A Theoretical and Experimental Approach to Understanding Neuronal Network Computational Abilities for Deployed Machine Learning

Motivation: **According to the National Institute of Standards and Technology (NIST), the human brain has an equivalent computational capacity as a 20 MW, 300 ton supercomputer [1].** Neurons achieve remarkable computational efficiency, learning from few examples while fusing information across modalities. However, our understanding of how neuronal networks achieve this performance is incomplete, limiting our ability to harness their efficiency.

Therefore, I propose analyzing the suitability of neuronal networks within a known scheme, reservoir computing, through a combined theoretical and experimental methodology. RCs provide a mathematically grounded way to analyze nonlinear, recurrent dynamics, similar to those in biological networks. This project will clarify the computational mechanisms that govern neuronal networks while directly connecting these findings to defense-relevant computing needs that highlight how the brain's efficiency can be utilized in next-generation neuromorphic hardware.

This project aligns with the Army Research Lab (ARL)'s interests shown in ARL-BAA-0016, "Neurophysiology of Cognition" and the Department of Defense (DoD)'s long-term goals for next-generation compute systems with low size, weight, and power (SWaP) requirements.

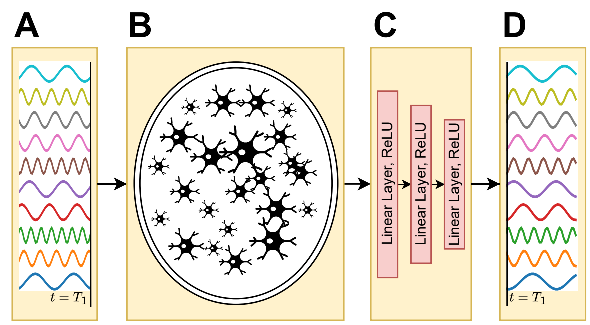

Figure 1: Overview of proposed neuronal network RC. **A**: multimodal dynamic input is recorded. **B**: The input is shown to living neurons, using either optogentics or multi-electrode arrays. The neurons pass the inputs between themselves, creating novel nonlinear representations and combining them with past inputs. **C**: A simple to train output layer captures the correct combinations of nonlinearities to create extrapolations (**D**).

Reservoir Computing and Biological Networks: Following reservoir computers' (RC) introduction as neuroscience models, the two fields continue to be intricately linked. Authors have simulated neurons as RCs, like Yamazaki and Tanaka and Mijalkov et al., while Sumi et al. demonstrated that living neurons display reservoir dynamics in-vitro, establishing feasibility for neuronal RCs [2–4]. This proposal builds directly on Sumi et al.'s living-neuron reservoir while extending it through a theoretical and experimental framework.

Theoretical Frameworks for RCs: The novelty of using living neuronal networks as reservoirs means that there is no comprehensive framework for understanding of their theoretical or realized performance. Current theories of RCs focus on three theoretical goals: RCs as universal function approximators, the memory and computational capacity of an RC, and probably approximately correct (PAC) learnability within RCs.

Prior work in Gonon, Gonon and Ortega, and Monzani and Prati has shown that all popular digital reservoir paradigms, including echo state networks, the most common type of RC proposed by Jaeger, are universal function approximators [5–8].

Memory and computational capacity characterize the temporal and nonlinear limits of any RC system. Proposed by Jaeger, memory capacity is a measure of how much past input information a reservoir can preserve [9]. Dambre et al. defines computational capacity, a measure of how much of the function space the reservoir spans [10].

PAC learnability is a theoretical approach to understanding a statistical learning algorithm. Gonon et al. uses PAC learning to define risk bounds for both generic reservoirs and random reservoirs, giving a bound on the minimum number of samples required to achieve a set generalization error [11].

Finally, researchers have studied the Vapnik–Chervonenkis (VC) dimensions of RC schemes. VC dimensions are tightly linked to PAC learnability, since finite VC dimensions guarantee learnability. Maass has argued that some reservoir have strong generalization capabilities because of especially small VC dimensions [12].

These existing theoretical guarantees for digital reservoirs rest on assumptions about dynamics, stability, and representation that have never been evaluated in living neuronal reservoirs. A core contribution of my work is to systematically test these assumptions, establishing which theoretical properties hold or break. This will help reveal the computational principles underlying biological networks.

Experimental Protocols for Novel Computing Systems: Theoretical bounds alone cannot establish the realizable performance of neuronal reservoirs. Therefore, I will also consider the rich literature on benchmarking RCs and novel computing paradigms.

Works like Maksymov, and Watt and Kostylev unify theoretical and experimental methodologies by calculating the memory capacity of unique physical reservoirs [13, 14]. Wringe et al. proposes a survey of benchmarks that can be applied to both digitial and physical RCs [15].

Energy is a critical factor for DoD edge computing, making neuron-level efficiency a key evaluation dimension. Belostoski et al. and Houshmand et al. propose frameworks for measuring how analog accelerators differ in energy usage, latency, and precision to their classical counterparts [16, 17]. Finally, my ongoing work, described by Regli et al. (in progress), proposes the energy-to-solution (ETS) metric for analyzing the energy used by an analog solver to find a solution to a single satisfiability (SAT) problem [18]. This metric includes an essential feature: in addition to energy used by the solver, it tracks the energy used by the classical I/O pathway.

Proposed Research Direction: I will integrate theoretical results and experimental methodologies to build a complete understanding of neuronal RCs' theoretical and real performance. First, it will be necessary to apply the accepted theories from reservoir computing to RCs with neuronal reservoirs.

These theories make key assumptions to prove the characteristics of digital reservoirs. **Therefore, I will test whether these theories can apply to neuronal RCs by verifying their assumptions are still valid for living reservoirs.** These assumptions, and how I plan to address them, are

- *Fading memory*: The fading memory assumption states that the system's memory fades over time. I plan to demonstrate this by measuring the reservoir's response to inputs u and v, where $u_{1:T} = v_{1:T}$ but $u_0 \neq v_0$, and showing that the response to the different initial input decreases over time.
- *Memory and computational capacity*: These measures can be verified simply because Dambre et al. proposes a model independent methodology for calculating them by observing the RC stimulated with random inputs [10].
- *Establishing a state map*: I can assess whether a state map, or function that assigns values to inputs, approximately exists by showing repeated input patterns to the reservoir, and showing that there is some F that, given inputs u, v, $u_{t-K:t} \neq v_{t-K:t} \implies F_t(u) \neq F_t(v)$.
- *Activation boundedness*: Activation boundedness requires that, for any input u, the reservoir's output is bounded with $F_t(u) < \infty$. Activation boundedness follows directly from the biophysical constraints of living cells.
- *Small-perturbation stability*: I will evaluate whether neuronal reservoirs have small-perturbation stability, or stability to small differences in the input, by showing that, for responses $F_t(u)$ to inputs u, v, where $||u-v|| \leq \delta$, $||F_t(u) - F_t(v)|| < \varepsilon$ for sufficiently small δ and ε.

These assumptions can be evaluated by observing the neuronal RCs under different input schemes, but some tests, such as establishing a state map, are nontrivial due to biological nonstationarity and noise. Neurons are measured using multi-electrode arrays or calcium imaging, which are coarse in the spatial or time domain, respectively, and neurons are inherently noisy, spiking unprompted or generating large "global bursts." Furthermore, when presented with any inputs, neurons adapt, potentially violating classical RC assumptions [19]. If any assumptions are not validated, I will revise RC learning theory to accurately reflect the living reservoirs.

Integration of Theoretical and Experimental Results: To ground my theoretical conclusions, I will pair them with experimental benchmarking, showing how neuronal RCs perform on real world tasks, and assess their usability. Firstly, I plan to conduct the benchmarks provided by Wringe et al., which test the RCs' computational abilities, and I will compare the results to the theoretical conclusions [15]. I will also reference Maksymov, and Watt and Kostylev to calculate the realized memory capacity [13, 14]. These experiments are an essential chance for me to demonstrate whether the provided assumptions are enough for generalization of the theoretical results.

Energy Usage as a Measure of Usefulness: Quantifying whether the neuronal reservoirs have lower size, weight and power (SWaP) requirements than their digital counterparts is the final step towards demonstrating their usefulness. Therefore, I will build on my ongoing work to map ETS to RCs. I propose a new measure, energy-to-generalization-level (EGL), relating the energy required by the RC to the RC's error on real-world datasets. This directly links energy to operational usefulness, which is key to demonstrating their potential SWaP advantages.

Timeline: **In year 1, I will establish a reproducible protocol for creating living reservoirs** by replicating and validating Sumi et al.'s work [4]. I will also quantify reservoir variability, a key challenge when using living cells.

In year 2, I will characterize their theoretic potential by stimulating the neurons with random or almost-random inputs and analyzing their responses against the theoretical assumptions of digital reservoirs. **In year 3, I will apply the RCs to DoD-relevant real-world datasets**, using metrics such as EGL to assess their robustness and operational value for edge-deployable neuromorphic systems. Deliverables include a validated reservoir-creation protocol (Year 1), theoretical reservoir potential, including capability measures (Year 2), and EGL benchmarks on DoD-relevant baselines (Year 3).

I work within a multidisciplinary lab with established protocols for culturing active and responsive neurons, making rigorous replication of Sumi et al. feasible [4]. The lab maintains the required stimulation and recording technology, and the analysis tools for the proposed measurements. My research centers on understanding computation in dynamical systems, ranging from concept drift detection to analog SAT solvers to neuromorphic computing, making this project a good fit for my skillset.

Impact and Contributions: Neurons pack incredible computational capabilities into our brain with low SWaP demands. This project quantifies and characterizes that advantage from a theoretical and experimental perspective. *Scientific Contributions:* **This work advances and formalizes our understanding of how neurons learn, store, and process information within a controlled computational framework.** By comparing their performance to understood schemes, such as ESN, we will be able to verify or correct assumptions around their processing. The insights from this project will refine our theoretical models of neuronal computation, ultimately realizing next-generation computing schemes that replicate the brain's computational efficiency.

Finally, this project extends our understanding of reservoir models. Using experimental and theoretical results, it will extend the framework of reservoir computing to reservoirs whose internal mechanisms are not fully understood or are inherently adaptive, such as living neuronal networks. Current theories of RCs largely assume stable, noise-limited systems with fixed connectivity, but real neurons are noisy and continuously adapt their connections and activity patterns in unknowable ways. These results will help clarify the principles that allow both neuronal network and digital reservoirs to perform efficient, low-sample learning, and will help identify what structural and dynamic properties make a reservoir computationally powerful.

Relevance to DoD Computational Goals: **This project directly aligns with the Army Research Lab's interests outlined by ARL-BAA-0016, "Neurophysiology of Cognition." The proposal's focus on characterizing and benchmarking living neuronal networks as computational reservoirs directly supports the ARL's Neural Computation thrust by uncovering the mechanisms that make the brain an incredibly efficient and powerful computer.** By developing theoretical frameworks for understanding neuronal network reservoirs, including PAC learnability, VC dimensions, and memory capacity, this work contributes to ARL's goal of establishing mathematical and computational models for closed-loop prediction and control of neural dynamics. The proposed experiments further link biological signals, such as neuronal spikes, to computational outputs, creating a bridge between neuronal networks and digital computation. Finally, by integrating theory and experimentation within a neuronal network RC, this project exemplifies the type of high-risk, high-reward foundational research ARL aims to support.

Furthermore, the ability to perform calculations on low SWaP devices is central to the DoD's goals outlined in the "Outside the Continental United States (OCONUS) Cloud Strategy" document [20]. In it, the DoD proposes that users in-theatre must have access to edge computing, including the ability to use and retrain ML algorithms. However, forward-deployed units that face severe SWaP constraints make the deployment of GPU-based accelerators infeasible. Therefore, this project's success could lead to the development of neuronal network RC accelerators, providing the computational efficiency of the brain to these units. Since the only trained part of an RC is the small readout model, an RC using neuronal networks could be retrained quickly with few examples for any novel scenario the unit encounters.

In summary, this research connects fundamental neuroscience to future defense applications by bringing the brain's low SWaP, efficiency, adaptability, and robustness to next-generation computing hardware. It will use theory and experimentation to show how neuronal networks can potentially serve as highly-capable, low-swap computers that give the DoD supercomputer-like intelligence without immense energy burdens.

References:

1 A. Madhavan. *Brain-Inspired Computing Can Help Us Create Faster, More Energy-Efficient Devices — If We Win the Race — nist.gov*. [Accessed 31-10-2025].

2 T. Yamazaki and S. Tanaka. “The cerebellum as a liquid state machine”. In: *Neural Networks* 20.3 (2007). Echo State Networks and Liquid State Machines, pp. 290–297.

3 M. Mijalkov et al. “Computational memory capacity predicts aging and cognitive decline”. In: *Nature Communications* 16.1 (Mar. 2025).

4 T. Sumi et al. “Biological neurons act as generalization filters in reservoir computing”. In: *Proceedings of the National Academy of Sciences* 120.25 (June 2023).

5 L. Gonon and J.-P. Ortega. “Reservoir Computing Universality With Stochastic Inputs”. In: *IEEE Transactions on Neural Networks and Learning Systems* 31.1 (2020), pp. 100–112.

6 L. Gonon and J.-P. Ortega. “Fading memory echo state networks are universal”. In: *Neural Networks* 138 (2021), pp. 10–13.

7 F. Monzani and E. Prati. *Universality conditions of unified classical and quantum reservoir computing*. 2024.

8 H. Jaeger. “The “echo state” approach to analysing and training recurrent neural networks-with an erratum note”. In: *Bonn, Germany: German national research center for information technology gmd technical report* 148.34 (2001), p. 13.

9 H. Jaeger. “Short term memory in echo state networks. gmd-report 152”. In: *GMD-German National Research Institute for Computer Science (2002)* (2002).

10 J. Dambre et al. “Information processing capacity of dynamical systems”. en. In: *Sci. Rep.* 2.1 (July 2012), p. 514.

11 L. Gonon et al. “Risk bounds for reservoir computing”. In: *Journal of Machine Learning Research* 21.240 (2020), pp. 1–61.

12 W. Maass. “Liquid State Machines: Motivation, Theory, and Applications”. In: *Computability in Context*, pp. 275–296.

13 I. S. Maksymov. “Analogue and Physical Reservoir Computing Using Water Waves: Applications in Power Engineering and Beyond”. In: *Energies* 16.14 (2023).

14 S. Watt and M. Kostylev. “Reservoir Computing Using a Spin-Wave Delay-Line Active-Ring Resonator Based on Yttrium-Iron-Garnet Film”. In: *Phys. Rev. Appl.* 13 (3 Mar. 2020), p. 034057.

15 C. Wringe et al. “Reservoir computing benchmarks: a tutorial review and critique”. In: *International Journal of Parallel, Emergent and Distributed Systems* 40.4 (2025), pp. 313–351.

16 L. Belostotski et al. “A Survey of Analog Computing for Domain-Specific Accelerators”. In: *Electronics* 14.16 (2025).

17 P. Houshmand et al. *Benchmarking and modeling of analog and digital SRAM in-memory computing architectures*. 2023.

18 W. Regli et al. *SAT Benchmarks to Assess Quantum-Inspired Solvers*. `https://github.com/UMD-ARLIS/QuICC-SAT-Datasets`. Manuscript in preparation. 2025.

19 S. Löwel and W. Singer. “Selection of Intrinsic Horizontal Connections in the Visual Cortex by Correlated Neuronal Activity”. In: *Science* 255.5041 (1992), pp. 209–212.

20 U.S. Department of Defense Chief Information Officer. *DoD OCONUS Cloud Vision*. Tech. rep. Dec. 2023.