

Quantum reservoir computing using arrays of Rydberg atoms

Rodrigo Araiza Bravo^{1,*}, Khadijeh Najafi^{1,2,†}, Xun Gao¹, and Susanne F. Yelin¹

¹*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*

²*IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, NY 10598 USA*

(Dated: May 30, 2022)

Quantum computing promises to provide machine learning with computational advantages. However, noisy intermediate-scale quantum (NISQ) devices pose engineering challenges such advantages. Recently, a series of quantum machine learning computational models inspired by the noise-tolerant dynamics on the brain have emerged as a means to circumvent the hardware limitations of NISQ devices. In this article, we introduce a quantum version of a recurrent neural network (RNN), a well-known model for neural circuits in the brain. Our quantum RNN (qRNN) makes use of the natural Hamiltonian dynamics of an ensemble of interacting spin-1/2 particles as a means for computation. In the limit where the Hamiltonian is diagonal, the qRNN recovers the dynamics of the classical version. Beyond this limit, we observe that the quantum dynamics of the qRNN provide quantum computational features that can aid in computation. To this end, we study a qRNN based on arrays of Rydberg atoms, and show that the qRNN is indeed capable of replicating the learning of several cognitive tasks such as multitasking, decision making, and long-term memory by leveraging several key features of this platform such as interatomic interactions, and quantum many-body scars.

I. INTRODUCTION

Quantum computing promises to enhance machine learning algorithms. However, implementing these advantages often relies on either fault-tolerant quantum computers not yet available [1–5], or on decoherence-limited, variational quantum circuits which may experience training bottlenecks [6, 7]. Thus, currently available noisy intermediate-scale quantum (NISQ) devices thwart quantum advantages in machine learning algorithms.

Recently, to counteract these challenges, several quantum machine learning architectures have emerged inspired by models for computation in the brain [8–10]. These brain-inspired algorithms are motivated by the inherent robustness of input- and hardware-noise in brain-like computation, and by the possibility to use the analogue dynamics of controllable, many-body quantum systems for computation without relaying on a digital circuit architecture. Broadly speaking, these brain-inspired algorithms can be put into two categories. The first of which encompasses systems quantizing the dynamics of biological computational models at the single-neuron level. Thus, the dynamics of single qubits or groups of qubits resemble the dynamics of a neurons in a neural circuit of interest. Examples of these include quantum memristors [11], which are electrical circuits with a history-dependent resistance, quantum versions of the biologically realistic Hodgkin-Huxley model for single neurons [12, 13], and unitary adiabatic quantum perceptron [14].

The second category of brain-inspired algorithms are those that rely on a macroscopic resemblance between many-body quantum systems and neural circuits. In this regard, the algorithms that has received the most

attention are quantum reservoir computers. Quantum reservoir computers use ensembles of quantum emitters with fixed interactions as means to perform versatile machine learning tasks relying on the complexity of the unitary evolution of the system. Since these systems can couple with both classical and quantum devices, which may encode the tasks' input, quantum reservoirs have been used for time-series prediction [15–17], entanglement measurement [18, 19], quantum state preparation [20], continuous-variable computation [21] which can be made universal [22], reduction of depths in quantum circuit [23], ground state finding [24], and for long-term memory employing ergodicity-breaking dynamics [25–27]. A comprehensive review of quantum reservoir computing can be found in [10].

In both categories, however, a thorough understanding of the potential computational advantages, and their origins is slowly emerging. In this article we contribute to this direction by proposing a quantum extension of a well-known neural circuit model called recurrent neural networks (RNNs), of which reservoir computers are a special case [28]. Our extension makes use of the Hamiltonian dynamics of ensembles of two-level systems, and in the limit where the Hamiltonian is diagonal, we recover the classical single-neuron dynamics naturally encoding RNNs into quantum hardware. Recently, another natural encoding of a reservoir computer has been proposed using superconducting qubits [29]. In our case, the general dynamics of the quantum RNN (qRNN) present several new features that can aid in the computation of both classical and quantum tasks. In particular, a qRNN used to simulating stochastic dynamics can exhibit speedups compared to classical RNNs.

To show that our scheme is experimentally realizable, we propose that arrays of Rydberg atoms can be used as qRNNs (Sec. IV). Although our Rydberg qRNNs have restricted connectivity, we are motivated to use Rydberg arrays due to recent studies with equally restricted

* oaraizabravo@g.harvard.edu

† knajafi@ibm.com

qRNNs which show significant computational capacity when driven near criticality [17, 24]. Moreover, recent experiments using optical tweezers [30–37] have catapulted the community’s interest in Rydberg arrays as they exhibit long coherence times, controllable and scalable geometries, and increasing levels of single-atom control [38]. Additionally, Rydberg arrays can be used for novel, programmable quantum simulations and universal computations [30, 39–43].

We numerically implement Rydberg qRNNs, and we successfully perform cognitive tasks even when a few atoms are available (Sec. V). The success of these tasks is explained by the physics of Rydberg atoms. For example, our Rydberg qRNNs excel at learning to multi-task since it can naturally encode RNNs with inhibitory and excitatory neurons which are vitals for many cognitive tasks [44]. This encoding relies on the different types of interactions between Rydberg atoms with different principal quantum numbers [45]. Likewise, a Rydberg qRNN exhibits long-term memory due to the weak-ergodicity breaking dynamics of many-body quantum scars [35, 46, 47]. Lastly, we discuss possible further research directions in Sec. VI.

We remark that the notion of qRNNs has been previously coined relying on universal quantum circuits and using measurements to implement the nonlinear dynamics of an RNN [48]. Instead, what we define as a “quantum RNN” leverages the inherent unitary dynamics of ensembles of two-level systems to compute, deviating from the quantum digital circuit model for computation.

II. CLASSICAL RECURRENT NEURAL NETWORKS

We begin by reviewing an archetypal RNN consisting of N binary neurons. Each neuron is in one of two possible states $s_n(t) \in \{-1, 1\}$ and is updated from the time-step t to $t + 1$ following the update rule

$$\begin{aligned} s_n(t+1) &= \text{sign}(h_n(t)s_n(t)), \\ h_n(t) &\equiv -\Delta_n(t) + \sum_m J_{nm}s_m(t), \end{aligned} \quad (1)$$

where $J_{nm} = J_{mn}$ are symmetric synaptic connections between neurons n and m . The time-dependent biases $\Delta_n(t)$ encode the RNN’s inputs. In order to avoid memorization during a learning task with inputs $u_n^{\text{task}}(t)$, the RNN receives Gaussian-whitened inputs

$$\Delta_n(t) = u_n^{\text{task}}(t) + \xi_n, \quad (2)$$

where ξ_n is a zero-mean Gaussian random variable with variance $\sigma_{\xi_n}^2$, making the evolution of the RNN stochastic. In RNNs, the value of $\sigma_{\xi_n}^2$ is proportional to the value of the tasks’ inputs u_n^{task} .

When studying learning tasks similar to those in the mammalian cortex [44] one turns to a continuous version

of the rule in (1) obtained in the case that the time-interval τ in which neurons are updated is small compared to J_{nm} . In this limit,

$$\tau \dot{s}_n(t) = -s_n(t) + \text{sign}(h_n(t)s_n(t)). \quad (3)$$

Thus, the RNN obeys a system of nonlinear differential equations. Note that (3) imply that $s_n \in [-1, 1]$ is a continuous and bounded variable [28].

A third way to describe an RNN is via the probability distribution $p_t(\mathbf{s})$ of observing each of the 2^N different configurations $\mathbf{s}$ at the t^{th} time-step. Due to the noise in the inputs Δ_n , the dynamics of the distribution follows a Markov chain description [28]. This description is particularly useful for analyzing the stochastic dynamics simulatable by an RNN. As we shall see in Sec III A, this representation will be useful in explaining how, relative to classical RNNs, the unitary dynamics of a qRNN can speed-up stochastic process simulations.

Lastly, we describe how an RNN is used for computation. After the RNN evolves for a time t_f , a subset of M neurons are used to collect the vector $\mathbf{r}(t_f) = (s_{n_1}(t_f), \dots, s_{n_M}(t_f), 1)$ with the last entry accommodating for a bias. The $N - M$ neurons which are not used for readout, are called *hidden neurons*. The RNN’s output is obtained via a linear transformation $\mathbf{y}^{\text{out}} = W^{\text{out}}\mathbf{r}(t_f)$ where W^{out} is a real-valued matrix. Thus, the computational complexity of the RNN comes from the nonlinear activation function in (1) which enables $\mathbf{y}^{\text{out}}$ to be a nonlinear function of the inputs.

In a learning task with a target output $\mathbf{y}^{\text{targ}}$, the RNN is trained by minimizing a loss function $\mathcal{L}(\mathbf{y}^{\text{out}}, \mathbf{y}^{\text{targ}})$ with respect to the network parameters such as W^{out} , J_{nm} , etc. subject to the task-determined inputs in (2). We choose the square-loss

$$\mathcal{L}(\mathbf{y}^{\text{out}}, \mathbf{y}^{\text{targ}}) = \frac{1}{N_s} \sum_{i=1}^{N_s} \|\mathbf{y}_i^{\text{targ}} - \mathbf{y}_i^{\text{out}}\|^2, \quad (4)$$

where i labels the N_s different input instances. For the tasks in Sec. V, we fix the connections J_{nm} , such that our qRNNs more closely resemble quantum reservoir computers.

III. QUANTUM RECURRENT NEURAL NETWORKS

A. Quantum update rule

Let us now extend the classical RNN in (1) to the quantum setting. We replace each of the N neurons by a spin-1/2 particle for which a spin measurement along the z -axis yields the values $\{-1, 1\}$. Thus, each neuron n is in a normalized quantum state in the Hilbert space $\mathcal{H}_n$ with basis vectors $\{|-1\rangle_n, |1\rangle_n\}$ which are eigenstates of the Pauli-Z operator $\sigma_n^z = |1\rangle\langle 1|_n - |-1\rangle\langle -1|_n$. The state of the composite system lives in the product Hilbert space $\mathcal{H} = \bigotimes_{n=1}^N \mathcal{H}_n$.




FIG. 2. Detection of a Z-error on three spins $L_{1,2,3}$ using a quantum RNN. A Z-error is conjugated into a bit-flip-like error using a Hamiltonian generating a rotation along the x-axis where $t = \pi/2\Omega$ and Ω is the dominant field of the Hamiltonian. The state of each of the L_i after the rotation (orange region) depends on whether a Z-error occurs, as it's illustrated at the bottom of the Fig.. As exemplified here for L_3 , a Z-error results on a spin flipping from what we would expect in the absence of errors. To detect the Z-error, a set of auxiliary qubits $A_{1,2}$ is brought in to perform a parity measurements of pairs (L_1, L_2) and (L_2, L_3) . Since under no Z-error the parity measurements must match, the parity measurements allow us to detect the location of the Z-error as specified in Table I.

freedom results in the ability to rotate each quantum neuron which can be used as means for computing in different basis. and thus revealing the quantum correlations which enhance the performance of a qRNN relative to its classical counterpart. In this section, we show how to use the qRNN's evolution to change the basis in which an error occurs. This freedom can be used to detect a Z-error, an error proper to quantum computation.

Consider the repetition code $|0_L\rangle = |-y\rangle^{\otimes 3}$ and $|1_L\rangle = |+y\rangle^{\otimes 3}$ on qubits labeled $L_{1,2,3}$ where $|\pm y\rangle = \frac{1}{\sqrt{2}}(|-1\rangle \pm i|1\rangle)$. Suppose we prepare the state $|\psi\rangle = a|0_L\rangle + b|1_L\rangle$, and consequently a Z-error occurs, we can detect the error by rotating all three spins $L_{1,2,3}$ using (5) with the dominant field being Ω for a time $t = \pi/2\Omega$. Note that the rotation conjugates the Z-error by

$$e^{-i\pi\sigma^x/4}\sigma^z e^{i\pi\sigma^x/4} \propto \sigma^y \quad (7)$$

where $\sigma_n^y = i|-1\rangle\langle 1| - i|1\rangle\langle -1|$ is like a bit-flip error except for a state-dependent phase. A bit-flip error can then be detected by bringing two extra spins $A_{1,2}$ and performing parity measurements of the pairings (L_1, L_2) , and (L_2, L_3) as described in Sec. III A. Using Table I, the final parity of (L_1, L_2) , and (L_2, L_3) gives the mea-

surement results a_1 and a_2 which can be used to discern where the Z-error occurred.

As an example, Fig. 2 illustrates the two final states of L_3 if no error occurs (bottom left), and if a Z-error occurs on L_3 (bottom right).

Detecting the Z-error hinges on (7) which can be achieved by using the qRNNs evolution to rotate the measurement basis. Note that rotation allows us to measure the error-syndrome of the stabilizer state $|\psi\rangle$, bringing out the quantum correlations of the state. Thus, the qRNN's native evolution can be used to perform quantum computational tasks. After the error is detected on spin L_i , all qubits are rotated again by $U^\dagger$ and σ_i^z can be applied to correct the error. We note that using a repetition code for error detection is a well known technique to the quantum computing community.

The previous two quantum features show that qRNNs are naturally suited to solve important problems in machine learning and quantum computing. Recently, qRNNs have been used to compress quantum circuits [23]. However, studies on using qRNNs for error correction in circuit-like quantum computing are warranted, and left for further studies.

$a_1 \backslash a_2$	-1	+1
-1	Error in L_2	Error in L_1
+1	Error in L_3	No error

TABLE I. Results of parity measurements for detection of a Z-error. Measuring spin A_i results in the outcome a_i . By comparing the outcomes, one can detect the location of the Z-error.




FIG. 3. Comparing a classical and a quantum RNN to stochastically evolve a distribution p_{t_f} from an initial distribution p_0 . Here, we consider $p_0(\mathbf{s}) = p_{t_f}(\mathbf{s}')$ when $s_n = -s'_n$ for all n . In this case, the RNN needs to produce a stochastic process matrix L^{targ} that flips all the spins through several time-steps. The classical RNN (top) requires $\mathcal{O}(2^{N-m})$ time-steps (i.e applications of P) while using m hidden neurons. A qRNN (bottom) requires one time-step and no hidden neurons.

Quantum feature 3: stochastic processes accessible to a qRNN

We now explore how a qRNN can be used to stochastically evolve a probability distribution faster than any classical RNN. Firstly, we note that if we initialize an RNN according to an initial distribution $p_0(\mathbf{s})$, the dynamics in (1) dictate that for $t > 0$ the RNN obeys a distribution given by the Markov-chain dynamics

$$p_t(\mathbf{s}) = \sum_{\mathbf{s}'} P(\mathbf{s}|\mathbf{s}') p_{t-1}(\mathbf{s}') \quad (8)$$

where $P(\mathbf{s}|\mathbf{s}')$ is the transition probability between states $\mathbf{s}'$ and $\mathbf{s}$, which particular value is given by (1) [28] (see Appendix A for details).

Given this observation, we see that an RNN can be used for the task evolving a probability distribution p_0 into $p_f = L^{\text{targ}} p_0$ by a series of stochastic transition matrices $L^{\text{out}} = P^{t_f}$. The goal is to adjust the parameters of the RNN (i.e. biases and connection weights) in order to simulate the stochastic matrix encoded in $L^{\text{out}} \approx L^{\text{targ}}$ in as few steps as possible. One may then ask if a qRNN can do this more efficiently than any RNN.

We answer this in the positive. It is worth noting that not all stochastic transition matrices L^{targ} are embeddable in a Markov process (for a review of classical and quantum embeddability see Appendix A). To simulate a stochastic system's future behavior, information about its past must be stored, and thus memory is a key resource. Quantum information processing promises a memory advantage for stochastic simulation [50]. In simulating stochastic evolution with classical resources there is a trade-off between the temporal and physical resources needed [51], and it's been shown that certain stochastic evolutions, when simulated with quantum hardware, may not suffer from such trade-off since the evolution arising from quantum Lindbladian dynamics are far more general than classical Markovian evolution [52]. That is, there exist matrices L^{targ} that are quantum embeddable but not classically embeddable. Moreover, even if L^{targ} is embeddable, the quantum evolution can lower the number of steps needed to produce L^{targ} due to the fact that the unitary dynamics of a quantum system allows a simultaneous, continuous, and coherent update of every neuron. This separation in capabilities illustrates the computational advantages from quantizing an RNN.

Let us now give an example of a matrix L^{targ} that can be achieved exponentially faster in a qRNN. Consider the task of realizing a transformation F corresponding to a global "spin-flip"

$$F_{\mathbf{s}|\mathbf{s}'} = \begin{cases} 1 & \text{if } \forall_n s_n \neq s'_n \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

Realizing F on N neurons using a classical Markov process requires a number of time steps of order $\mathcal{O}(2^{N-m})$ where m is the number of hidden neurons (for details, see

Sec. III.A in Ref. [52]). In other words, a classical RNN cannot produce F efficiently when all available neurons must be flipped. This is a result of (1), and the fact that flipping neuron n is done by ensuring that there is another neuron m in the opposite state so that $J_{nm} > 0$ dominates h_n .

On the other hand, a qRNN can perform F in a single step regardless if all neurons need to be flipped. To see this, one can consider the case of (5) with $\Omega \gg h_n$. In this case, all neurons simultaneously flip with no hidden neurons and in a single time-step evolving under a unitary U . That is, if $|\psi_0\rangle = \sum_{\mathbf{s}} \sqrt{p_0(\mathbf{s})} |\mathbf{s}\rangle$, then

$$|\psi_f\rangle = U|\psi\rangle = \sum_{\mathbf{s}} \sqrt{p_f(\mathbf{s})} |\mathbf{s}\rangle. \quad (10)$$

While the realization of the matrix F via (5) signal a quantum advantage, we highlight that this advantage is extremely sensitive to the decoherence arising from spontaneous emission (i.e. spontaneous relaxations from $|1\rangle$ to $|-1\rangle$), a main source of noise in NISQ devices (see Appendix A). It remains an open problem whether there exist stochastic processes enabled by (5) which are robust to noise, and in the future we hope to explore how to shield unitary stochastic processes against noise in experimentally realizable NISQ devices.

The spin-flip process F can be efficiently simulated using a classical computer. However, F exemplifies the qRNN's ability to access stochastic processes inaccessible to classical RNNs without hidden neurons. This implies that if an RNN is employed simulate evolving p_0 to p_{t_f} stochastically by passing it through a number of linear transformations, there are instances where the qRNN requires exponentially fewer steps. Stochastic simulation, of course, has applications in finance, biology, ecology, among other fields. As an example, Ref. [53] used this quantum advantage to propose a quantum circuit algorithm for stochastic process characterization and presented applications in finance and correlated random walks. The separation above illustrates the computational advantages from quantizing an RNN.

B. qRNNs under spontaneous emission

Having seen how (5) recovers the discrete update rule (1), we now show that a qRNN under dissipation naturally evolves under continuous-time dynamics analogous to (3). This establishes only mathematical similarities between the evolution of NISQ devices and neural circuits, allowing us to use available quantum hardware for cognitive tasks, an idea that we explore further in Sec. V.

Consider the qRNN in (5) under spontaneous emission where a spin relaxes from $|1\rangle$ to $|-1\rangle$ at a rate γ . In order to extract the dynamics of continuous variables, we focus on the dynamics of the expectation values of local Pauli operators. The expectation value of an observable A is

$\langle A \rangle = \text{Tr}(A\rho)$ where ρ is the density matrix describing the system. In particular, we focus on the expectations of the operators σ_n^x , and $\sigma_n^y = i|1\rangle\langle 1|_n - i|1\rangle\langle -1|_n$. If we start the qRNN at a state for which $\langle \sigma_n^z(0) \rangle = -1$ then (see Appendix B)

$$\begin{aligned} \langle \dot{\sigma}_n^y \rangle = & -\frac{1}{\tau} \langle \sigma_n^y(t) \rangle - \frac{\Omega}{2\gamma} \sum_m J_{nm} \langle \sigma_n^x(t) \sigma_m^y(t) \rangle \\ & + \Delta_n(t) \langle \sigma_n^x(t) \rangle, \end{aligned} \quad (11)$$

where we have defined the neural time-scale $\tau^{-1} = \gamma/2 + \Omega^2/4\gamma$ which is different than that in (3) but bears the analogous significance of the time-scale in which $\langle \sigma_n^y \rangle$ decays.

Differently than (3), notice that the dynamics of $\langle \sigma_n^y \rangle$ are influenced by the spin's value along the x -axis, a consequence of the nontrivial commutation relation of spin variables. The commutation relations also make (11) quadratic, and therefore nonlinear. The quadratic term in (11) is analogous to the nonlinear term that give RNNs their computational power.

In Appendix B, we explore the dynamics of $\langle \sigma_n^x \rangle$ as well and show that together with $\langle \sigma_n^y \rangle$, we recover dynamics analogous to integrate-and-fire RNN model [54], a more realistic model of neural networks in the brain than the one in (3).

IV. QRNNs USING RYDBERG ATOMS: AN EXPERIMENTAL PROPOSAL

The similarities between (11) and the evolution of RNNs suggest the ability of qRNNs to emulate neurological learning. To explore neurological learning in qRNNs, we propose to fix the architecture of the qRNN coupling constants J_{nm} based on optical-tweezers arrays of Rydberg atoms.

The natural Hamiltonian of a Rydberg array closely resembles the one in (5). A Rydberg atom is a single valance-electron atom that can be coherently driven between an atomic ground state $|g\rangle$ and a highly excited state $|r\rangle$ with a much larger principal quantum number. These states can represent our $|-1\rangle$ and $|1\rangle$ neuronal states respectively. A Rydberg atoms in its excited state exhibits a large electronic dipole moment, and, consequently, a collection of Rydberg atoms interact via a $1/R^6$ van der Waals potential where R denotes the physical distance between the two atoms. For an array of Rydberg atoms where the atoms are at fixed positions, the Hamiltonian of the system is [35]

$$H_{Ryd} = \Delta \sum_n \hat{n}_n + \frac{\Omega}{2} \sum_n \sigma_n^x + \sum_{nm} \frac{V}{R_{nm}^6} \hat{n}_n \hat{n}_m \quad (12)$$

where $\hat{n}_n = |1\rangle\langle 1|_n$, Ω is the coherent Rabi drive coupling the $|-1\rangle$ and $|1\rangle$ states, $\Delta < 0$ is a global drive frequency mismatch to the atomic spacing of the atoms,




FIG. 4. Schematic picture of RNNs with classical and quantum neurons. (A) Classical RNN. The inputs are local biases, and the inter-neuronal connections J_{nm} are arbitrary. A set of neurons is used for readout to produce the output $\mathbf{y}^{\text{out}} = W^{\text{out}} \mathbf{r}$. (B) qRNN made from Rydberg atoms which restrict the connections to $J_{nm} \sim 1/R_{nm}^6$ where R_{nm} is the physical distance between atoms n and m . Here we depict interactions between nearest and next-nearest neighbors, however each neuron interacts with all others in the chain via $J_{nm} \sim 1/R_{nm}^6$. Local expectation values of a subset of atoms are used for readout. (C) A schematic of the array of Rydberg atoms of the RNN in (B) using optical tweezers. Each atom experiences a Rabi-drive Ω , and a local detuning Δ_n which encodes the RNN's inputs. One of the main sources of decoherence in Rydberg atoms is spontaneous emission at a rate γ .

and V is the nearest neighbor interaction strength. Using acousto-optical deflectors (AOD) and spatial light modulator (SLM), one can create spatially depending light-shifts resulting in site and time-dependent detunings $\Delta_n(t) = \Delta + \alpha(t)\Delta_n$ where $\alpha(t)$ is a time-dependent envelope. With this in mind, the Hamiltonian in (12) can be mapped to a Hamiltonian like that in (5) with $J_{nm} = V/R_{nm}^6$ since $\hat{n}_n = (\sigma_n^z + 1)/2$. In this paper, for concreteness, we compare our numerics against the experimental realization of Rydberg arrays in Ref. [35, 38], where the rates Ω, Δ_n, V are all in units of mega-Hertz, while time constants are in units of microseconds. In these experiments, an off-resonance intermediate state, $|6P_{3/2}, F=3, M_F=-3\rangle$, is used to couple $|g\rangle = |5S_{1/2}, F=2, m_F=-2\rangle$ and $|r\rangle = |70S_{1/2}, m_J=-1/2, m_I=-3/2\rangle$ of Rubidium-87 atoms through a two-photon transition. Thus, photon-scattering off the intermediate state is the dominant source of decoherence. As we show in Appendix D, we can model this with a modified spontaneous emission process given by the jump operator

$$L^+ = \sqrt{\gamma}|g\rangle\langle r| + \beta|g\rangle\langle g| \quad (13)$$

instead of the typical $\sqrt{\gamma}|g\rangle\langle r|$ jump operator. In the equation above, $\gamma = 2\pi/(20 \mu\text{s})$, and $(\alpha, \beta) = (0.05, 0.16)$ for the realistic settings we simulate. With the full unitary and dissipative dynamics, we can think of an array of Rydberg atoms as a quantum analogue of a continuous-time RNN. Fig. 4 compares the architecture of a classical RNN in Fig. 4A, and a Rydberg RNN in Figs. 4B-C.

We note that training RNNs can be unstable as that often relies on (truncated) back-propagation through time or real time recurrent learning. One way to circumvent this problem is by keeping the system's parameters fixed, and instead only training the output filter W^{out} . This easier training schedule motivated the introduction of reservoir computers [55] and their quantum analogues [10, 15–25, 27]. Thus, in the following numerical experiments we fix the position of the atoms in either a 1D chain or a 2D square lattices and train only W^{out} and some temporal parameters depending on the task. Logically, successful performance on the tasks here presented sufficiently shows qRNNs computational ability. While we include the effect of small imperfections on the positions of the atoms, we see no significant effect on the performance of the tasks after averaging our results over 10 realizations of the atom's positions. We leave full optimization of the qRNN for future work.

Lastly, several features of the many-body dynamics of arrays of Rydberg atoms are particularly well suited for emulating biological tasks. In Sec. V A, we show how Rydberg arrays can be used to implement inhibitory and excitatory neurons which are vital in many biological tasks such as multitasking [56]. The key idea behind encoding inhibitory neurons will be leveraging positive and negative interactions between Rydberg atoms with different principal quantum numbers [45]. Additionally, in Sec. V D we show that Rydberg arrays can store long-term memory by taking advantage of the weak-ergodicity breaking dynamics of quantum many-body scars [35, 46, 47].

V. LEARNING BIOLOGICAL TASKS VIA QRNNs

We focus on analyzing the Rydberg qRNNs' potential to learn biologically plausible tasks. In the tasks analyzed, we fixed the geometry of the atoms depending on the task at hand, which makes our qRNNs into quantum reservoir computers. As a proof of principle, we focus on four simple neurological task which indicate good performance even with a small number of atoms. We show that a Rydberg qRNN can encode inhibitory and excitatory neurons vital for successful multitasking. Likewise, we show that Rydberg qRNNs can learn to make decision by distinguishing properties of stimuli, have working memory, and exhibit long-term memory enhanced by quantum many-body scars. Simulation details of each task can be found in Appendix D.




FIG. 5. Encoding inhibitory neurons using Rydberg atoms and using them for multitasking. Multitasking consists in fixing the qRNN's parameters, and training W^{out} to produce three conflicting outputs. (A) Shows the scheme for encoding inhibitory neurons. Rydberg atoms with different principal quantum numbers are used such that pairs $(n_Q)(n_Q')$ interact attractively while $(n_Q')(n_Q')$ and $(n_Q)(n_Q)$ pairs interact repulsively. The network receives two binary inputs x, y . (B) Square error for learning the functions XOR, OR, and AND on the inputs with different numbers of inhibitory neurons. Better performance is observed when 1 in every 4 neurons is inhibitory. (C)-(E) Example of learnt functions using eight neurons and two inhibitory neuron, which results on performing 40% better than without inhibitory neurons.

A. Multitasking

A hallmark of classical RNNs is their ability to multitask by learning several output functions at once. Dale's principle defines an inhibitory neuron, indexed by n , as one with a negative sign in its interactions with all other neurons [57]

$$J_{nm} \leq 0 \quad \forall m. \quad (14)$$

On the other hand, a neuron i is excitatory if $J_{nm} \geq 0$ for all m . In an RNN, it is conjectured that multitasking requires that inhibitory neurons are present at a ratio of 1:4 [56]. In this section, we show that Rydberg atoms offer with a simple way to implement networks with a mixture of inhibitory and excitatory neurons, thereby demonstrating the versatility of neutral atoms at mimicking the behavior of realistic neural circuits.

Two Rydberg atoms with different principal quantum numbers n_Q , and n_Q' and angular momentum quantum numbers the same can interact with a $1/r^6$ attractive potential $V_{n_Q, n_Q'}$ [45]. Using the Python package `PairInteraction` [58], we note that if n_Q represents the

state $|r\rangle = |70S_{1/2}, m_j = -1/2, m_I = -3/2\rangle$, and n'_Q represents $|r'\rangle = |73S_{1/2}, m_j = -1/2, m_I = -3/2\rangle$, then the interaction $V_{n_Q, n'_Q} = V \approx -V_{n_Q, n'_Q}$ where V is the strength between atoms with principal quantum numbers n_Q (see Appendix D). We can use this fact to encode inhibitory neurons. We restrict the concentration of n'_Q Rydberg atoms to be sparse such that pairs of n'_Q atoms are placed as far as possible at a distance d_{max} in a 1D chain arrangement. We choose the field strength V so that $V/d_{max}^6 = 10^{-2}$, and as a result we can neglect the interactions between pairs of n'_Q atoms, but not the interactions between pairs $(n_Q)(n'_Q)$ and $(n_Q)(n_Q)$. This amounts to saying that if atom n is driven to n'_Q then for all m $J_{nm} \lesssim 0$ as in (14). By implementing this in on our reservoir we can learn XOR, AND, and OR simultaneously for different concentrations of inhibitory neurons as illustrated in Fig. 5(A).

Fig. 5(B) shows the errors of simultaneously learning XOR, OR, and AND as a function of the system size N for different number of inhibitory neurons in the array. The network is initialized in the state $|g\rangle^{\otimes N}$, and the network receives two binary inputs $x, y \in \{0, 1\}$ (in units of MHz) for a time Δt (in units of μs) with input noise $\sigma_{in} = 0.1$. Afterwards, the network is interrogated to give XOR(x, y), OR(x, y), and AND(x, y). W^{out} is trained using the loss in (4). The errors shown in Fig. 5(B) are the minimum achieved over a wide range of choices of interaction time $\Delta t \in [0, 5] \mu s$. This shows that in some cases our qRNN can benefit from having a connectivity matrix J_{nm} with both positive (excitatory) and negative (inhibitory) values, analogously to the mammalian brain. For small system sizes, it seems that a ratio of 1:4 inhibitory neurons better the learning performance, similar to the results in [56]. This is supported by the performance at 4 and 8 neurons in Fig. 5(B). Particularly, $N = 8$ neurons, two of which are inhibitory, results in a 40% decrease of the loss. Nonetheless, we observe that having no inhibitory neurons is best when dealing with $N = 6$ and 10 neurons. No inhibitory neurons is never the worse choice. Fig. 5(C)-(E) shows the results of learning XOR, OR, and AND simultaneously using $N = 8$ and two inhibitory neuron. Note that the network is fully capable of classification with errors well below the input noise threshold σ_{in} .

Lastly, while this task shows the success of the qRNN at approximating Boolean functions of the input, we note that one may also want to calculate different nonlinear functions of the input. We remark that our qRNN can approximate biologically relevant nonlinear functions such as ReLu and sigmoid.

B. Decision making

One of the great successes of classical RNNs is their ability to integrate sensory stimuli in order to choose between two actions. Here, we present the Rydberg qRNN




FIG. 6. Decision making task using a Rydberg qRNN. **(A)** Schematic of the input stimuli as a pair of time-dependent detunings on two atoms. The stimuli are turned on for a normally distributed time Δt with standard deviation $\sigma_{in} = 0.1$. The network decides on a relaxation time t_{out} to output the decision sign $(\Delta_1^{in} - \Delta_2^{in})$. **(B)** The psychometric response of the decision making task which maps the accuracy towards deciding that Δ_1^{in} is largest as a function of the inputs' difference. The simulated response (dotted) is well fitted by a sigmoid function (solid curve).

with a variant of the dot motion decision making task initially studied in monkeys in which several inputs are analyzed to produce a scalar nonlinear function of them which represents a decision [59]. This task shows the Rydberg qRNN's ability to produce nonlinear functions of the input and perform simple cognitive tasks, a feature of most qRNNs proposed thus far [60].

In this task, a qRNN is presented with two inputs Δ_1^{in} and Δ_2^{in} , and the goal is to train the network to choose which input is the largest. That is,

$$y^{\text{targ}} = \text{sign}(\Delta_1^{in} - \Delta_2^{in}). \quad (15)$$

The stimuli, which in the case of a qRNN are local detunings on a pair of atoms, are turned on for a normally distributed time Δt with variance also $\sigma_{in} = 0.1$ and mean $\langle \Delta t \rangle = 0.1 \mu s$ (see Fig. 6(A)). The stimuli are then turned off, and the network chooses a relaxation time t_{out} after which it “makes a decision” by approximating (15). This is known as the fixed-duration protocol since the experimentalist fixes the stimulation period, and the subject, the qRNN in this case, learns to choose a response time t_{out} .

In the brain, we expect the performance of a decision making task to follow a sigmoidal psychometric response [44, 59]. A psychometric response maps out the accuracy of a decision making task as a function of stimuli distinguishability. As an example of a psychometric response, the reader could think about paying a routine visit to the eye-doctor and having to discern letters “b” and “p” written on the wall. If the letters are large enough, they become distinguishable, and if the letters are too small one often fails to make out the right letter.

Classically, a decision-making task benefits from connectivity between all neurons. Since our connectivity is limited by physical constraints, a 2D square lattice structure was chosen to prevent neurons from being isolated

from the rest. Moreover, a 2D square lattice is experimentally friendly. We set up a Rydberg qRNN of 3×2 atoms with two input atoms and two different output atoms (for details see Appendix D). The qRNN is then trained by optimizing over t_{out} , and W^{out} such that the qRNN's output approximates (15) while keeping the network parameters J_{nm} , Ω and Δ_n^{in} fixed. We observe that $t_{out} \approx 1 \mu s$ is regularly obtained as this is the time scale in which the information about $\Delta_{1,2}^{in}$ propagates through the network. In our case, $c_1 = \Delta_1^{in} - \Delta_2^{in}$ is a natural choice for a measure of stimuli distinguishability. Fig. 6(B) shows the psychometric response of the task which is qualitatively similar to the ones obtained in classical RNNs [44]. Moreover, we see in Fig. 6(B) that if $|c_1| \geq \sigma_{in}$ such that it is above the input noise level our network success more than 80% of the time. The success of this task shows the Rydberg qRNN's ability to emulate simple cognitive tasks.

C. Parametric working memory

Our next neurological task is that of parametric working memory. Working memory, which is one of the most important cognitive functions, deals with the brain's ability to retain and manipulate information for the later execution of a task. Here, we train a network to perform a task based on the decision making task in Sec. VB but with two temporally separate stimuli (see Fig. 7(A)). We use the fixed-time protocol where the separation between stimuli, denoted by t_{delay} , is fixed by us. The stimuli are both turned on for a time Δt , and after the second input the network is left to relax for a time t_{out} before two output neurons are used to approximate (15). To avoid overfitting, we add Gaussian noise to the times Δt , t_{out} , and t_{delay} with zero mean and standard deviation $\sigma_{in} = 0.1$. The network optimizes over W^{out} . Thus, the network has to retain information about Δ_1^{in} for a few “seconds” to then compare against Δ_2^{in} and make a decision.

We set a Rydberg qRNN of 3×2 atoms with two input atoms and two different output atoms (for details see Appendix D). Fig. 7(B) shows the loss of the network as a function of the total time the inputs are injected into the network ($\tau = 2\Delta t + t_{delay}$). We note that the loss function is high for small τ since it takes the input neurons to correlate with the rest of the qRNN. Accordingly, in Fig. 7(B) we show that growth of the entanglement entropy of the input qubits accompanies a decrease in the loss function. For Fig. we fixed $t_{out} = 0.1$, a choice which, according to Fig. 7(C), has little effect on the qRNN's performance.

In Fig. 7(C) we show the accuracy of the qRNN at reproducing (15) as a function of the time the inputs are turned on (Δt) and for different choices of t_{out} . For these plots $t_{delay} = 0.1$ is fixed. We notice that the accuracy is largely invariant to our sampled choices of t_{out} .

Lastly, in Fig. 7(D), we probe the qRNN's accuracy




FIG. 7. Working memory of a Rydberg RNN. (A) Schematic of the network's inputs where two atoms are detuned for a time Δt but temporally separated by a time t_{delay} . Two different output neurons are used for readout at a time t_{out} after the second input is turned off. (B) Loss of the working memory task as a function of the total input time $2\Delta t + t_{delay}$ (gray). Entanglement entropy between the input qubits and the rest of the qRNN as a function of $2\Delta t + t_{delay}$ (blue). Here, the mean value of t_{out} is 0.1. The loss stays large for small input times until the input qubits start entangling with the rest of the qRNN. (C) Accuracy as a function of the time the input are turned on (Δt) for four different choices of t_{out} and with fixed $t_{delay} = 0.1$. These curves show that accuracy is largely independent of t_{out} and Δt as long as $\Delta t < 0.3$. (D) Accuracy of the working memory task at $\Delta t = 0.15$ and $t_{out} = 0.5$ as a function of t_{delay} . The blue curve is the performance when $V > \Omega$ which puts the qRNN in the Rydberg blockaded regime, while the red curve is the performance when $V < \Omega$ which puts the qRNN in the disordered regime. This plots show that when $V > \Omega$, the Rydberg qRNN can hold memory for later manipulation better than when $V < \Omega$. Shaded regions indicate error bars.

as a function of t_{delay} . For these experiments, we fix $t_{out} = 0.5$ and $\Delta t = 0.15$. Importantly we set $V = 2\pi \times 10$ MHz and $\Omega = 2\pi \times 4.2$ MHz such that $V > \Omega$ and neighboring Rydberg excitations are off-resonance putting our qRNN in the so-called *blockaded regime* [61, 62]. While one initially might expected the accuracy to decrease for increasing t_{delay} , we found that this is not the case and instead the accuracy oscillates persistently reaching high accuracies as shown in Fig. 7(D) blue curve. Interestingly, this behavior disappears when the coupling $V = 2\pi \times 0.1$ MHz such that $V < \Omega$ as shown in the red curve in Fig. 7(D), although the performance is statistically significant even for long t_{delay} with an accuracy greater than 50%. We can conclude that, in the blockaded regime, the qRNN is able to hold information for longer periods of time. We can understand this dependence on V/Ω as follows. In the disordered regime atoms are mostly uncorrelated and the atoms are allowed to freely oscillate with the dynamics being dominated by




FIG. 8. A state encoding a memory m is prepared. The state evolves under its natural Hamiltonian before being interrogated via local measurements in order to retrieve m . If the evolution time is short, the system is yet out of equilibrium and remembers its initial condition. Thus, m can be retrieved. On the other hand, after a long time the system may thermalize and local measurements fail to provide information of the initial state. Thus, the memory retrieval time is upper bounded by the thermalization time of the initial state $|\psi_m(0)\rangle$ under the system's dynamics. In the example in Sec. VD, the system is a chain of Rydberg atoms, and final measurements are performed on a single atom which are then linearly post-processed so to retrieve m . In this case, a thermal state can be observed by measuring if the entanglement entropy of ring obeys a volume-law. If the dynamics can be stabilized against thermalization, the memory can be retrieved at larger times.

the drive Ω . Thus, after a short period of time, the inputs coming through a z -field are largely irrelevant and the network is unable to hold the information about the first input. On the other hand, when $V > \Omega$ the atoms are largely correlated since neighboring excitations of Rydberg atoms are blockaded and the dynamics are slowed down. These slow dynamics in the system allow for longer memory times. In Sec. VD, we will explore the longer-term memory in the blockaded regime and show that long-term memory in a qRNN can be stabilized due to the presence of quantum many-body scars.

D. Long-term Memory via Quantum Many-body Scars

Finally, we turn to examine a qRNNs ability to encode long-term memory. The task consists of encoding a classical bit m in the initial state of a qRNN $|\psi_m(0)\rangle$ so that after the system is left to evolve under its inherent dynamics for a time T , local measurements of the state $|\psi_m(T)\rangle$ are used to recover m . However, m cannot be recovered from local measurements if the dynamics obey the eigenstate-thermalization hypothesis (ETH) [63]. Instead, local measurements of $|\psi_m(T)\rangle$ obey thermal statistics described by the energy spectrum of the Hamiltonian and bear no information of the initial con-

dition $|\psi_m(0)\rangle$. Thus, qRNNs which violate the ETH are naturally suited for memory tasks, since they can locally retain information about their initial state. Indeed this notion has began to be studied in quantum reservoirs [25, 27]. Recent experiments using quench dynamics in arrays of Rydberg atoms have revealed quantum many-body scarring behavior [35], which can be stabilized [46, 47] to delay the thermalization of the system. Here, we use these results to enlarge the memory lifetime of a qRNN. Simulation details are found in Appendix D.

In the case of a kicked ring of Rydberg atoms experiencing nearest-neighbor blockade, the dynamics are captured by the so-called PXP-model [35, 47, 64, 65]

$$H(t) = H_{\text{PXP}} + \hat{N} \sum_{k \in \mathbb{Z}} \theta_k \delta(t - k\tau) \quad (16)$$

$$H_{\text{PXP}} = \Omega \sum_{n=1}^N P_{n-1} \sigma_n^x P_{n+1} \quad \hat{N} = \sum_n \hat{n}_n \quad (17)$$

where $P_n = |g\rangle\langle g|_n$ projects the atom at the n^{th} site onto the ground state and we choose periodic boundary conditions to mitigate edge effects. In (16) we let $\theta_k = \pi + \epsilon_k$ where ϵ_k is a Gaussian random variable with mean ϵ and variance σ_{in}^2 . That is, ϵ_k plays the role of added noise in the qRNN. For this discussion we let $\gamma = 0$ since we know from experiments that the quantum scarring behavior is robust to the atom's decoherence, and the choice to work with the Hamiltonian evolution helps speed up the acquisition of numerical data.

We denote $\chi_\tau = \exp(-i\pi\hat{N}) \exp(-i\tau H_{\text{PXP}})$. It has been empirically observed that χ_τ approximately exchanges the Neel states $|AF\rangle = |1010\dots\rangle$ and $|AF'\rangle = |0101\dots\rangle$ for $\tau \approx 1.51\pi\Omega^{-1}$ [35]. Note that $\chi_\tau \chi_\tau = \mathbb{1}$, and so under no noise, any state $|\psi\rangle$ is recovered after a cycle of evolution of 2τ . However, the noise ϵ_k destroys the revival of all initial states except for $|AF\rangle$ and $|AF'\rangle$ (see Appendix C). This leads to many-body quantum scars stabilized by the operator $\exp(-i\pi\hat{N})$ [46, 47].

Given the dynamics in (16), we propose the following scheme for encoding a binary memory $m \in \{0, 1\}$. We choose a reference state $|\psi\rangle$, and let $|\psi_0(0)\rangle = |\psi\rangle$ and $|\psi_1(0)\rangle = \chi_\tau |\psi\rangle$. Subsequently, the state $|\psi_m(0)\rangle$ is left to evolve for n cycles of duration $2\tau = 2(1.51\pi)$ after which the populations $\mathbf{r}_m(n) = (P_g(2n\tau|m), P_r(2n\tau|m))$ of the single-atom reduced density matrix are used to retrieve m . The retrieval is done by training a vector W_n^{out} on M instances of $\mathbf{r}_m(n)$ in order to minimize (4) with $\mathbf{y}^{\text{targ}} = \mathbf{m}$ the binary vector of memories and $\mathbf{y}^{\text{out}}(n) = W_n^{\text{out}} \mathbf{r}(n)$ our networks' output after n cycles.

To quantify the quality of the memory retrieval $R(n)$, we use the squared Pearson's r -factor

$$R(n) = \frac{\text{cov}^2(\mathbf{m}, \mathbf{y}^{\text{out}}(n))}{\sigma^2(\mathbf{m})\sigma^2(\mathbf{y}^{\text{out}}(n))}. \quad (18)$$

Fig. 9(a) shows the memory retrieval error as a function of the number of cycles for three different choices of reference states. Fig. 9(b) shows the average entanglement

entropy ($\bar{S}_E$) of the left-most atom in the ring. Saturation of $\bar{S}_E$ signals a growth in the memory retrieval error as the state “forgets” the initial condition. From other studies, we see that memory is retrieved at longer times due to the slow thermalization of the Neel states due to quantum many-body scars [35, 46, 47, 64, 65]. The time-crystalline nature of the qRNN using $|\psi\rangle = |AF\rangle$ signals long-time correlations, and thus the qRNN can be used to encode and predict series with long-time correlations [17].

The Neel states exhibit long-term memory due to the evolution’s scarring behavior. This can be understood by analyzing the average evolution produced by a single cycle. Up to second order in ϵ_k , the state at time $2\tau n$, $\rho(n)$, evolves to the state at time $2\tau(n+1)$, $\rho(n+1)$, where (see Appendix C)

$$\begin{aligned} \rho(n+1) = & \rho(n) - i\epsilon[H^+, \rho(n)] \\ & + \sigma_{in}^2 \left(H^+ \rho(n) H^+ - \frac{1}{2} \{H^+ H^+, \rho(n)\} \right) \\ & + \sigma_{in}^2 \left(H^- \rho(n) H^- - \frac{1}{2} \{H^- H^-, \rho(n)\} \right). \end{aligned} \quad (19)$$

Here, $H^\pm = \hat{N} \pm \chi_\tau \hat{N} \chi_\tau$ are Hermitian operators. We can rewrite (19) as $\rho(n+1) = \rho(n) + \mathcal{L}_{\epsilon,\sigma}(\rho(n))$. Since $[H^+, \chi_\tau] = 0$, the operator H^+ has an emergent $\mathbb{Z}_2$ symmetry which means that the ground states of H^+ are well approximated by the states $|\pm\rangle = \frac{1}{\sqrt{2}}(|AF\rangle \pm |AF'\rangle)$ [47]. Note that

$$H^+|+\rangle \approx N|+\rangle, \quad H^-|+\rangle \approx 0, \quad (20)$$

$$H^+|-\rangle \approx N|-\rangle, \quad H^-|-\rangle \approx 0, \quad (21)$$

where N is the size of the system. We conclude that if $\rho(n) = |AF\rangle\langle AF|$ then $\rho(n+1) \approx \rho(n)$ as this state is (approximately) in the kernel of $\mathcal{L}_{\epsilon,\sigma}$. Therefore, the Neel states are suitable memory states.

Equation (19) also tell us that any density matrix in the kernel of $\mathcal{L}_{\epsilon,\sigma}$ may also serve as a memory state since it is a steady states of the evolution. This would allow us to enlarge the number of memories accessible in a qRNN. In Appendix C we show the existence of a large number of steady states, and we present a scheme to prepare a number of them. It’s worth noting, however, that these memory may have to be distinguished from one another via global measurements. The questions of how to efficiently prepare and distinguish these memory states remain importantly both opened and key in telling us if a memory quantum advantage can be claimed in qRNNs. As it stands, using quantum scars signals that Rydberg-inspired RNNs may present enhanced memory since quantum scars are classically simulatable due to their low entanglement entropy. However, it’s unclear whether the system can be classically simulated at late times due to the onset of the thermalization. These questions are left for future studies.




FIG. 9. Dependence of memory retrieval on different reference states. We use a ring of $N = 8$ Rydberg atoms with $\epsilon = \sigma = 0.1$, and $M = 100, 30$ samples for the training and testing sets respectively. The memories are chosen from a balanced Bernoulli distribution. (a) Shows the memory retrieval error for three different choices of reference state, $|AF\rangle = |grgrgr\rangle$, $|gg\rangle = |gg\dots g\rangle$ and $|d_2\rangle = |grgggr\rangle$. Due to the scarring behavior of $|AF\rangle$, the memory length is greatly improved. (b) Shows the left-most atom’s entanglement entropy averaged over the M memory instances ($\bar{S}_E$). Saturation of $\bar{S}_E$ signals the thermalization of the system and thus a decrease in R .

Quite recently, another proposal to enlarge the number of memories accessible in a qRNN has been introduced using the emergent scale-free network dynamics of a melting discrete time-crystal in an Ising chain [25]. The proposal in [25] can be seen as generalization of the qRNN presented in (16) by dropping the constraint of the Rydberg blockade. Our results, as well as those in [25], pose the possibility of having an RNN with a memory capacity that outpaces that of classical RNNs such as the Hopfield network [66].

VI. CONCLUSIONS AND OUTLOOK

In this article, we present a quantum extension of a classical RNN on binary neurons. This implies a deep connection between controllable many-body quantum systems and brain-inspired computational models. Our qRNN facilitates the ability to employ the analogue dynamics of quantum systems for computation instead of the circuit-based paradigm. We show how features of the quantum evolution of our qRNN can be used for quantum learning tasks, and to speedup simulation of stochastic dynamics. We implement a qRNN using arrays of Rydberg atoms and show how Rydberg atoms analogously perform biological tasks even in the presence of a few atoms. This can be explained via the physics of the

system. For example, we showed how weak-ergodicity breaking collective dynamics in Rydberg atoms can be employed for long-term memory.

While this article takes a first step forward in connecting controllable quantum systems and neural networks from a fundamental perspective, several questions remain unanswered. Firstly, from the first two quantum features hereby presented, studies of how qRNNs can be used for quantum error correction in circuit-like quantum computing are warranted. Directly from this work, investigations into advantageous stochastic processes in qRNNs which are robust to decoherence are enticing. These advantages will likely emerge from collective behavior of quantum neurons. Therefore, the field will soon require a thorough understanding of the collective dissipative dynamics of neurons in qRNNs, which would also shed light into rigorous studies of the computational power of these architectures. Guided by the fact that neural networks become universal approximators by interconnecting many neurons, one may also consider the spatial and control requirements necessary for universal brain-inspired quantum machine learning.

Given the vast number of classical computational models for the brain, there are several immediate research directions. One of these is the exploration of a systematic way to quantize more biologically realistic models of neural circuit. A possible starting point for translating different neural circuits would be to exploit key engineering and fundamental features of different NISQ platforms. For example, recent experiments using Rydberg atoms in photonic cavities may provide us with the ability to capture neural plasticity on qRNNs by arbitrarily tuning the inter-neural interactions [67]. Likewise, superconducting circuits have lately been used to encode biologically realistic single-neuron models [13]. Along these explorations,

it will be imperative to establish a variety of methods to analyze how quantum neural networks recover the classical protocols in certain limits, as well as the source and extend of the quantum advantages that each platform can offer.

Lastly, while our memory encoding scheme in Sec. V D offers a possibility to encode a binary memory, whether a higher number of memories can be encoded efficiently remains an important open question. In Appendix C we offer a proposal based on the steady states of the effective dissipative evolution in the pre-thermalization regime introduced by the noise in the qRNN. This already shows a theoretical number of memories greater than those attainable by the vanilla Hopfield network [66]. However, distinguishing these memories, or producing Hamiltonians with a desired memory state in mind, is left for future research. It is clear, however, that memory in a quantum reservoir relies on ergodicity breaking dynamics [25, 27]. Hamiltonian engineering techniques, together with more general driven Hamiltonians such as those in [25], may pave the way towards programmable memories in a qRNN.

ACKNOWLEDGMENTS

The authors thank Mikhail D. Lukin and Nishad Maskara for insightful discussion. RAB acknowledges support from NSF Graduate Research Fellowship under Grant No. DGE1745303, as well as funding from Harvard University's Graduate Prize Fellowship. XG acknowledges support from Quantum Science of the Harvard-MPQ Center for Quantum Optics, the Templeton Religion Trust grant TRT 0159, and by the Army Research Office under Grant W911NF1910302 and MURI Grant W911NF-20-1-0082. SFY acknowledges funding from NSF and AFOSR.

-
- [1] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
 - [2] A. W. Harrow, A. Hassidim, and S. Lloyd, Quantum algorithm for linear systems of equations, *Physical Review Letters* **103**, 10.1103/physrevlett.103.150502 (2009).
 - [3] N. Wiebe, D. Braun, and S. Lloyd, Quantum algorithm for data fitting, *Physical Review Letters* **109**, 10.1103/physrevlett.109.050505 (2012).
 - [4] G. H. Low, T. J. Yoder, and I. L. Chuang, Quantum inference on bayesian networks, *Physical Review A* **89**, 10.1103/physreva.89.062315 (2014).
 - [5] S. Lloyd, M. Mohseni, and P. Rebentrost, Quantum principal component analysis, *Nature Physics* **10**, 631 (2014).
 - [6] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nature communications* **9**, 1 (2018).
 - [7] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, *Nature communications* **12**, 1 (2021).
 - [8] D. Marković and J. Grollier, Quantum neuromorphic computing, *Applied Physics Letters* **117**, 150501 (2020).
 - [9] B. Kiraly, E. J. Knol, W. M. J. van Weerdenburg, H. J. Kappen, and A. A. Khajetoorians, An atomic Boltzmann machine capable of self-adaption, *Nature Nanotechnology* 10.1038/s41565-020-00838-4 (2021).
 - [10] P. Mujal, R. Martínez-Peña, J. Nokkala, J. García-Bení, G. L. Giorgi, M. C. Soriano, and R. Zambrini, Opportunities in quantum reservoir computing and extreme learning machines, *arXiv preprint arXiv:2102.11831* (2021).
 - [11] P. Pfeiffer, I. Egusquiza, M. Di Ventra, M. Sanz, and E. Solano, Quantum memristors, *Scientific reports* **6**, 1 (2016).
 - [12] T. Gonzalez-Raya, X.-H. Cheng, I. L. Egusquiza, X. Chen, M. Sanz, and E. Solano, Quantized single-ion-channel hodgkin-huxley model for quantum neurons, *Physical Review Applied* **12**, 014037 (2019).
 - [13] T. Gonzalez-Raya, E. Solano, and M. Sanz, Quantized three-ion-channel neuron model for neural action potentials, *Quantum* **4**, 224 (2020).

- [14] E. Torrontegui and J. J. García-Ripoll, Unitary quantum perceptron as efficient universal approximator, *EPL (Europhysics Letters)* **125**, 30004 (2019).
- [15] K. Fujii and K. Nakajima, Quantum reservoir computing: a reservoir approach toward quantum machine learning on near-term quantum devices, *arXiv preprint arXiv:2011.04890* (2020).
- [16] K. Nakajima, K. Fujii, M. Negoro, K. Mitarai, and M. Kitagawa, Boosting computational power through spatial multiplexing in quantum reservoir computing, *Phys. Rev. Applied* **11**, 034021 (2019).
- [17] A. Kutvonen, K. Fujii, and T. Sagawa, Optimizing a quantum reservoir computer for time series prediction, *Scientific Reports* **10**, 14687 (2020).
- [18] S. Ghosh, A. Opala, M. Matuszewski, T. Paterek, and T. C. H. Liew, Quantum reservoir processing, *npj Quantum Information* **5**, 10.1038/s41534-019-0149-8 (2019).
- [19] S. A. Khan, F. Hu, G. Angelatos, and H. E. Türeci, Physical reservoir computing using finitely-sampled quantum systems, *arXiv preprint arXiv:2110.13849* (2021).
- [20] S. Ghosh, T. Paterek, and T. C. H. Liew, Quantum neuromorphic platform for quantum state preparation, *Phys. Rev. Lett.* **123**, 260404 (2019).
- [21] L. Govia, G. Ribeill, G. Rowlands, H. Krovi, and T. Ohki, Quantum reservoir computing with a single nonlinear oscillator, *Physical Review Research* **3**, 013077 (2021).
- [22] J. Nokkala, R. Martínez-Peña, G. L. Giorgi, V. Parigi, M. C. Soriano, and R. Zambrini, Gaussian states of continuous-variable quantum systems provide universal and versatile reservoir computing, *Communications Physics* **4**, 1 (2021).
- [23] S. Ghosh, T. Krisnanda, T. Paterek, and T. C. Liew, Realising and compressing quantum circuits with quantum reservoir computing, *Communications Physics* **4**, 1 (2021).
- [24] P. Mujal, Quantum reservoir computing for speckle-disorder potentials, *arXiv preprint arXiv:2201.11096* (2022).
- [25] A. Sakurai, M. P. Estarellas, W. J. Munro, and K. Nemoto, Quantum reservoir computation utilising scale-free networks, *arXiv preprint arXiv:2108.12131* (2021).
- [26] W. Xia, J. Zou, X. Qiu, and X. Li, The reservoir learning power across quantum many-body localization transition, *Frontiers of Physics* **17**, 1 (2022).
- [27] R. Martínez-Peña, G. L. Giorgi, J. Nokkala, M. C. Soriano, and R. Zambrini, Dynamical phase transitions in quantum reservoir computing, *Phys. Rev. Lett.* **127**, 100502 (2021).
- [28] A. Coolen, Statistical mechanics of recurrent neural networks I—statics, in *Handbook of biological physics*, Vol. 4 (Elsevier, 2001) Chap. 14, pp. 553–618.
- [29] Y. Suzuki, Q. Gao, K. C. Pradel, K. Yasuoka, and N. Yamamoto, Natural quantum reservoir computing for temporal information processing, *Scientific Reports* **12**, 1 (2022).
- [30] M. Saffman, T. G. Walker, and K. Mølmer, Quantum information with Rydberg atoms, *Rev. Mod. Phys.* **82**, 2313 (2010).
- [31] B. J. Lester, N. Luick, A. M. Kaufman, C. M. Reynolds, and C. A. Regal, Rapid production of uniformly filled arrays of neutral atoms, *Physical review letters* **115**, 073003 (2015).
- [32] D. Barredo, S. De Léséleuc, V. Lienhard, T. Lahaye, and A. Browaeys, An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays, *Science* **354**, 1021 (2016).
- [33] M. Endres, H. Bernien, A. Keesling, H. Levine, E. R. Anschuetz, A. Krajenbrink, C. Senko, V. Vuletic, M. Greiner, and M. D. Lukin, Atom-by-atom assembly of defect-free one-dimensional cold atom arrays, *Science* **354**, 1024 (2016).
- [34] H. Labuhn, D. Barredo, S. Ravets, S. De Léséleuc, T. Macrì, T. Lahaye, and A. Browaeys, Tunable two-dimensional arrays of single rydberg atoms for realizing quantum ising models, *Nature* **534**, 667 (2016).
- [35] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, *Nature* **551**, 579 (2017).
- [36] A. Cooper, J. P. Covey, I. S. Madjarov, S. G. Porsev, M. S. Safronova, and M. Endres, Alkaline-earth atoms in optical tweezers, *Physical Review X* **8**, 041055 (2018).
- [37] J. Wilson, S. Saskin, Y. Meng, S. Ma, R. Dilip, A. Burgers, and J. Thompson, Trapping alkaline earth rydberg atoms optical tweezer arrays, *Physical Review Letters* **128**, 033201 (2022).
- [38] S. Ebadi, T. T. Wang, H. Levine, A. Keesling, G. Semeghini, A. Omran, D. Bluvstein, R. Samajdar, H. Pichler, W. W. Ho, *et al.*, Quantum phases of matter on a 256-atom programmable quantum simulator (2020), *arXiv:2012.12281 [quant-ph]*.
- [39] L. Isenhower, E. Urban, X. Zhang, A. Gill, T. Henage, T. A. Johnson, T. Walker, and M. Saffman, Demonstration of a neutral atom controlled-not quantum gate, *Physical review letters* **104**, 010503 (2010).
- [40] H. Pichler, S.-T. Wang, L. Zhou, S. Choi, and M. D. Lukin, Quantum optimization for maximum independent set using Rydberg atom arrays (2018), *arXiv:1808.10816 [quant-ph]*.
- [41] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. Wang, S. Ebadi, H. Bernien, A. Zibrov, H. Pichler, S. Choi, *et al.*, Generation and manipulation of Schrödinger cat states in Rydberg atom arrays, *Science* **365**, 570 (2019).
- [42] L. Henriët, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak, Quantum computing with neutral atoms, *Quantum* **4**, 327 (2020).
- [43] S. R. Cohen and J. D. Thompson, Quantum computing with circular rydberg atoms, *PRX Quantum* **2**, 030322 (2021).
- [44] H. F. Song, G. R. Yang, and X.-J. Wang, Training excitatory-inhibitory recurrent neural networks for cognitive tasks: A simple and flexible framework, *PLOS Computational Biology* **12**, 1 (2016).
- [45] J. Han and T. F. Gallagher, Millimeter-wave rubidium Rydberg van der waals spectroscopy, *Phys. Rev. A* **79**, 053409 (2009).
- [46] D. Bluvstein, A. Omran, H. Levine, A. Keesling, G. Semeghini, S. Ebadi, T. T. Wang, A. A. Michailidis, N. Maskara, W. W. Ho, *et al.*, Controlling quantum many-body dynamics in driven rydberg atom arrays, *Science* **371**, 1355 (2021).
- [47] N. Maskara, A. A. Michailidis, W. W. Ho, D. Bluvstein, S. Choi, M. D. Lukin, and M. Serbyn, Discrete time-crystalline order enabled by quantum many-body

- scars: entanglement steering via periodic driving, arXiv preprint arXiv:2102.13160 (2021).
- [48] J. Bausch, Recurrent quantum neural networks, arXiv preprint arXiv:2006.14619 (2020).
- [49] Y. Wu and X. Yang, Strong-coupling theory of periodically driven two-level systems, *Phys. Rev. Lett* **98**, 013601 (2007).
- [50] F. Ghafari, N. Tischler, J. Thompson, M. Gu, L. K. Shalm, V. B. Verma, S. W. Nam, R. B. Patel, H. M. Wiseman, and G. J. Pryde, Dimensional quantum memory advantage in the simulation of stochastic processes, *Phys. Rev. X* **9**, 041013 (2019).
- [51] D. H. Wolpert, A. Kolchinsky, and J. A. Owen, A space-time tradeoff for implementing a function with master equation dynamics, *Nature communications* **10**, 1 (2019).
- [52] K. Korzekwa and M. Lostaglio, Quantum advantage in simulating stochastic processes, *Phys. Rev. X* **11**, 021019 (2021).
- [53] C. Blank, D. K. Park, and F. Petruccione, Quantum-enhanced analysis of discrete stochastic processes, *npj Quantum Information* **7**, 1 (2021).
- [54] A. N. Burkitt, A review of the integrate-and-fire neuron model: I. Homogeneous synaptic input, *Biological Cybernetics* **95**, 1 (2006).
- [55] H. Jaeger, Short term memory in echo state networks, in *GMD - German National Research Institute for Computer Science* (GMD-Report 152, 2002).
- [56] V. Capano, H. J. Herrmann, and L. de Arcangelis, Optimal percentage of inhibitory synapses in multi-task learning, *Scientific Reports* **5**, 10.1038/srep09895 (2015).
- [57] J. C. Eccles, P. Fatt, and K. Koketsu, Cholinergic and inhibitory synapses in a pathway from motor-axon collaterals to motoneurons, *The Journal of Physiology* **126**, 524 (1954).
- [58] S. Weber, C. Tresp, H. Menke, A. Urvoy, O. Firstenberg, H. P. Büchler, and S. Hofferberth, Tutorial: Calculation of Rydberg interaction potentials, *J. Phys. B: At. Mol. Opt. Phys.* **50**, 133001 (2017).
- [59] J. D. Roitman and M. N. Shadlen, Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task, *The Journal of Neuroscience* **22**, 9475 (2002).
- [60] L. Govia, G. Ribeill, G. Rowlands, and T. Ohki, Non-linear input transformations are ubiquitous in quantum reservoir computing, *Neuromorphic Computing and Engineering* **2**, 014008 (2022).
- [61] E. Urban, T. A. Johnson, T. Henage, L. Isenhower, D. Yavuz, T. Walker, and M. Saffman, Observation of rydberg blockade between two atoms, *Nature Physics* **5**, 110 (2009).
- [62] A. Gaetan, Y. Miroshnychenko, T. Wilk, A. Chotia, M. Viteau, D. Comparat, P. Pillet, A. Browaeys, and P. Grangier, Observation of collective excitation of two individual atoms in the rydberg blockade regime, *Nature Physics* **5**, 115 (2009).
- [63] L. D'Alessio, Y. Kafri, A. Polkovnikov, and M. Rigol, From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics, *Advances in Physics* **65**, 239 (2016).
- [64] P. Fendley, K. Sengupta, and S. Sachdev, Competing density-wave orders in a one-dimensional hard-boson model, *Physical Review B* **69**, 075106 (2004).
- [65] I. Lesanovsky and H. Katsura, Interacting fibonacci anyons in a rydberg gas, *Physical Review A* **86**, 041601 (2012).
- [66] V. Folli, M. Leonetti, and G. Ruocco, On the maximum storage capacity of the hopfield model, *Frontiers in computational neuroscience* **10**, 144 (2017).
- [67] A. Periwal, E. S. Cooper, P. Kunkel, J. F. Wienand, E. J. Davis, and M. Schleier-Smith, Programmable interactions and emergent geometry in an atomic array, arXiv preprint arXiv:2106.04070 (2021).
- [68] G. S. Goodman, An intrinsic time for non-stationary finite Markov chains, *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* **16**, 165 (1970).
- [69] F. Reiter and A. S. Sørensen, Effective operator formalism for open quantum systems, *Physical Review A* **85**, 032111 (2012).
- [70] J. A. Nelder and R. Mead, A simplex method for function minimization, *The computer journal* **7**, 308 (1965).
- [71] M. Serbyn, D. A. Abanin, and Z. Papić, Quantum many-body scars and weak breaking of ergodicity, *Nature Physics* **17**, 675 (2021).

Appendix A: Probability transformations using qRNNs

In the case of the RNN presented in (1), using Ref. [28] we can derive that $P(\mathbf{s}|\mathbf{s}')$ in (8) is given by

$$P(\mathbf{s}|\mathbf{s}') = \prod_{i=1}^N \frac{1}{2} (1 + s_i g [h_i(\mathbf{s}')/\sigma_{in}^2 - 1])$$

where $g[x] = \text{Erf}[z/\sqrt{2}]$ is the error function due to the Gaussian noise. Regarding the task in Sec. III A of flipping all neurons at once, one could naively think that this can be done classically by taking the inputs $\Delta_n \rightarrow \infty$, however, since the noise's strength σ_{in}^2 scales as the size of the inputs, one obtains $P(\mathbf{s}|\mathbf{s}') \rightarrow \prod_{i=1}^N \frac{1}{2} (1 + s_i/2)$ which is a completely random update independent of the original state.

A transition matrix L obeys $L_{\mathbf{s}'|\mathbf{s}} \geq 0$ and $\sum_{\mathbf{s}'} L_{\mathbf{s}'|\mathbf{s}} = 1$. L is said to be *classically embeddable* if it can be generated by a continuous Markov process via

$$\frac{d}{dt}P(t) = K(t)P(t), \quad P(0) = \mathbf{1}, \quad P(t_f) = L, \quad (\text{A.1})$$

where K is called a generator matrix that preserves the positive nature of P via the constraint $K_{\mathbf{s}|\mathbf{s}'} \geq 0$ for $\mathbf{s} \neq \mathbf{s}'$, and normalization via the constraint $\sum_{\mathbf{s}} K_{\mathbf{s}|\mathbf{s}'} = 0$. Applied to our setup, a classically embeddable stochastic process is one that can transform $p_{t_f} = Lp_0$ via an RNN without employing any hidden neurons (i.e. $M = N$ neurons are used for readout), and in a single step. In general, determining if a matrix L is embeddable is an open question, but any embeddable matrix must necessarily satisfy [68]

$$\prod_{\mathbf{s}} L_{\mathbf{s}|\mathbf{s}} \geq \det L \geq 0. \quad (\text{A.2})$$

From (A.2), it immediately follows that the global “spin-flip” matrix F defined in (9) is not classically embeddable. That is, $\det F = 1$ and $\prod_{\mathbf{s}} F_{\mathbf{s}|\mathbf{s}} = 0$, violating

(A.2). Notice that the impossibility of performing F without hidden neurons is quite general, and it is not limited to the stochastic process allowed by (1). Moreover, the number of time-steps needed to achieve F using m hidden neurons is of order $\mathcal{O}(2^{N-m})$ (for details, see Sec. III.A in Ref. [52]).

Similar definitions of embeddability exist in the quantum setting. A stochastic process L is said to be *quantum embeddable* if there exist a Markovian quantum channel $\mathcal{E}$ such that

$$L_{\mathbf{s}'|\mathbf{s}} = \langle \mathbf{s}' | \mathcal{E}(|\mathbf{s}\rangle\langle \mathbf{s}|) | \mathbf{s}' \rangle. \quad (\text{A.3})$$

A Markovian quantum channel $\mathcal{E}$ is a channel arising from the time-evolution under a master equation, and thus $\mathcal{E}$ may include unitary and dissipative terms. As pointed out in Ref. [52], all classically embeddable stochastic processes. Moreover, permutations such as F in (9) are quantum embeddable since all permutations are unitary operators.

We highlight that realizing F is extremely sensitive to the decoherence arising from spontaneous emission, a main source of noise in NISQ devices. If γ is the decay rate at which spin $|1\rangle$ relaxes to $|-1\rangle$, one can show that the unitary evolution leads to the stochastic process F^γ where $\det F^\gamma = e^{-\mathcal{O}(2^N)}$. Notice that whether F^γ violates (A.2) becomes rapidly inconclusive with an increasing system size.

Appendix B: Continuous-time dynamics for a qRNN

A successful neural circuit model is the integrate and fire RNN (IF-RNN). In an IF-RNN each of the N neurons is influenced by pre-synaptic firing rates and produces a post-synaptic firing rate as an output. Each neuron is endowed with a firing rate $s_n(t)$, where n denotes the n^{th} neuron. The pre-synaptic firing rates arriving at the n^{th} neuron are integrated to produce a pre-synaptic current $I_n(t)$. In turn, the neuron produces a firing rate s_n influenced by its current and the firings of other neurons. Additionally, each neuron can receive a temporal input stimulus $\Delta_n^{\text{in}}(t)$ which affects both the currents and the firing rates. Generally, the firing rates and currents are described by non-linear, coupled differential equations of the form

$$\dot{I}_n = -\tau_I^{-1} I_n + G_n(\mathbf{s}(t), \mathbf{I}(t), J_{nm}, \Delta^{\text{in}}(t)) \quad (\text{B.1})$$

$$\dot{s}_n = -\tau_s^{-1} s_n + F_n(\mathbf{s}(t), \mathbf{I}(t), J_{nm}, \Delta^{\text{in}}(t)) \quad (\text{B.2})$$

where $\tau_{I,r}$ are relaxation time-constants for the currents and firing rates respectively. The vector $\mathbf{s}(t)$ is defined as $\mathbf{s}(t) = (s_1(t), \dots, s_N(t))$, with $\mathbf{I}(t)$, and $\Delta^{\text{in}}(t)$ defined analogously. The functions G and F ensure the dynamics are non-linear which gives RNNs their vast computational complexity. The specific forms of G and F depend on the application and relation between the currents and the firing rates one is trying to capture by the model.

The qRNN in Sec. IIIB follows the Heisenberg-Langevine equations of motion

$$\begin{aligned} \dot{A} = & i[H, A] + \sum_n \left(\frac{\gamma}{2} \sigma_n^+ + f_n^\dagger \right) [A, \sigma_n^-] \\ & + \sum_n [A, \sigma_n^+] \left(\frac{\gamma}{2} \sigma_n^- + f_n \right) \end{aligned} \quad (\text{B.3})$$

for any operator A . In (B.3), $\sigma_n^+ = |1\rangle\langle -1|_n$, $\sigma_n^- = (\sigma_n^+)^\dagger$, and f_n is a Langevin noise operator with Gaussian statistics $\langle f_n(t) \rangle = 0$ and $\langle f_n(t) f_m^\dagger(t') \rangle \propto \delta_{mn} \delta(t - t')$. In the equation above $[A, B] \equiv AB - BA$ stands for the commutator between matrices A and B .

In order to extract the statistics of the system, one may choose to look at the dynamics of two different local observables' expectation values. For example, the equations of motion for expectations of the local Pauli operators $\sigma_n^x = |-1\rangle\langle 1|_n + |1\rangle\langle -1|_n$, and $\sigma_n^y = i|-1\rangle\langle 1|_n - i|1\rangle\langle -1|_n$ are given by

$$\langle \dot{\sigma}_n^x \rangle = -\frac{\gamma}{2} \langle \sigma_n^x \rangle + i \langle [H(t), \sigma_n^x] \rangle \quad (\text{B.4})$$

$$\langle \dot{\sigma}_n^y \rangle = -\frac{\gamma}{2} \langle \sigma_n^y \rangle + i \langle [H(t), \sigma_n^y] \rangle \quad (\text{B.5})$$

with $H(t)$ specified by (5). The expectation values are calculated in the quantum-mechanical sense such that for an operator A , $\langle A \rangle = \text{Tr}(A\rho)$, and terms linear in f_n cancel out. Notice that the commutators in equations (B.4)-(B.5) play the role of the functions G and F in (B.1)-(B.2).

For σ_n^z , (B.3) gives

$$\dot{\sigma}_n^z = -\gamma/2 \sigma_n^z - \frac{\Omega}{2} \sigma_n^y + \gamma \mathbb{1}/2 - 2f_n^\dagger \sigma_n^- \quad (\text{B.6})$$

This can be integrated out to give

$$\begin{aligned} \sigma_n^z(t) - \sigma_n^z(0) = & \int_0^t dt' e^{-\gamma/2(t-t')} \left(-\frac{\Omega}{2} \sigma_n^y(t') \right. \\ & \left. - 2f_n^\dagger(t') \sigma_n^-(t') + \gamma/2 \mathbb{1} \right) \end{aligned} \quad (\text{B.7})$$

We choose to start the network at $\langle \sigma_n^z \rangle = -1$ for all n . We plug this back into (B.5), and we take the expectation values to eliminate terms linear in f_n . We obtain

$$\begin{aligned} \langle \dot{\sigma}_n^y \rangle = & -\frac{\gamma}{2} \langle \sigma_n^y \rangle + \frac{\Omega}{2} \sum_{m=1}^N J_{nm} \int_0^t dt' e^{-\gamma(t-t')} \langle \sigma_n^x(t) \sigma_m^y(t') \rangle \\ & + \Delta_n(t) \langle \sigma_n^x \rangle - \frac{\Omega^2}{4} \int_0^t \langle \sigma_n^y(t') \rangle e^{-\gamma(t-t')} dt'. \end{aligned} \quad (\text{B.8})$$

Similar equations can be found for $\langle \sigma_n^x \rangle$. Equation (B.8) tells us that $\langle \sigma_n^y \rangle$ depends on past statistics, and thus our network has a memory time bounded by $1/\gamma$. Let J denote the matrix J_{nm} . For $\gamma t \gg 1$, we can extend the

lower bound of integration to $-\infty$. Using the approximation $\int_{-\infty}^t e^{-\gamma(t-t')} f(t') dt' \approx -\gamma^{-1} f(t)$, we obtain

$$\begin{aligned} \langle \dot{\sigma}_n^x \rangle &= -\frac{\gamma}{2} \langle \sigma_n^x \rangle - \Delta_n^{in} \langle \sigma_n^y \rangle \\ &\quad - \frac{\Omega}{2\gamma} \sum_m J_{nm} \langle \sigma_n^y \sigma_m^y \rangle \end{aligned} \quad (\text{B.9})$$

$$\begin{aligned} \langle \dot{\sigma}_n^y \rangle &= -\left(\frac{\gamma}{2} + \frac{\Omega^2}{4\gamma}\right) \langle \sigma_n^y \rangle + \Delta_n^{in} \langle \sigma_n^x \rangle \\ &\quad - \frac{\Omega}{2\gamma} \sum_m J_{nm} \langle \sigma_n^x \sigma_m^y \rangle \end{aligned} \quad (\text{B.10})$$

thus leading to (11). In (B.10)-(B.9), the time dependence is implied.

Let us now define $s_n(t) \equiv \langle \sigma_n^y(t) \rangle$ and $I_n(t) \equiv \langle \sigma_n^x(t) \rangle$ so that $\mathbf{s}(t) = (s_1(t), \dots, s_N(t))$ and $\mathbf{I}(t) = (I_1(t), \dots, I_N(t))$. We see that (B.10)-(B.9) match (B.1)-(B.2) where

$$\begin{aligned} G_n &= -\Delta_n s_n - \frac{\Omega}{2\gamma} \sum_m J_{nm} \langle I_n I_m \rangle, \quad \tau_I^{-1} = \frac{\gamma}{2}, \quad (\text{B.11}) \\ F_n &= \Delta_n I_n - \frac{\Omega}{2\gamma} \sum_m J_{nm} \langle I_n s_m \rangle, \quad \tau_s^{-1} = \frac{\gamma}{2} + \frac{\Omega^2}{4\gamma}. \end{aligned} \quad (\text{B.12})$$

Equations (B.9)-(B.10) allow us to naturally interpret $\langle \sigma_n^y \rangle$ as the firing rate of the n^{th} neuron, and $\langle \sigma_n^x \rangle$ as the current. That is, the rate of the pre-synaptic neuron $\langle \sigma_k^y \rangle$ amounts to a current in the post-synaptic neuron $\langle \sigma_n^x \rangle$ that drives its rate $\langle \sigma_n^y \rangle$.

Equations (B.9)-(B.10) comprise a system of coupled quadratic differential equations, where the quadratic terms arise from the nontrivial commutation relation of the Pauli-operators $[\sigma_n^\alpha, \sigma_m^\beta] = i\delta_{\alpha\beta}\epsilon_{\alpha\beta\gamma}\sigma_n^\gamma$ where $\epsilon_{\alpha\beta\gamma}$ is the Levi-Cevita symbol. These quadratic terms in (B.9)-(B.10) make a qRNN a powerful computational system similar to how the functions G and F make an RNN a powerful computational system.

Appendix C: Memory and quantum many-body scars

As described in the main text, and more thoroughly discussed in Ref. [47], the scarring behavior of the kicked PXP-model is robust to fixed imperfections in the drive. The robustness persist even for random noise. Fig. 10 exemplifies the overlap with the initial condition for a noisy, kicked PXP-model for different values of ϵ and σ_{in} , which is a natural extension of the model in [47]. The Neel state $|AF\rangle$ exhibits robust revivals invariant of σ_{in}^2 . This fact can be explained with the effective theory presented below.

To understand the robustness of the quantum scarring behavior in the qRNN it is instructive to seek an effective description of the system's evolution. Recall that a




FIG. 10. Fidelities with the initial state after evolving for $n = 100$ cycles of noisy, kicked dynamics. The fidelity is defined as $F \equiv |\langle \psi | \psi(2n\tau) \rangle|^2$ where $|\psi\rangle$ is the initial state. Here, we used $L = 8$ Rydberg atoms and define $|AF\rangle = |grgrgr\rangle$, $|gg\rangle = |gggggggg\rangle$, and $|d_2\rangle = |grgggrgr\rangle$. The Neel state $|AF\rangle$ is robust to the noise in the drive since this state is invariant to decoherence up to second order in ϵ_i .

cycle is defined as two imperfect applications of χ_τ . The Hamiltonian in (16) produces the single-cycle unitary

$$U_\tau(\epsilon_1, \epsilon_2) = e^{-i\epsilon_2 \hat{N}} \chi_\tau e^{-i\epsilon_1 \hat{N}} \chi_\tau = e^{-i\epsilon_2 \hat{N}} e^{-i\epsilon_1 \chi_\tau \hat{N} \chi_\tau} \quad (\text{C.1})$$

where we use the fact that χ_τ is both Hermitian and unitary. Using the Baker-Campbell-Hausdorff formula to second order in ϵ_i , we can rewrite (C.1) as

$$U_\tau(\epsilon_1, \epsilon_2) \approx e^{-i(\epsilon_2 \hat{N} + \epsilon_1 \chi_\tau \hat{N} \chi_\tau)}. \quad (\text{C.2})$$

A state $\rho(n)$ evolves to $\rho(n+1) = U_\tau(\epsilon_1, \epsilon_2) \rho(n) U_\tau^\dagger(\epsilon_1, \epsilon_2)$ after a cycle. Expanding this to second order in ϵ_k and using the fact that $\langle \epsilon_k \rangle = \epsilon$ and $\langle \epsilon_k \epsilon_l \rangle = \sigma_{in}^2 \delta_{kl}$, we obtain the average evolution of the state

$$\begin{aligned} \rho(n+1) - \rho(n) &= -i\epsilon[H^+, \rho(n)] \\ &\quad + \sigma_{in}^2 \left(\hat{N} \rho(n) \hat{N} - \frac{1}{2} \{\hat{N}^2, \rho(n)\} \right) \\ &\quad + \sigma_{in}^2 \left(\chi_\tau \hat{N} \chi_\tau \rho(n) \chi_\tau \hat{N} \chi_\tau \right. \\ &\quad \left. - \frac{1}{2} \{\chi_\tau \hat{N}^2 \chi_\tau, \rho(n)\} \right). \end{aligned} \quad (\text{C.3})$$

Here, $\{A, B\} = AB + BA$ denote commutators and anti-commutators respectively. We define $H^+ = \hat{N} + \chi_\tau \hat{N} \chi_\tau$. For times $T \gg 2\tau$, we can take (C.3) to be a Lindbladian evolution since the noise satisfies the Markovian properties. We can rewrite (C.3) as

$$\dot{\rho} = \mathcal{L}_{\epsilon, \sigma}(\rho) \quad (\text{C.4})$$

$$\mathcal{L}_{\epsilon, \sigma}(\cdot) = -i\frac{\epsilon}{2\tau}[H^+, \cdot] + \frac{\sigma_{in}^2}{2\tau} D^+(\cdot) + \frac{\sigma_{in}^2}{2\tau} D^-(\cdot) \quad (\text{C.5})$$

$$D^\pm(\cdot) = H^\pm \cdot H^\pm + \frac{1}{2} \{H^\pm H^\pm, \cdot\} \quad (\text{C.6})$$

where $H^- = \hat{N} - \chi_\tau \hat{N} \chi_\tau$. For $\tau = 1.51\pi$, the Neel states are approximately simultaneous eigenstates of $\chi_\tau \hat{N} \chi_\tau$ and $\hat{N}$ with eigenvalues N for a system of size N . Thus, they are simultaneous eigenstates of $H^\pm$ and thus

$$\mathcal{L}_{\epsilon, \sigma}(|AF\rangle\langle AF|) \approx 0, \quad \mathcal{L}_{\epsilon, \sigma}(|AF'\rangle\langle AF'|) \approx 0. \quad (\text{C.7})$$




FIG. 11. Number of zero eigenvalues of the super-operator $\mathcal{L}_{\epsilon, \sigma}$ as a function of the system size. $\mathcal{L}_{\epsilon, \sigma}$ describes the effective dynamics of a qRNN composed of kicked Rydberg atoms. The number of zeros surpasses the linear number of memories available in the Hopfield network.




FIG. 12. Empirical memory states $\rho_{ss}(s)$ obtained from evolving the initial states $|s\rangle$ which are basis states of the Rydberg blockaded Hilbert space. N_m^e denotes the number of memories found using this procedure. The different plots show the fidelities $F(\rho_{ss}(s), \rho_{ss}(s'))$ between different steady states. The red squares delimit the basis states with different number of Rydberg excitations starting with the zero excitation sector on the top-left square and ending with $N/2$ excitations sector on the bottom-right square. Red arrows denote initial configurations for each of the N_m^e memories found empirically. While this procedure produces a number of memory states smaller than the number of zeros of $\mathcal{L}_{\epsilon, \sigma}$, $N_m^e > N$, a bound unattainable by common classical RNNs.

Therefore, the Neel states are steady states of the evolution. It is worth noting that $\mathcal{L}_{\epsilon, \sigma}$ captures the pre-thermal evolution. Ultimately, higher order effects in ϵ_k take over and lead to the thermalization of the Neel states similar to the results in [47] and as seen in Fig. 9. Nonetheless, the thermalization of the Neel states is delayed relative to other states due to (C.7).

Moreover, any density matrix ρ_{ss} in the kernel of $\mathcal{L}_{\epsilon, \sigma}$ can be used as a memory state. Expressing $\mathcal{L}_{\epsilon, \sigma}$ as a super-operator on density matrices, we can look at its spectrum which is in general complex. Fig. 11 shows the number of zero eigenvalues of $\mathcal{L}_{\epsilon, \sigma}$ for different system sizes N . It's evident that the number of zeros scales larger than linearly on N . Therefore, a qRNN evolving under $\mathcal{L}_{\epsilon, \sigma}$ may have a larger number of memory states

than a classical RNN. To prepare these states, we propose to initialize the qRNN on different string configurations $|s\rangle$ satisfying the Rydberg blockade constraint. For example, one can have $s = rgg..g$ while $s = rrg...g$ is not allowed. The system is left to evolve for some time T_{ss} to reach a steady state $\rho_{ss}(s)$ which can then be used as memory. Different initial strings can lead to different steady states as exemplified in Fig. 12. Fig. 12 shows the fidelity between $\rho_{ss}(s)$ and $\rho_{ss}(s')$ defined by the trace norm

$$F(\rho_{ss}(s), \rho_{ss}(s')) = \left(\text{Tr} \sqrt{\sqrt{\rho_{ss}(s)} \rho_{ss}(s') \sqrt{\rho_{ss}(s)}}} \right)^2. \quad (\text{C.8})$$

The red arrows in Fig. 12 indicates the different memory states obtained by this scheme. It's worth noting that this scheme offers us an empirical number of memories N_m^e that scales at most as ϕ^N , where $\phi \approx 1.62$ is the Golden ratio, since that's the number of basis states respecting the Rydberg blockade. We see that $N_m^e > N$ in all instances, a bound unattainable by classical RNNs such as the Hopfield network [66]. However, this scheme relies on an efficient way to recognize the different memory states through measurements, a question that we leave for future investigations.

Appendix D: Experimental values, and numerical simulations

In this section we outline the details for the experimental values used for the numerical simulation of Sec. V. Firstly, for simulating Rydberg atoms we use the experimental values in Ref. [38] for concreteness (see Fig. 13). In this experimental platform, a two-photon transition couples $|g\rangle = |5S_{1/2}\rangle$ and $|r\rangle = |50S_{1/2}\rangle$ via an off-resonance state $|6P_{3/2}\rangle$. For this setup and for short periods of time for the simulation ($< 10 \mu\text{s}$), the dominant source of decoherence are photon-scattering processes out of the intermediate state. Using the fact that the intermediate state is off-resonance, we can adiabatically eliminate it to produce an effective decay operator (see Sec. IV.B in Ref. [69])

$$\sigma_{eff}^- = \frac{\sqrt{\gamma_{420}}}{2\delta} |g\rangle (\Omega_{420}\langle g| + \Omega_{1013}\langle r|) \quad (\text{C.1})$$

which is an effective spontaneous emission from $|r\rangle$ to $|g\rangle$ accompanied with decoherence on the ground state.

We chose $\Omega = 4.2 \text{ MHz}$. Additionally, a pair of $|r\rangle$ atoms interact with a strength $C_6 = 862.9 \text{ GHz}(\mu\text{m})^6$. We used the `PairInteraction` python package from [58] to determine that a pair of $|r\rangle = |70S_{1/2}\rangle$ and $|r'\rangle = |73S_{1/2}\rangle$ has a similar interaction strength of $C_6^{rr'} = -836.6 \text{ GHz}(\mu\text{m})^6 \approx -C_6$. We used this interaction to model the inhibitory and excitatory neurons in Sec. V A ($V_{n_Q, n_Q} = V$, $V_{n_Q, n'_Q} = -V$). We denote $V = C_6/a_0^6$ where a_0 is tuned to give us different nearest neighbor interaction strengths.




FIG. 13. Schematic of Rydberg atoms as used in Ref. [38]. The ground state $|g\rangle = |5S_{1/2}\rangle$, and the Rydberg state $|r\rangle = |70S_{1/2}\rangle$ are coupled via a two-photon transition. An off-resonance 420 nm laser ($\Omega_{420} = 2\pi \times 160$ MHz, $\delta = 2\pi \times 1$ GHz) couples $|g\rangle$ with the intermediate $|6P_{3/2}\rangle$ state, and a 1013 nm laser ($\Omega_{1013} = 2\pi \times 50$ MHz) couples the intermediate state and $|r\rangle$ creating an effective drive between $|g\rangle$ and $|r\rangle$ at rate $\Omega = \frac{\Omega_{420}\Omega_{1013}}{\delta} = 2\pi \times 4.2$ MHz. Four spontaneous emission processes are at play: emission to nearby Rydberg atoms due to black-body-radiation at a rate $\gamma_{BBR} = 2\pi/(250 \mu s)$, photon-scattering out of the intermediate state into the ground state at rate $\gamma_{420} = 2\pi/(20 \mu s)$ and into the Rydberg state at rate $\gamma_{1013} = 2\pi/(150 \mu s)$, and spontaneous emission from $|r\rangle$ to $|g\rangle$ at rate $\gamma_{SE} = 2\pi/(375 \mu s)$. Since $\gamma_{BBR} + \gamma_{SE} + \gamma_{1013} = 2\pi/(75 \mu s)$ is smaller than γ_{420} , the leading source of decoherence for short periods of time ($< 10 \mu s$) is due to the γ_{420} decay.

Next, we explain and report the numerical parameters chose for each of the biological tasks.

1. Multitasking

Our scheme to encode inhibitory and excitatory neurons relies on approximating (13), and as a result one needs the “inhibitory neurons” to be as far away as possible from each other such that they do not interact positively with each other. For this reason, this task uses a 1D open chain of atoms separated by a distance a_0 with the inhibitory neurons being at opposite ends of the chain and in the bulk with maximum spacing from each other. The input neurons are chosen to be the two at one end of the chain, while the output neuron is chosen to be at the opposite end of the chain. This choice was made to ensure that the input neurons interact with the whole chain before readout.

The inputs are uniformly sampled from $\{0, 2\pi\}$ MHz with added Gaussian noise $\sigma_{in} = 0.1$, and all Δt sampled from a Gaussian with average $\langle \Delta t \rangle \in [0, 5]$ (μs) and standard deviation σ_{in} . For each size of the network and number of inhibitory neurons, We choose a_0 such that the separation between inhibitory neurons d_{max} results in $V/d_{max}^2 = 10^{-2}$. For example, for the case of 4 neurons and two inhibitory neurons on either end, note that one

needs $V/3^6 = 10^{-2}$, which amounts to choosing $V = 7.2$ MHz. Note that this value of V is of the order of magnitude of $\Omega = 4.2$ MHz and so the qRNN in this case is well into the non-classical regime.

The learnt parameters in the output linear map W^{out} which in this case is a matrix in $\mathbb{R}^{3+1 \times 1}$ with the last row representing a bias term. Note that the dimension of the map is so because only one neuron is measured but three functions have to be fitted.

2. Decision making

In classical RNNs tasks such as decision making and working memory require connectivity between all neurons. Since our connectivity is limited by physical constraints, an open 2D square lattice structure was chosen to prevent neurons from being isolated from the rest. Moreover, a 2D square lattice is experimentally friendly. In our case we use an open 2×3 lattice with the two input neurons being at the top-left corner of the chain, and two output neurons being the bottom-right corner. Again, this architecture was chosen so that the input neurons have to interact with the rest of the system before readout. We use $V = 2\pi \times 10$ MHz for our simulations, and choose $\Delta t = 2\pi/V$ as the time the inputs are turned on as that’s the timescale in which the input atoms entangle with the rest of the chain.

The inputs are uniformly sampled from $\{0, \pi/2, \pi, 3\pi/2, 2\pi\}$ (MHz) with added Gaussian noise $\sigma_{in} = 0.1$. In this task, the time that the stimuli are turned on Δt is fixed to a mean of $\langle \Delta t \rangle = 0.1 \mu s$ and with added Gaussian noise $\sigma_{in} = 0.1$. In this task we optimize over the linear output map W^{out} , a matrix in $\mathbb{R}^{1+1,2}$ since one function is fitted and two neurons are measured. Additionally, we train the output time t_{out} after the stimuli are turned off and before the network is probed to come up with an input that satisfied (15). To do the optimization, we make use of the Nelder-Mead algorithm [70].

In order to compute the psychometric response plotted in Fig. 6B, we measure the expectation values on the two output neurons and produce the vector $\mathbf{r}(\Delta_1^{in}, \Delta_2^{in}) = (\langle \sigma_{out1}^y \rangle, \langle \sigma_{out2}^y \rangle, 1)$ which depends on the inputs $\Delta_{1,2}^{in}$, as well as the temporal parameters $(\Delta t, t_{out})$. We then compute $y^{out}(\Delta_{1,2}^{in}) = W^{out} \cdot \mathbf{r}(\Delta_{1,2}^{in})$ and (W^{out}, t_{out}) are optimized such that $y^{out}(\Delta_{1,2}^{in}) \approx y^{targ}$ in 15. The optimization is done by generating about 40,000 different values of $\Delta_{1,2}^{in}$ of different levels of contrast $|\Delta_1^{in} - \Delta_2^{in}|$ ranging from 0 to 1 MHz. Once the optimization is done, we look at the loss towards Δ_2^{in} , which is obtained as the error in classifying Δ_2^{in} as greater than Δ_1^{in} when indeed $\Delta_2^{in} > \Delta_1^{in}$. The error is quantified using the mean-square loss in (4).

3. Working memory

This task's setup is identical to the decision making task except that the two inputs are separated by a delay time t_{delay} . The values of the interaction strength V used for Fig. 7 are $V = 2\pi \times 10$ MHz and $V = 2\pi \times 0.1$ MHz corresponding to $V/\Omega > 1$ and $V/\Omega < 1$ respectively. The former of which sets us in the Rydberg blockaded regime while the later is not. In this task, the times Δt and t_{delay} are fixed up to an added Gaussian with noise $\sigma_{in} = 0.1$. In this task we optimize over the linear output map W^{out} , a matrix in $\mathbb{R}^{1+1,2}$ since one function is fitted and two neurons are measured.

4. Long-term memory

Although quantum scars are known to exist in other geometries and dimensions [71], for this task we use a

1D chain of Rydberg atoms since for this case quantum many-body scars have been experimentally observed [35, 46]. Furthermore, our chain has periodic boundary conditions to avoid edge-effects. Since we know that scars are robust to decoherence, we set $\gamma_{420} = 0$ so that we can evolve our states for longer periods of time. The number of cycles n in Fig. 9 corresponds to n evolutions under the PXP hamiltonian for a time $2\tau = 1.51 \times \pi\Omega^{-1}$. In this case, we take $V \gg \Omega$ and renormalized $\Omega = 1$. The noisy field in (16) is sampled according to $\epsilon_k \sim N(\mu = 0.1, \sigma = 0.1)$. The input m is sampled as a fair random coin. Lastly, after each number of cycles n , the only trained parameter is $W_n^{out} \in \mathbb{R}^{1+1 \times 1}$ since only one atom is probed to calculate an answer as to the input m .