir computing (RC) by leveraging a network of cultured biological neurons as the computational substrate. Unlike traditional artificial reservoirs, our bio-hybrid system exploits

the intrinsic dynamics of living neurons, which operate with remarkable energy efficiency and rich nonlinear behavior. These properties make biological reservoirs a compelling alternative for neuromorphic computing, particularly in scenarios where low-power, adaptive computation is desirable. Furthermore, this line of research contributes to a deeper understanding of how biological neural circuits process information, offering insights into the mechanisms underlying cognition and neural computation.

To validate our approach, we designed and implemented a custom stimulation and recording protocol using high-density multi-electrode array (HD-MEA) technology. Our experiments focused on a digit classification task involving 10 input categories (digits 0 through 9), each mapped to a specific spatial pattern of electrical stimulation. The recorded neuronal responses were used to train a simple linear readout layer, and the classification results demonstrated that the biological reservoir was capable of generating high-dimensional feature embeddings sufficient to support accurate pattern recognition. These findings provide strong evidence that cultured neural networks can be repurposed as effective computational modules for RC applications, even in non-trivial tasks such as handwritten digit classification.

While the current study focused on a single benchmark dataset to assess system performance, future work will aim to generalize these findings by testing the BRC system on more diverse and complex input patterns. This will help evaluate the scalability and generalization capability of the biological reservoir.

Overall, this work lays the foundation for a new class of biologically grounded reservoir computing architectures. By uniting experimental neuroscience with machine learning principles, we move closer to realizing energy-efficient, adaptive systems that bridge the gap between artificial and biological intelligence.

Acknowledgements

This work has been supported by: Matteo Caleo Foundation; Scuola Normale Superiore (FC); PRIN project "AICult" (grant #2022M95RC7) from the Italian Ministry of University and Research (MUR) (FC); PNRR project "Tuscany Health Ecosystem - THE" (CUP B83C22003930001) funded by the European Union - NextGenerationEU.

References

- [1] L. F. Abbott and Carl van Vreeswijk. Asynchronous states in networks of pulse-coupled oscillators. *Phys. Rev. E*, 48: 1483–1490, 1993. 2
- [2] Ahmed Badar, Arnav Varma, Adrian Staniec, Mahmoud Gamal, Omar Magdy, Haris Iqbal, Elahe Arani, and Bahram Zonooz. 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*, pages 515–531. Springer, 2021. 1
- [3] Filippo Maria Bianchi, Simone Scardapane, Sigurd Løkse, and Robert Jenssen. Reservoir computing approaches for representation and classification of multivariate time series. *IEEE Trans. Neural Networks Learn. Syst.*, 32(5):2169–2179, 2021. 1
- [4] P. Bonifazi and P. Fromherz. Silicon chip for electronic communication between nerve cells by non-invasive interfacing and analog–digital processing. *Advanced Materials*, 14(17):1190–1193, 2002. 1
- [5] Hongwei Cai, Zheng Ao, Chunhui Tian, Zhuhao Wu, Hongcheng Liu, Jason Tchieu, Mingxia Gu, Ken Mackie, and Feng Guo. Brain organoid reservoir computing for artificial intelligence. *Nature Electronics*, 6(12):1032–1039, 2023. 2
- [6] Stuart M Chambers, Christopher A Fasano, Eirini P Papetrou, Mark Tomishima, Michel Sadelain, and Lorenz Studer. Highly efficient neural conversion of human es and ips cells by dual inhibition of smad signaling. *Nature Biotechnology*, 27(3):275–280, 2009. 1
- [7] Ilaria Chiaradia and Madeline A Lancaster. Brain organoids for the study of human neurobiology at the interface of in vitro and in vivo. *Nature neuroscience*, 23(12):1496–1508, 2020. 1, 3
- [8] Luca Ciampi, Gabriele Lagani, Giuseppe Amato, and Fabrizio Falchi. Biologically-inspired semi-supervised semantic segmentation for biomedical imaging. *CoRR*, abs/2412.03192, 2024. 2
- [9] Luca Ciampi, Gabriele Lagani, Giuseppe Amato, and Fabrizio Falchi. A biologically-inspired approach to biomedical image segmentation. In *Computer Vision - ECCV 2024 Workshops - Milan, Italy, September 29-October 4, 2024, Proceedings, Part XIV*, pages 158–171. Springer, 2024. 2
- [10] Andrea Cossu, Davide Bacciu, Antonio Carta, Claudio Gallicchio, and Vincenzo Lomonaco. 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. 2
- [11] Peter U. Diehl and Matthew Cook. Unsupervised learning of digit recognition using spike-timing-dependent plasticity. *Frontiers Comput. Neurosci.*, 9:99, 2015. 2
- [12] Daniel Durstewitz and Gustavo Deco. Computational significance of transient dynamics in cortical networks. *European Journal of Neuroscience*, 27(1):217–227, 2008. 1
- [13] José Manuel Ferrández, Victor Lorente, Félix de la Paz, and Eduardo Fernández. Training biological neural cultures: Towards hebbian learning. *Neurocomputing*, 114:3–8, 2013. 2
- [14] Paul Ferré, Franck Mamalet, and Simon J. Thorpe. Unsupervised feature learning with winner-takes-all based STDP. *Frontiers Comput. Neurosci.*, 12:24, 2018. 2
- [15] Claudio Gallicchio and Alessio Micheli. Deep echo state network (deepesn): A brief survey. *CoRR*, abs/1712.04323, 2017. 1, 2
- [16] Wulfram Gerstner and Werner M. Kistler. *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, 2002. 2
- [17] Anubhuti Goel and Dean V. Buonomano. Temporal interval learning in cortical cultures is encoded in intrinsic network dynamics. *Neuron*, 91(2):320–327, 2016. 2
- [18] Demis Hassabis, Dharshan Kumaran, Christopher Summerfield, and Matthew Botvinick. Neuroscience-inspired artificial intelligence. *Neuron*, 95(2):245–258, 2017. 2
- [19] Ludovico Iannello, Luca Ciampi, Gabriele Lagani, Fabrizio Tonelli, Eleonora Crocco, Lucio Maria Calcagnile, Angelo Di Garbo, Federico Cremisi, and Giuseppe Amato. From neurons to computation: Biological reservoir computing for pattern recognition. *arXiv preprint arXiv:2505.03510*, 2025. 2
- [20] Ludovico Iannello, Fabrizio Tonelli, Federico Cremisi, Lucio Maria Calcagnile, Riccardo Mannella, Giuseppe Amato, and Angelo Di Garbo. Criticality in neural cultures: Insights into memory and connectivity in entorhinal-hippocampal networks. *Chaos, Solitons & Fractals*, 194:116184, 2025. 7
- [21] Bernd Illing, Wulfram Gerstner, and Johanni Brea. Biologically plausible deep learning - but how far can we go with shallow networks? *Neural Networks*, 118:90–101, 2019. 2
- [22] Takuya Isomura, Kiyoshi Kotani, and Yasuhiko Jimbo. Cultured cortical neurons can perform blind source separation according to the free-energy principle. *PLoS Comput. Biol.*, 11(12), 2015. 2
- [23] Herbert Jaeger. 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. 1, 2
- [24] Fahad Javed, Qing He, Lance E Davidson, John C Thornton, Jeanine Albu, Lawrence Boxt, Norman Krasnow, Marinos Elia, Patrick Kang, Stanley Heshka, and Dymna Gallagher. 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. 1
- [25] Adrien Journé, Hector Garcia Rodriguez, Qinghai Guo, and Timoleon Moraitis. Hebbian deep learning without feedback. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. 2
- [26] Brett J. Kagan, Andy C. Kitchen, Nhi T. Tran, Forough Habbollahi, Moein Khajehnejad, Bradyn J. Parker, Anjali Bhat, Ben Rollo, Adeel Razi, and Karl J. Friston. In vitro neurons learn and exhibit sentience when embodied in a simulated game-world. *Neuron*, 110(23):3952–3969.e8, 2022. 2
- [27] Douglas Kline and Victor L. Berardi. Revisiting squared-error and cross-entropy functions for training neural network classifiers. *Neural Comput. Appl.*, 14(4):310–318, 2005. 4
- [28] Dmitry Krotov and John J. Hopfield. Unsupervised learning by competing hidden units. *Proc. Natl. Acad. Sci. USA*, 116(16):7723–7731, 2019. 2
- [29] Gabriele Lagani, Fabrizio Falchi, Claudio Gennaro, and Giuseppe Amato. 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*, pages 25–40. Springer, 2021. 2
- [30] Gabriele Lagani, Raffaele Mazzitti, Fabrizio Falchi, Claudio Gennaro, Guido Marco Cicchini, Tommaso Pizzorusso, Federico Cremisi, and Giuseppe Amato. 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*, pages 726–729. IEEE, 2021. 2
- [31] Gabriele Lagani, Davide Bacciu, Claudio Gallicchio, Fabrizio Falchi, Claudio Gennaro, and Giuseppe Amato. 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*, pages 136–141. ACM, 2022. 2
- [32] Gabriele Lagani, Fabrizio Falchi, Claudio Gennaro, and Giuseppe Amato. Comparing the performance of hebbian against backpropagation learning using convolutional neural networks. *Neural Comput. Appl.*, 34(8):6503–6519, 2022.
- [33] Gabriele Lagani, Claudio Gennaro, Hannes Fassold, and Giuseppe Amato. 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*, pages 251–264. Springer, 2022. 2
- [34] Gabriele Lagani, Fabrizio Falchi, Claudio Gennaro, and Giuseppe Amato. Spiking neural networks and bio-inspired supervised deep learning: A survey. *CoRR*, abs/2307.16235, 2023. 2
- [35] Gabriele Lagani, Fabrizio Falchi, Claudio Gennaro, and Giuseppe Amato. Synaptic plasticity models and bio-inspired unsupervised deep learning: A survey. *CoRR*, abs/2307.16236, 2023. 2
- [36] Gabriele Lagani, Fabrizio Falchi, Claudio Gennaro, Hannes Fassold, and Giuseppe Amato. Scalable bio-inspired training of deep neural networks with fasthebb. *Neurocomputing*, 595:127867, 2024. 2
- [37] Mantas Lukosevicius and Herbert Jaeger. Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.*, 3(3):127–149, 2009. 1, 2
- [38] Wolfgang Maass, Thomas Natschläger, and Henry Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Comput.*, 14(11):2531–2560, 2002. 1, 2
- [39] Adam H. Marblestone, Greg Wayne, and Konrad P. Körding. Toward an integration of deep learning and neuroscience. *Frontiers Comput. Neurosci.*, 10:94, 2016. 2
- [40] Vito Paolo Pastore, Paolo Massobrio, Aleksandar Godjovski, and Sergio Martinoia. Identification of excitatory-inhibitory links and network topology in large-scale neuronal assemblies from multi-electrode recordings. *PLoS Comput. Biol.*, 14(8), 2018. 2
- [41] Blake A Richards, Timothy P Lillicrap, Philippe Beaudoin, Yoshua Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, Rui Ponte Costa, Archy de Berker, Surya Ganguli, Colleen J Gillon, Danijar Hafner, Adam Kepecs, Nikolaus Kriegeskorte, Peter Latham, Grace W Lindsay, Kenneth D Miller, Richard Naud, Christopher C Pack, Panayiota Poirazi, Pieter Roelfsema, João Sacramento, Andrew Saxe, Benjamin Scellier, Anna C Schapiro, Walter Senn, Greg Wayne, Daniel Yamins, Friedemann Zenke, Joel Zylberberg, Denis Therien, and Konrad P Körding. A deep learning framework for neuroscience. *Nature Neuroscience*, 22(11):1761–1770, 2019. 2
- [42] Maria Elisabetta Ruaro, Paolo Bonifazi, and Vincent Torre. Toward the neurocomputer: image processing and pattern recognition with neuronal cultures. *IEEE Trans. Biomed. Eng.*, 52(3):371–383, 2005. 2, 3, 4
- [43] Daniele Di Sarli, Claudio Gallicchio, and Alessio Micheli. 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*, pages 1–5. IEEE, 2020. 2
- [44] Simone Scardapane, John B. Butcher, Filippo Maria Bianchi, and Zeeshan Khawar Malik. Advances in biologically inspired reservoir computing. *Cogn. Comput.*, 9(3):295–296, 2017. 2
- [45] Goded Shahaf and Shimon Marom. Learning in networks of cortical neurons. *Journal of Neuroscience*, 21(22):8782–8788, 2001. 2
- [46] Sen Song, Kenneth D. Miller, and L. F. Abbott. Competitive hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*, 3(9):919–926, 2000. 2
- [47] Zhe Sun, Vassilis Cutsuridis, Cesar F. Caiafa, and Jordi Solé-Casals. Brain simulation and spiking neural networks. *Cogn. Comput.*, 15(4):1103–1105, 2023. 2
- [48] Josh Tenenbaum. Building machines that learn and think like people. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, page 5. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM, 2018. 2
- [49] Fabrizio Tonelli, Ludovico Iannello, Stefano Gustincich, Angelo Di Garbo, Luca Pandolfini, and Federico Cremisi. Dual inhibition of mapk/erk and bmp signaling induces entorhinal-like identity in mouse esc-derived pallial progenitors. *Stem Cell Reports*, 0(0), 2024. 4, 7
- [50] Qian Wang and Peng Li. 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*, pages 2652–2657. IEEE, 2016. 2
- [51] Yoshihiro Yonemura and Yuichi Katori. Dynamical predictive coding with reservoir computing performs noise-robust multi-sensory speech recognition. *Frontiers Comput. Neurosci.*, 18, 2024. 1
- [52] Yong Zhang, Peng Li, Yingyezhe Jin, and Yoonsuck Choe. 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. 1, 2