

Stochastic Quantum Spiking Neural Networks with Quantum Memory and Local Learning

Jiechen Chen, *Member, IEEE*, Bipin Rajendran, *Senior Member, IEEE*, and Osvaldo Simeone, *Fellow, IEEE*

Abstract—Neuromorphic and quantum computing have recently emerged as promising paradigms for advancing artificial intelligence, each offering complementary strengths. Neuromorphic systems built on spiking neurons excel at processing time series data efficiently through sparse, event-driven computation, consuming energy only upon input events. Quantum computing, on the other hand, operates on state spaces that grow exponentially in dimension with the number of qubits — as a consequence of tensor-product composition — with quantum states admitting superposition across basis states and entanglement between subsystems. Hybrid approaches combining these paradigms have begun to show potential, but existing quantum spiking models have important limitations. Notably, they implement classical memory mechanisms on single qubits, requiring repeated measurements to estimate firing probabilities, while relying on conventional backpropagation for training. In this paper, we propose a novel stochastic quantum spiking (SQS) neuron model that addresses these challenges. The SQS neuron uses multi-qubit quantum circuits to realize a spiking unit with internal quantum memory, enabling event-driven probabilistic spike generation in a single shot during inference. Furthermore, we study networks of SQS neurons, dubbed SQS neural networks (SQSNN), and demonstrate that they can be trained via a hardware-friendly local learning rule, eliminating the need for global classical backpropagation. The proposed SQSNN model is shown via experiments with both conventional and neuromorphic datasets to improve over previous quantum spiking neural networks, as well as over classical counterparts, when fixing the overall number of trainable parameters, highlighting its potential for event-driven applications such as neuromorphic integrated sensing and communications (N-ISAC).

Index Terms—Quantum neuromorphic computing, spiking neural networks, quantum machine learning, probabilistic spiking models, local learning rules.

I. INTRODUCTION

A. Context and Motivation

Over the past few decades, *quantum computing* and *neuromorphic computing* have emerged as two prominent paradigms for advancing computation beyond the limitations of conventional transistor-based digital architectures [1, 2]. Each of these technologies brings complementary advantages. *Quantum computing* operates on multi-qubit state spaces whose dimension scales exponentially with the number of qubits due to tensor product composition (see, e.g., [3]). Within

this exponentially large space, superposition and entanglement underlie transformations that may not be efficiently simulable on classical computers [4, 5].

Neuromorphic computing, on the other hand, is inspired by the operational principles of biological neural systems. It implements networks of spiking neurons that communicate via discrete, asynchronous events and remain inactive in the absence of input [6]. This event-driven, sparse mode of computation can lead to highly energy-efficient processing, particularly for time-series and streaming data [7]. Neuromorphic computing is highly synergistic with event-driven sensing, an energy-efficient form of sensing that produces information only when sufficiently large changes are detected in the scene. Furthermore, spiking neural networks (SNNs) naturally support online and continual learning, adapting to new information without retraining from scratch [6, 8–10].

Given these complementary strengths, a compelling question is whether combining quantum and neuromorphic approaches can yield new computational models that inherit the advantages of both domains [2, 11]. An ideal version of this integration would leverage quantum systems to provide exponentially large signal spaces or dynamical richness, while adopting neuromorphic principles to endow the computation with robustness, sparsity, and event-driven efficiency.

Recent works have begun exploring this intersection. For instance, quantum reservoir computing (QRC) leverages the complex dynamics of quantum systems as a reservoir for temporal information processing [12, 13]. For example, the work showed that a small network of only 6–7 qubits with disordered quantum evolution can achieve sequence learning performance on par with a classical recurrent neural network of hundreds of neurons [14–16].

Another line of work has adapted classical spiking neuron models to quantum hardware. Notably, prior work has introduced the quantum leaky integrate-and-fire neuron (QLIF) as a unit for the design of quantum spiking neural networks [11]. The QLIF neuron is implemented as a single-qubit that mimics the leaky integrate-and-fire behavior of classical spiking neurons [6]. By cascading QLIF neurons, this work demonstrated the first *quantum SNN* (QSNN), which was able to perform simple image recognition tasks such as MNIST classification.

This seminal work proved the viability of spiking neural networks on gate-based quantum computers. However, the QLIF neuron has several limitations from its design:

- *Single qubit and classical memory*: First, each QLIF neuron is essentially a classical LIF unit emulated on a single qubit, which restricts it to limited state dynamics characterized by classical state variables.

J. Chen is with the Department of Engineering, King's College London, London, WC2R 2LS, UK (email: jiechen.chen@kcl.ac.uk). B. Rajendran and O. Simeone are with the Institute for Intelligent Networked Systems, Northeastern University London, One Portsoken Street, London, E1 8PH, UK (email: {b.rajendran, o.simeone}@nulondon.ac.uk). This work was supported by the European Research Council (ERC) under the European Union's Horizon Europe Programme (grant agreement No. 101198347), by an Open Fellowship of the EPSRC (EP/W024101/1), and by the EPSRC project (EP/X011852/1).

The authors are grateful to Andrea Gentile (Pasqal) for the useful feedback and advice on hardware implementation.




Fig. 1. Schematic of an SQS neuron with $N = 3$ input-output qubits (green) and $N_{\text{mem}} = 3$ memory qubits (orange), receiving input from P presynaptic neurons. At each time t , incoming presynaptic spikes are weighted into input currents $\mathbf{z}_t = [z_t^1, \dots, z_t^N]$. The currents determine rotation angles ϕ_t applied to the input qubits (blue gates). A parameterized quantum circuit (PQC, gray box) with trainable parameters θ entangles input and memory qubits, updating the neuron's state. The N input qubits are then measured (meter symbols) to produce output spikes $\mathbf{s}_t$ (red arrows), after which those qubits are reset to the ground state $|0\rangle$. The memory qubits are not measured, preserving quantum information (orange shading) that influences subsequent time steps. This internal quantum memory allows the SQS neuron to exhibit complex temporal dynamics, while not requiring multi-shot measurements.

- *Multi-shot measurements:* The QLIF neuron's internal state, analogous to a membrane potential, is encoded as the excitation probability of one qubit, and a threshold rule is used to determine spike emission. In practice, determining whether the neuron should fire requires estimating this excitation probability at each time step via repeated quantum measurements (multiple shots) and comparing against a threshold. This multi-shot measurement scheme adds significant overhead, as the quantum state must be prepared and measured many times to obtain the spiking probability with high confidence. This means the spiking outcome is not generated within a single quantum evolution, but rather via a classical post-processing of measurement statistics.
 - *Classical training via backpropagation:* Third, the QLIF network training process remains classical, leveraging backpropagation through time on a classical simulation of the network. This reliance on classical, global learning rules makes it challenging to deploy the model directly on quantum hardware for online learning as is the case for classical neuromorphic models [6].
- Overall, existing hybrid models like QRC and QLIF validate the concept of quantum neuromorphic systems, but they either treat the quantum device as a black-box dynamical system or implement spiking neurons in a way that closely resembles classical models. There is a clear need for new designs that fully exploit quantum dynamics while supporting efficient operation and learning on hardware.
- ### B. Main Contributions
- In this work, we address the limitations of the QLIF neuron discussed above by proposing the *Stochastic Quantum Spiking* (SQS) neuron model illustrated in Fig. 1, along with the corresponding *SQS neural networks* (SQSNNs). The main features of SQS models and SQSNNs are as follows:
- *Multi-qubit neuron architecture with quantum memory:* Each SQS neuron is realized using a multi-qubit quantum circuit, allowing the neuron to exploit entangled states and richer internal dynamics. This design increases the expressive capacity of each neuron and ensures that the model truly leverages the unique features of quantum hardware. Specifically, the SQS neuron employs an internal quantum memory mechanism inspired by *quantum recurrent neural networks* (QRNNs) [17, 18]. Accordingly, the neuron's multi-qubit circuit carries over quantum information from one time instant to the next, preserving relevant memory of the neuron's previous inputs and activity. Unlike standard QRNNs, however, the proposed SQS neuron uses quantum memory to realize stochastic spike generation from presynaptic spike events, enabling event-driven communication and local learning. Furthermore, SQSNNs support arbitrary connectivity among SQS neurons, making it possible to deploy flexible architectures compatible with state-of-the-art classical SNNs.
 - *Event-driven, single-shot spike generation:* The SQS neurons produces output spikes *probabilistically* in an event-driven manner based on a *single-shot measurement* of the multi-qubit state. In contrast to QLIF's multi-shot probability estimation, an SQS neuron emits spikes within one quantum evolution, with the spike event occurring stochastically according to the Born rule. This event-driven operation aligns with neuromorphic principle that no computation should be carried out when no events occurs, while avoiding costly repeated measurements.
 - *Hardware-compatible local learning rule:* We develop a novel learning rule for networks of SQS neurons that updates synaptic weights using only locally available information at each SQS neuron, thus avoiding the need for backpropagation on a classical simulator. This bio-inspired learning scheme is designed to run natively on

quantum hardware, adjusting connection strengths based on stochastic spike correlations in a manner akin to three-factor learning rules in neuromorphic computing [19].

Overall, the SQS neuron model harnesses quantum resources, including multiple qubits, entanglement, and inherent measurement randomness, to implement a spiking neuron that operates in an event-driven, probabilistic fashion like a biological neuron. In this regard, SQSNNs can be viewed as examples of dynamic PQC, which have been recently shown to have significant advantages in terms of expressivity and trainability [20].

At a practical level, the proposed neural architecture is particularly well suited for hardware platforms characterized by: (i) *low repetition rate*, i.e., high shot-to-shot latency, favoring single-shot measurements [21]; and (ii) support for *mid-circuit measurements* [22], which are required to implement the QRNN-based quantum memory and readout mechanisms. Both properties apply to leading candidate platforms for scalable quantum computing, namely trapped ions and neutral atoms [23–26]. Furthermore, in practice, the implementation of SQSNNs is subject to the connectivity constraints of the underlying quantum hardware. For example, superconducting quantum processors typically require hardware-aware routing and compilation strategies due to fixed qubit connectivity [27], whereas trapped-ion and neutral-atom platforms offer more flexible qubit interactions and reconfiguration capabilities [28, 29].

Experimental results compare different models by fixing the same number of trainable parameters across all models. This has become a standard benchmarking methodology in quantum machine learning (see, e.g., [30–32]), as the number of parameters directly relates to key model properties such as trainability, complexity, and expressivity [31]. Our results demonstrate the effectiveness of SQSNNs, and of the local training rule on quantum hardware, as well as its superior accuracy compared to QLIF on a larger-scale dataset trained using backpropagation on a classical simulator.

The rest of the paper is organized as follows. Section II reviews the QLIF neuron proposed by reference [11]. The SQS neuron is proposed in Section III, while Section IV presents a local training rule of SQS-based neural networks. Experimental setting and results are described in Section V. Finally, Section VI concludes the paper.

II. PRIOR WORK: QUANTUM LEAKY INTEGRATE-AND-FIRE NEURON

In this section, we review the *quantum leaky integrate-and-fire* (QLIF) neuron proposed by [11]. As illustrated in Fig. 2, a QLIF neuron is modeled via an internal single qubit, which evolves over discrete time $t = 1, 2, \dots$ through a *measure-and-prepare* mechanism characterized by *multi-shot* measurements. The state of the qubit serves as the inner state of the spiking neurons. Furthermore, a rotation angle φ_t determines whether the neuron produces a spike or not at time t through a measurement probability evaluated via multiple shots. Consider a QLIF neuron that is fed by P pre-synaptic neurons, at each time t , the QLIF neuron receives

an input binary signal $\mathbf{x}_t = [x_t^1, \dots, x_t^P]^T \in \{0, 1\}^P$, with $x_t^p \in \{0, 1\}$ representing the binary activation of the p -th presynaptic neuron. A signal $x_t^p = 1$ indicates the presence of a spike, and $x_t^p = 0$ indicates that the p -th presynaptic neuron does not spike at time t .

The QLIF neuron connects to P presynaptic neurons via two vectors of weights $\mathbf{w}^s = [w_1^s, \dots, w_P^s]^T$ and $\mathbf{w}^o = [w_1^o, \dots, w_P^o]^T$. Specifically, the pre-synaptic signals are aggregated into the input signals

$$z_t^s = \sum_{p=1}^P w_p^s x_t^p, \quad (1)$$

and

$$z_t^o = \sum_{p=1}^P w_p^o x_t^p. \quad (2)$$

As in classical neuromorphic computing, the complexity of evaluating the signals (1) and (2) is proportional to the number of pre-synaptic spikes. Therefore, sparser input signals reduce the complexity of computing (1) and (2).

The signals (1) and (2) are used to determine the state of the internal qubit of the QLIF neuron through a rotation angle φ_t . In particular, at each time t , the internal qubit state $|\psi_t\rangle$ is prepared from the ground state $|0\rangle$ via the Pauli X-rotation gate

$$R_X(\varphi_t) = \begin{bmatrix} \cos(\varphi_t/2) & -i \sin(\varphi_t/2) \\ -i \sin(\varphi_t/2) & \cos(\varphi_t/2) \end{bmatrix} \quad (3)$$

as $|\psi_t\rangle = R_X(\varphi_t)|0\rangle$, where the angle φ_t depends on the input signals (1)-(2) as discussed below. A standard measurement of the qubit yields output 1 with probability

$$\alpha_t = |\langle 1 | \psi_t \rangle|^2 = (\sin(\varphi_t/2))^2. \quad (4)$$

This probability is estimated by preparing and measuring state $|\psi_t\rangle$ multiple times. The probability (4) represents the counterpart of the membrane potential of classical spiking neurons [19, 33]. In fact, using the estimated probability α_t , a spike is generated if the probability α_t exceeds a predetermined threshold $\vartheta \in [0, 1]$, i.e.,

$$s_t = \begin{cases} 0 \text{ (no spike)}, & \text{if } \alpha_t \leq \vartheta, \\ 1 \text{ (spike)}, & \text{if } \alpha_t > \vartheta. \end{cases} \quad (5)$$

This signal serves as an input to the post-synaptic neurons connected to the given QLIF neuron for time $t + 1$.

We now describe the evaluation of the phase φ_t used in the rotation (3). First, at time t , the phase φ_t is first initialized as

$$\varphi_t^o = \begin{cases} 2 \arcsin(\sqrt{\alpha_t}), & \text{if } s_{t-1} = 0, \\ 0, & \text{if } s_{t-1} = 1. \end{cases} \quad (6)$$

Using (4), this operation recovers the phase φ_{t-1} if no spike is generated at time $t - 1$, i.e., if $s_{t-1} = 0$. Otherwise, if the QLIF had spiked at the previous time, i.e., $s_{t-1} = 1$, the phase is reset to zero, i.e., $\varphi_t^o = 0$. Having initialized the phase φ_t^o , the rotation angle φ_t is computed as

$$\varphi_t = \varphi_t^o + \phi_t, \quad (7)$$




Fig. 2. Schematic of the QLIF neuron with P presynaptic inputs [11]. A single qubit (green) serves as the neuron’s internal state, evolving in discrete time steps via a measure-and-prepare scheme. At each time t , input spikes from P presynaptic neurons are weighted (blue arrows) and encoded as a rotation ϕ_t on the qubit. The qubit is measured multiple times to estimate a spike probability (red arrow), which is compared against a threshold to decide if the neuron fires. This design requires repeated quantum measurements and uses a classical memory (accumulated threshold voltage), mimicking a leaky integrate-and-fire process.

where the angle ϕ_t depends on the signals (1)-(2) computing from the input spikes. Note that we set $\phi_0^o = 0$, so that the qubit state at time $t = 1$ is $|\psi_1\rangle = R_X(\phi_1)|0\rangle = R_X(\phi_1)|0\rangle$.

When the input current z_t^s is positive, the rotation gate applies a positive angle $\phi_t > 0$, exciting the membrane potential and increasing the spiking probability (4), while otherwise a rotation with a negative angle $\phi_t \leq 0$ is applied, reducing the spiking probability (4). Thus a positive input current z_t^s causes the QLIF neuron to integrate its inputs, while a negative input current $z_t^s \leq 0$ entails the “leaking” of the membrane potential (4). Specifically, given the previous spiking probability α_{t-1} , the angle ϕ_t is determined as

$$\phi_t = \begin{cases} f_s(z_t^s), & \text{if } z_t^s > 0, \\ f_o(z_t^o), & \text{if } z_t^s \leq 0, \end{cases} \quad (8)$$

with

$$f_s(z_t^s) = \pi z_t^s \quad (9)$$

and

$$f_o(z_t^o) = -2 \arcsin(\sqrt{\alpha_{t-1} \exp(-\beta |z_t^o|/T_1)}), \quad (10)$$

where $\beta > 0$ is a scaling factor and T_1 is the parameter that controls the rate of decay towards the ground state $|0\rangle$. Referring to [11] for details, function $f_s(z_t^s)$ describes a linear integration of positive current into the rotation (8), whereas function $f_o(z_t^o)$ in (10) implements leakage by enforcing an exponential decay of the previous excitation probability α_{t-1} with time constant controlled by parameter T_1 . Reference [11] optimizes the weights w^s and w^o by using the arctan surrogate gradient for the spike generation process (5) via backpropagation on a classical computer.

III. STOCHASTIC QUANTUM SPIKING NEURON MODEL

As illustrated in Fig. 1, a SQS neuron can process and produce N spikes, acting as a multiple-input multiple-output unit. As the QLIF neuron reviewed in the previous section, the SQS neuron is composed of a classical logic module that encodes incoming data into quantum rotation gates. However,

rather than relying on a classical memory, an SQS neuron encompasses a QRNN module [17, 18] that plays a role of *quantum memory* through an event-driven parameterized quantum circuit (PQC). Furthermore, unlike the QLIF neuron, the SQS neuron only requires single-shot measurements, thus foregoing the need for multiple measurement-and-prepare steps. This enables the construction of quantum SNN models in which inference requires single-shot measurements at all neurons, with the neurons communicating via spiking signals.

A. Architecture of the SQS Neuron

As illustrated in Fig. 1, the SQS neuron encompasses two groups of qubits: N input-output qubits and N_{mem} memory qubits for a total of $N_{\text{tot}} = N + N_{\text{mem}}$ qubits. The input-output qubits are reset to the ground state $|0\rangle^{\otimes N}$ at the beginning of each time step t , processing the incoming signals from the pre-synaptic neurons and producing output spikes via measurement. The memory qubits evolve over time jointly with the input-output qubits to serve as memory for the SQS neuron. Unlike LIF or QLIF spiking neurons, which maintain a scalar membrane potential governed by first-order integrate-and-leak dynamics, the proposed SQS neuron efficiently evolves a state vector on a $2^{N_{\text{tot}}}$ -dimensional Hilbert space through a trainable unitary transformation, potentially enabling a richer class of temporal state transitions. In this sense, the SQS neuron is aligned with recent advances in the design of classical spiking neurons that support complex dynamics via the integration with state-space models (SSMs) [34].

As shown in Fig. 1, at each discrete time $t = 1, \dots, T$, a PQC defined by a parameterized unitary transformation $U(z_t, \theta)$ operates on the register of $N + N_{\text{mem}}$ qubits, where θ are parameters of the PQC to be optimized. Assuming that the neuron is connected to P pre-synaptic neurons, at each time t , the neuron processes NP incoming spiking signals $\mathbf{x}_t = [\mathbf{x}_{t,1}, \dots, \mathbf{x}_{t,P}]^T$ with $\mathbf{x}_{t,p} = [x_{t,p}^1, \dots, x_{t,p}^N]^T \in \{0, 1\}^N$ representing the N spiking signals produced by the p -th pre-synaptic neuron. The SQS neuron itself produces N spiking

signals $\mathbf{s}_t = [s_t^1, \dots, s_t^N]^T \in \{0, 1\}^N$ by measuring the N input-output qubits through the following steps:

- *Quantum embedding*: Convert the $NP \times 1$ input signals $\mathbf{x}_t$ into $N \times 1$ input currents $\mathbf{z}_t = [z_t^1, \dots, z_t^N]$, and embed the input currents $\mathbf{z}_t$ into the state of the input-output qubits via rotation gates;
- *Quantum memory retrieval*: Following the QRNN model, evolve jointly input-output qubits and memory qubits;
- *Single-shot measurements*: Measure the N input-output qubits to produce the $N \times 1$ spiking signals $\mathbf{s}_t$;

Importantly, while the spiking output of the QLIF neuron is determined on the basis of multi-shot measurements via the thresholding function (5), the SQS neuron is run only once, producing a random N -dimensional spiking output on the basis of a single-shot measurement.

These steps are elaborated as in the rest of this section.




Fig. 3. Illustration of the connection between a set of P presynaptic neurons and the SQS neuron of interest. The n -th input current z_t^n of the SQS neuron is computed as the weighted sum of the corresponding spiking signal $\{x_{t,p}^n\}_{p=1}^P$ from the P presynaptic neurons as in equation (11).

B. Quantum Embedding

As illustrated in Fig. 3, denote the weights between the n -th output of all pre-synaptic neurons and the neuron under study as $\mathbf{w}^n = [w_1^n, \dots, w_P^n]$, the spikes $\mathbf{x}_t = [x_{t,1}, \dots, x_{t,P}]$ from the P pre-synaptic neurons are weighed to produce the N input currents $\mathbf{z}_t = [z_t^1, \dots, z_t^N]^T$. Writing the P spiking signals from all pre-synaptic neurons' n -th output as $\mathbf{x}_t^n = [x_{t,1}^n, \dots, x_{t,P}^n]^T$, the n -th element of the current, z_t^n , is evaluated as

$$z_t^n = \mathbf{w}^n \mathbf{x}_t^n. \quad (11)$$

As for (1) and (2), the complexity of evaluating (11) is proportional to the number of spikes generated by the pre-synaptic neurons.

Furthermore, the input currents $\mathbf{z}_t$ are then embedded into rotation angles $\phi_t = [\phi_t^1, \dots, \phi_t^N]^T$, with the n -th angle given by

$$\phi_t^n = \begin{cases} f_s(z_t^n), & \text{if } z_t^n > 0, \\ 0, & \text{if } z_t^n \leq 0, \end{cases} \quad (12)$$

with $f_s(z_t^n) = \pi z_t^n$ for all $n = 1, \dots, N$. Each angle ϕ_t^n determines the rotation $R_X(\phi_t^n)$ applied to the n -th input-output qubit. The synaptic weights $\mathbf{w}^n$ are real-valued learnable parameters. Unlike for the QLIF neuron (see (6)), the angle ϕ_t^n of the rotations $R_X(\phi_t^n)$ does not implement any memory mechanisms, since memory effects are accounted for by the memory qubit, as discussed below.

By (12), when the current $z_t^n \leq 0$, the corresponding qubit does not apply any rotation, reducing the number of internal quantum operations required. Therefore, sparsity in the spiking signals $\mathbf{x}_t^n$ reduces the complexity of both the classical logic evaluating (12) and the quantum operations applied to the neuron's qubit.

C. Quantum Memory Retrieval

As discussed in the previous subsection, at each time t , the input-output qubits are initialized to the ground state $|0\rangle^{\otimes N}$ and are applied the rotation gates $\{R_X(\phi_t^n)\}_{n=1}^N$ to embed the input currents $\mathbf{z}_t$ via (11) and (12). These input-output qubits, along with the memory qubits in the second group, are processed by a PQC defined by a unitary matrix $V(\theta)$. Accordingly, the overall unitary transformation $U(\mathbf{z}_t, \theta)$ applied by the SQS neuron to the initial state of the $(N + N_{\text{mem}})$ -qubit register can be expressed as

$$U(\mathbf{z}_t, \theta) = V(\theta) \cdot \left(\bigotimes_{n=1}^N R_X(\phi_t^n) \otimes I \right), \quad (13)$$

where I is a $2^{N_{\text{mem}}} \times 2^{N_{\text{mem}}}$ identity matrix.

After joint transformation (13), the input-output qubits are measured to generate the output spiking signal $\mathbf{s}_t \in \{0, 1\}^N$. In contrast, the memory qubits are not measured, retaining information across time steps.

D. Single-Shot Measurements

At the initial time $t = 1$, all the N_{tot} qubits are all initialized to the ground state $|0\rangle$. Denote as ρ_t^{IM} the density matrix of the $N + N_{\text{mem}}$ qubits at the beginning of time interval t and as

$$\rho_t^{\text{I}} = \text{Tr}_{\text{M}}(\rho_t^{\text{IM}}) \quad (14)$$

the corresponding reduced density matrix for the input-output qubits.

By Born's rule, the probability of producing the output signals $\mathbf{s}_t \in \{0, 1\}^N$ is given by

$$p(\mathbf{s}_t | \rho_t^{\text{I}}) = \langle \mathbf{s}_t | \rho_t^{\text{I}} | \mathbf{s}_t \rangle, \quad (15)$$

where $|\mathbf{s}_t\rangle = |s_t^1, \dots, s_t^N\rangle$.

Furthermore, as the input-output qubits collapse onto the computational state $|\mathbf{s}_t\rangle$, the post-measurement state of the memory qubits is described by the density matrix

$$\sigma_t^{\text{mem}} = \frac{(\langle \mathbf{s}_t | \otimes I) \rho_t^{\text{IM}} (|\mathbf{s}_t\rangle \otimes I)}{\langle \mathbf{s}_t | \rho_t^{\text{I}} | \mathbf{s}_t \rangle}, \quad (16)$$

where I is a $2^{N_{\text{mem}}} \times 2^{N_{\text{mem}}}$ identity matrix. Accordingly, upon the application of the given unitary matrix (13), the density matrix of all the qubits at the following time step $t + 1$ is

$$\rho_{t+1}^{\text{IM}} = U(\mathbf{z}_t, \theta) (|0\rangle\langle 0|^{\otimes N} \otimes \sigma_t^{\text{mem}}) U(\mathbf{z}_t, \theta)^\dagger. \quad (17)$$

Through the relation (17), as anticipated, the state σ_t^{mem} of the second group of qubits serves as memory.

IV. LOCAL TRAINING OF SQS-BASED NEURAL NETWORKS

Similar to conventional neural models and to the QLIF neuron, the SQS neuron proposed in the previous section can be structured into a neural network that can be trained to perform specific tasks. In this section, we first describe a general network architecture based on SQS neurons, and then introduce a training rule designed to run effectively on quantum hardware. As in neuromorphic systems [6], the training rule is *local*, in the sense that updates to the parameters can be evaluated at each neuron based on information available locally, along with a scalar global feedback signal [19, 35]. Technically, this result is obtained by leveraging the probabilistic, autoregressive, structure of the model induced by a network of SQS neurons. The local training rule builds on standard zero-th order perturbation-based gradient estimates [36] within each neuron, while enforcing cooperation across all neurons via global feedback signals.

A. SQS-based Neural Network Model

As shown in Fig. 4, an *SQS-based neural network* (SQSNN) consists of a directed graph $(\mathcal{S}, \mathcal{E})$ encompassing a set $\mathcal{S}$ of SQS neurons. The neurons are connected via directed edges described by a set $\mathcal{E} \subseteq \{(i, j) \text{ with } i, j \in \mathcal{S} \text{ and } i \neq j\}$ of ordered pairs (i, j) . Accordingly, each neuron i within the SQSNN receives inputs from a set $\mathcal{P}_i$ of pre-synaptic neurons, with $\mathcal{P}_i = \{j \in \mathcal{S} : j \neq i \text{ and } (j, i) \in \mathcal{E}\}$. The number of pre-synaptic neurons for neuron i is thus $P_i = |\mathcal{P}_i|$.

Following the literature on neuromorphic systems (see, e.g., [19, 37]), we partition the SQS neurons in the SQSNN into two disjoint subsets: the subset $\mathcal{O}$ of output neurons and the subset $\mathcal{H}$ of hidden neurons, with $\mathcal{S} = \mathcal{O} \cup \mathcal{H}$. The spiking signals of the output neurons provide the outputs of the SQSNN. In contrast, the spiking signals of the hidden neurons serve as internal representations guiding the output neurons towards producing informative signals.

Furthermore, as illustrated by the example in Fig. 4, the output neurons do not have any outgoing edges, representing the last layer of the network. Mathematically, we have the condition

$$\mathcal{P}_i \cap \mathcal{O} = \emptyset, \quad (18)$$

or equivalently $\mathcal{P}_i \cap \mathcal{H} = \mathcal{P}_i$, for all neurons $i \in \mathcal{S}$.

Being based on an arbitrary directed graph connecting hidden and output neurons, the architecture of the proposed SQSNN is aligned and compatible with existing SNN architectures [19]. Accordingly, SQSNN models can be obtained by replacing classical LIF neurons with SQS neurons, while preserving the original network topology, thus suggesting a natural design strategy for SQSNN architectures.

As a final note, we observe that when the PQC of all neurons are fixed and not trained, the SQSNN architecture introduced here can be interpreted as a form of quantum reservoir computing (QRC) with mid-circuit measurements




Fig. 4. Illustration of an SQS-based neural network (SQSNN) with six neurons ($\mathcal{S} = \{1, \dots, 6\}$). Edges $(i \rightarrow j)$ indicate neuron i feeds into neuron j . Neurons $\{1, 2, 3, 4, 5\}$ are hidden, and neuron 6 is the only output neuron ($\mathcal{O} = \{6\}$). Output neurons have no outgoing edges. Each node is an SQS neuron as in Fig. 1, with a self-loop denoting dependence on its own past spikes. The network's joint spike probability factorizes along these connections: e.g., neuron 6's spiking at time t depends only on the past spikes of neurons 4, 5 (its parents, denoted as $\mathcal{P}_6$) and on its own past spikes.

[15, 16, 38]. In it, the neurons serve as module of the reservoir with classical communication between modules, defined via mid-circuit measurements.

B. Probabilistic Autoregressive Model

As described in Section IV-A, at each time t , each (hidden or output) neuron i , receives the spiking signals $\mathbf{s}_{\mathcal{P}_i, t} \in \{0, 1\}^{N_{\mathcal{P}_i}}$ from its pre-synaptic neurons in set $\mathcal{P}_i$, and produces a spiking vector-valued signal $\mathbf{s}_{i, t} \in \{0, 1\}^N$. Denote the past spiking output of any neuron $i \in \mathcal{S}$ up to time $t - 1$ as $\mathbf{s}_{i, t-1}^t = [\mathbf{s}_{i, 1}, \dots, \mathbf{s}_{i, t-1}]$, and the collection of all past spiking outputs of any subset $\mathcal{S}' \subseteq \mathcal{S}$ of neurons as $\mathbf{s}_{\mathcal{S}', t-1}^t = \{\mathbf{s}_{i, t-1}^t\}_{i \in \mathcal{S}'}$. Following the SQS model introduced in Section III, the signal $\mathbf{s}_{i, t}$ produced by neuron i at any time t depends on the past spiking outputs $\mathbf{s}_{\mathcal{P}_i \cup \{i\}, t-1}^t$ from itself and from the pre-synaptic neurons $\mathcal{P}_i$.

Specifically, the dependence of the spiking history $\mathbf{s}_{\mathcal{P}_i \cup \{i\}, t-1}^t$ is reflected by the expression (15) of the spiking probability as a function of the neuron's state $\rho_{i, t}^I$ of the input-output qubits of neuron i . In fact, by the dynamic update (17), the state $\rho_{i, t}^I$ of the input-output qubits of neuron i evolves as a function of the sequence $\mathbf{s}_{\mathcal{P}_i \cup \{i\}, t-1}^t$ through the currents (11). Accordingly, the spiking probability of neuron i at time t , when conditioned on the sequence $\mathbf{s}_{\mathcal{P}_i \cup \{i\}, t-1}^t$, can be written using (15) as

$$p_{\Theta_i}(\mathbf{s}_{i, t} | \mathbf{s}_{\mathcal{P}_i \cup \{i\}, t-1}^t) = \langle \mathbf{s}_{i, t} | \rho_{i, t}^I | \mathbf{s}_{i, t} \rangle. \quad (19)$$

The notation $p_{\Theta_i}(\mathbf{s}_{i, t} | \mathbf{s}_{\mathcal{P}_i \cup \{i\}, t-1}^t)$ emphasizes the dependence of the probability (19) on the local parameters $\Theta_i = \{\mathbf{w}_i, \boldsymbol{\theta}_i\}$ of neuron i . These include the synaptic weights $\mathbf{w}_i = \{\{w_{i, p}^n\}_{p=1}^{P_i}\}_{n=1}^N$ between the N input-output qubits of the P_i pre-synaptic neurons and the corresponding N qubits of neuron i , as well as the parameters $\boldsymbol{\theta}_i$ of the PQC in neuron i . Indeed, the state $\rho_{i, t}^I$ is obtained via (14) and (17), where

the unitary transformation $U(z_{i,t}, \theta_i)$ in (13) depends on the variational parameters θ_i .

Denote as $\Theta = \{\Theta_i\}_{i \in \mathcal{S}}$ the collection of all model parameters. By applying the chain rule, the joint spiking probability $\mathbf{s}^T = [s_1^T, \dots, s_{|\mathcal{S}|}^T]$ of all neurons in the SQSNN up to time T can be factorized over time index t and over neuron i as

$$\begin{aligned} p_{\Theta}(\mathbf{s}^T) &= \prod_{t=1}^T p_{\Theta}(\mathbf{s}_t | \mathbf{s}^{t-1}) \\ &= \prod_{t=1}^T \prod_{i \in \mathcal{S}} p_{\Theta_i}(s_{i,t} | \mathbf{s}_{\mathcal{P}_i \cup \{i\}}^{t-1}). \end{aligned} \quad (20)$$

This corresponds to a probabilistic autoregressive sequence model akin to probabilistic spiking models studied in, e.g., [19, 37].

Denote the spiking signals emitted by the output neurons as $\mathbf{o}_t = \{\mathbf{o}_{i,t}\}_{i \in \mathcal{O}}$, with $\mathbf{o}_{i,t} = [o_{i,t}^1, \dots, o_{i,t}^N]$, while the spiking signals emitted by the hidden neurons are written as $\mathbf{h}_t = \{\mathbf{h}_{i,t}\}_{i \in \mathcal{H}}$, with $\mathbf{h}_{i,t} = [h_{i,t}^1, \dots, h_{i,t}^N]$. As mentioned, the signals $\mathbf{o}_t$ correspond to the model output (see Fig. 4). We also denote as $\Theta^O = \{\Theta_i\}_{i \in \mathcal{O}}$ and $\Theta^H = \{\Theta_i\}_{i \in \mathcal{H}}$ the sets of local parameters for the observed and hidden neurons, respectively. With these definitions, given that, by (18), the output neurons have no outgoing edges, the probability (20) can be expressed as

$$p_{\Theta}(\mathbf{s}^T) = \prod_{t=1}^T p_{\Theta^H}(\mathbf{h}_t | \mathbf{h}^{t-1}) p_{\Theta^O}(\mathbf{o}_t | \mathbf{o}^{t-1}, \mathbf{h}^{t-1}), \quad (21)$$

where the spiking probability $p_{\Theta^H}(\mathbf{h}_t | \mathbf{h}^{t-1})$ of the hidden neurons does not depend on the spiking probability of the output neurons due to the assumption (18).

C. Maximum Likelihood Learning

To train an SQSNN, we assume the availability of a dataset $\mathcal{D} = \{\mathbf{o}^T\}$ of sequences $\mathbf{o}^T = [\mathbf{o}_1, \dots, \mathbf{o}_T]$, where each sequence $\mathbf{o}^T$ corresponds to an example of desired behavior for the output neurons. Note that the sequence $\mathbf{o}^T$ in the training dataset is used to supervise the spiking activity of the output neurons $i \in \mathcal{O}$, while the spiking behavior of the hidden neurons is left unspecified by the data. In practice, the data typically also specifies the (classical) inputs to a number of neurons [19, 37]. Examples can be found in [19, 37] and in Section V. The inclusion of inputs is not made explicit in set $\mathcal{D}$ to simplify the notation.

By the standard *maximum likelihood* (ML) criterion, we aim at minimizing the empirical average log-loss on data set $\mathcal{D}$, yielding the optimized model parameters

$$\Theta^* = \arg \min_{\Theta} \left\{ \mathcal{L}_{\mathcal{D}}(\Theta) = \sum_{\mathbf{o}^T \in \mathcal{D}} \mathcal{L}_{\mathbf{o}^T}(\Theta) \right\}, \quad (22)$$

with the per-data-point log-loss

$$\begin{aligned} \mathcal{L}_{\mathbf{o}^T}(\Theta) &= -\log p_{\Theta}(\mathbf{o}^T) \\ &= -\log \sum_{\mathbf{h}^T} p_{\Theta}(\mathbf{s}^T), \end{aligned} \quad (23)$$

where $\sum_{\mathbf{h}^T}$ denotes the sum over all possible $2^{|\mathcal{H}|}$ values $\mathbf{h}^T$ taken by the signals produced by hidden neurons. By (23), evaluating the log-loss $\mathcal{L}_{\mathbf{o}^T}(\Theta)$ requires marginalizing over the spiking signals of the hidden neurons.

A direct application of standard gradient-based methods in quantum machine learning to problem (22) is not feasible. In fact, this would require estimating the probabilities $p_{\Theta}(\mathbf{s}^T)$ in (23) for all possible realization $\mathbf{h}^T$ of the hidden neurons. To address this problem, we proceed in two steps. First, we introduce a bound on the log-loss (23) that enables Monte Carlo estimation, removing the need for the exponential number of network runs required to evaluate the sum in (23). Second, using the bound, we derive a local learning rule based on local perturbation-based gradient estimates and a global scalar feedback signal.

D. Estimating a Likelihood Bound

To derive the proposed training method, we start by introducing an upper bound on the log-likelihood. As in (23), marginalizing the joint distribution (21), the marginal likelihood over hidden trajectories can be written as

$$\begin{aligned} p_{\Theta}(\mathbf{o}^T) &= \sum_{\mathbf{h}^T} \prod_{t=1}^T p_{\Theta^H}(\mathbf{h}_t | \mathbf{h}^{t-1}) p_{\Theta^O}(\mathbf{o}_t | \mathbf{o}^{t-1}, \mathbf{h}^{t-1}) \\ &= \prod_{t=1}^T \mathbb{E}_{p_{\Theta^H}(\mathbf{h}_t | \mathbf{h}^{t-1})} [p_{\Theta^O}(\mathbf{o}_t | \mathbf{o}^{t-1}, \mathbf{h}^{t-1})]. \end{aligned} \quad (24)$$

As a result, the negative log-likelihood can be expressed as

$$\begin{aligned} \mathcal{L}_{\mathbf{o}^T}(\Theta) &= -\sum_{t=1}^T \log \mathbb{E}_{p_{\Theta^H}(\mathbf{h}_t | \mathbf{h}^{t-1})} [p_{\Theta^O}(\mathbf{o}_t | \mathbf{o}^{t-1}, \mathbf{h}^{t-1})] \\ &\leq -\sum_{t=1}^T \mathbb{E}_{p_{\Theta^H}(\mathbf{h}_t | \mathbf{h}^{t-1})} [\log p_{\Theta^O}(\mathbf{o}_t | \mathbf{o}^{t-1}, \mathbf{h}^{t-1})] \\ &= L_{\mathbf{o}^T}(\Theta), \end{aligned} \quad (25)$$

where the second line follows from Jensen's inequality [19]. Using the factorization of the conditional spike distribution in (20), the upper bound $L_{\mathbf{o}^T}(\Theta)$ in (25) can be rewritten as

$$\begin{aligned} L_{\mathbf{o}^T}(\Theta) &= -\sum_{t=1}^T \sum_{i \in \mathcal{O}} \mathbb{E}_{p_{\Theta^H}(\mathbf{h}_t | \mathbf{h}^{t-1})} [\log p_{\Theta_i^O}(\mathbf{o}_{i,t} | \mathbf{o}_{\mathcal{P}_i \cup \{i\}}^{t-1})] \\ &= -\sum_{t=1}^T \sum_{i \in \mathcal{O}} \mathbb{E}_{p_{\Theta^H}(\mathbf{h}_t | \mathbf{h}^{t-1})} [\log (\text{Tr}(\rho_{i,t}^{\mathbf{I}} O_{i,t}))], \end{aligned} \quad (26)$$

where $\text{Tr}(\cdot)$ is the trace operation, and the observable $O_{i,t}$ is defined as

$$O_{i,t} = |\mathbf{o}_{i,t}\rangle \langle \mathbf{o}_{i,t}|. \quad (27)$$

The key advantages of the upper bound $L_{\mathbf{o}^T}(\Theta)$ in (26) is that it can be approximated via a Monte-Carlo estimate obtained by running forward passes on the model, with each forward pass producing random samples $\mathbf{h}_t \sim p_{\Theta^H}(\mathbf{h}_t | \mathbf{h}^{t-1})$ for the spiking signals of the hidden neurons.

We are interested in minimizing the upper bound on the training loss $\mathcal{L}_{\mathcal{D}}(\Theta)$ in (22) given by $L_{\mathcal{D}}(\Theta) = \sum_{o^T \in \mathcal{D}} L_{o^T}(\Theta)$, i.e.,

$$\Theta^* = \arg \min_{\Theta} \left\{ L_{\mathcal{D}}(\Theta) = \sum_{o^T \in \mathcal{D}} L_{o^T}(\Theta) \right\}. \quad (28)$$

E. Local Zero-th Order Learning

Given the objective (28), a conventional-based scheme would estimate the gradient $\nabla_{\Theta} L_{o^T}(\Theta)$ via local perturbations of *all* parameters [39]. However, this would require a number of *global* forward passes through the entire network that grows *linearly* with the number $|\Theta|$ of parameters in Θ [39]. This becomes quickly impractical as the number of neurons grows.

To address the computational and scalability challenges of conventional zeroth-order training, in this subsection, we propose a local learning rule in which each neuron updates its parameters using a single set of M *global* forward trajectories, followed by *local* forward passes at each neuron for gradient estimation. Through this method, the number of global forward executions is reduced from a quantity proportional to the number of parameters $|\Theta|$ to the constant M . This relocation from global passes to local computations facilitates scalable training while preserving correct gradient estimation.

As illustrated in Fig. 5, given the current iterate Θ and a data point o^T , we begin by running the SQSNN with parameters Θ for M forward passes to collect M samples $h^T(1), \dots, h^T(M)$ of the hidden spiking signals. After this initial phase, each neuron updates its local parameters using zeroth-order perturbation methods that require only *local* forward passes within each neuron. Finally, the model parameters are updated at each neuron based on the results of the local forward passes and the *global feedback signals* $\ell(1), \dots, \ell(M)$.

Accordingly, unlike conventional zero-th order learning methods, which require a separate set of forward passes for each parameter (see e.g., [36, 40]), our method performs just one set of M forward runs, after which local updates are computed independently for each neuron.




Fig. 5. Illustration of the proposed local zeroth-order learning rule. At each iteration, the proposed local zeroth-order learning rule implements: ① M global forward passes, where M is a constant M independent of the number of parameters, ② the evaluation of M global feedback signals at the output SQS neurons, and ③ local perturbation-based, zeroth-order updates within each neuron.

To derive the local gradient updates within each neuron, we consider separately hidden and output neurons.

1) *Output neurons*: Using (26), the gradient with respect to the parameters Θ_i^O of an output neuron $i \in \mathcal{O}$ is given by

$$\nabla_{\Theta_i^O} L_{o^T}(\Theta) = - \sum_{t=1}^T \mathbb{E}_{p_{\Theta^H}(h_{\mathcal{P}_i,t} | h_{\mathcal{P}_i}^{t-1})} \left[\nabla_{\Theta_i^O} \log (\text{Tr}(\rho_{i,t}^I O_{i,t})) \right]. \quad (29)$$

This can be estimated by using the local hidden variables $\{h_{\mathcal{P}_i}^T(m)\}_{m=1}^M$ extracted from the corresponding global realizations $\{h^T(m)\}_{m=1}^M$ of the hidden neurons as

$$\hat{\nabla}_{\Theta_i^O} L_{o^T}(\Theta) = - \frac{1}{M} \sum_{m=1}^M \sum_{t=1}^T \nabla_{\Theta_i^O} \log (\text{Tr}(\rho_{i,t}^I(m) O_{i,t}(m))), \quad (30)$$

where $\rho_{i,t}^I(m)$ is the density matrix of the input-output qubit of neuron i during the m -th global forward pass, and $O_{i,t}(m) = |o_{i,t}(m)\rangle\langle o_{i,t}(m)|$ is the corresponding m -th observable. In this next subsection, we will discuss how to compute the local gradient $\nabla_{\Theta_i^O} \log (\text{Tr}(\rho_{i,t}^I(m) O_{i,t}(m)))$.

2) *Hidden neurons*: For any hidden neuron $i \in \mathcal{H}$, using the REINFORCE gradient [19], the gradient with respect to the parameter Θ_i^H is obtained as

$$\nabla_{\Theta_i^H} L_{o^T}(\Theta) = - \sum_{t=1}^T \mathbb{E}_{p_{\Theta^H}(h_t | h^{t-1})} \left[\left(\sum_{i \in \mathcal{O}} \log \text{Tr}(\rho_{i,t}^I O_{i,t}) \right) \times \nabla_{\Theta_i^H} \log (\text{Tr}(\rho_{i,t}^I H_{i,t})) \right], \quad (31)$$

where $H_{i,t} = |h_{i,t}\rangle\langle h_{i,t}|$ is the observable corresponding to the realization $h_{i,t}$ of hidden neuron $i \in \mathcal{H}$. Define as

$$\ell_t(m) = - \sum_{i \in \mathcal{O}} \log \text{Tr}(\rho_{i,t}^I(m) O_{i,t}(m)) \quad (32)$$

the log-loss for the current data point o^T evaluated using the m -th realization $h^T(m)$ of the hidden neuron. As seen in Fig. 5, the quantities $\{\ell(1), \dots, \ell(M)\}$, with $\ell(m) = (\ell_1(m), \dots, \ell_T(m))$, are fed back to all hidden neurons at the end of time T , providing global feedback signals to all hidden neurons. Using this information, each neuron i estimates the gradient $\nabla_{\Theta_i^H} L_{o^T}(\Theta)$ as

$$\hat{\nabla}_{\Theta_i^H} L_{o^T}(\Theta) = \frac{1}{M} \sum_{m=1}^M \sum_{t=1}^T \ell_t(m) \times \nabla_{\Theta_i^H} \log (\text{Tr}(\rho_{i,t}^I(m) H_{i,t}(m))). \quad (33)$$

F. Local Gradient Estimation

In this subsection, we focus on estimating the gradient $\nabla_{\Theta_i^O} \log (\text{Tr}(\rho_{i,t}^I(m) O_{i,t}(m)))$ for any output neuron $i \in \mathcal{O}$ in (30) and the gradient $\nabla_{\Theta_i^H} \log (\text{Tr}(\rho_{i,t}^I(m) H_{i,t}(m)))$ for any hidden neuron $i \in \mathcal{H}$ in (33). We write generically such gradients as $\nabla_{\Theta_i} \log (\text{Tr}(\rho_{i,t}^I S_{i,t}))$ with $S_{i,t} = |s_{i,t}\rangle\langle s_{i,t}|$. Recall that for each neuron $i \in \mathcal{S}$, the local parameters $\Theta_i = \{w_i, \theta_i\}$ include the synaptic weights w_i and the PQC parameters θ_i . Therefore, estimating the gradient $\nabla_{\Theta_i} \log (\text{Tr}(\rho_{i,t}^I S_{i,t}))$ requires evaluating the gradients

$\nabla_{\mathbf{w}_i} \log(\text{Tr}(\rho_{i,t}^I S_{i,t}))$ and $\nabla_{\boldsymbol{\theta}_i} \log(\text{Tr}(\rho_{i,t}^I S_{i,t}))$. These gradients depend on the realization $\mathbf{h}_{\mathcal{P}_i}^T$ of the signals of the parents $\mathcal{P}_i$ of node i . As described next, we propose to estimate the former using simultaneous perturbation stochastic approximation (SPSA) [41], while the latter is estimated using the parameter shift rule (PSR). The use of the PSR for PQC parameters is standard, while SPSA can be more efficient than individual per-weight perturbations for synaptic weights.

1) *Gradient with Respect to the Synaptic Weights:* The synaptic weights $\mathbf{w}_i = \{\{w_{i,p}^n\}_{p=1}^{P_i}\}_{n=1}^N$ associated with each neuron i pertain to the NP_i links between the N input-output qubits of the P_i pre-synaptic neurons and the corresponding N qubits of neuron i . To estimate the gradient $\nabla_{\mathbf{w}_i} \log(\text{Tr}(\rho_{i,t}^I S_{i,t}))$ with respect to synaptic weight $\mathbf{w}_i$, we apply SPSA.

Accordingly, given the spiking signals of the parent neurons, neuron i runs a number of local forward passes to estimate the expected value $\text{Tr}(\rho_{i,t}^I S_{i,t})$ of the observable $S_{i,t}$ obtained with synaptic weights $\mathbf{w}_i + \epsilon \Delta \mathbf{w}_{i,m}$ and $\mathbf{w}_i - \epsilon \Delta \mathbf{w}_{i,m}$, where $\Delta \mathbf{w}_{i,m}$ is the m -th random perturbation vector with components drawn independently from the Rademacher distribution, and $\epsilon > 0$ represents the magnitude of the perturbation. We generate M_{syn} such independent perturbations $\{\Delta \mathbf{w}_{i,m}\}_{m=1}^{M_{\text{syn}}}$ for each neuron i . Denote as $\rho_{i,t}^I(\mathbf{w}_i + \epsilon \Delta \mathbf{w}_{i,m})$ and $\rho_{i,t}^I(\mathbf{w}_i - \epsilon \Delta \mathbf{w}_{i,m})$ the states of the input-output qubits with perturbed synaptic weights $\mathbf{w}_i + \epsilon \Delta \mathbf{w}_{i,m}$ and $\mathbf{w}_i - \epsilon \Delta \mathbf{w}_{i,m}$, respectively. To estimate the corresponding expectations $\text{Tr}(\rho_{i,t}^I(\mathbf{w}_i + \epsilon \Delta \mathbf{w}_{i,m}) S_{i,t})$ and $\text{Tr}(\rho_{i,t}^I(\mathbf{w}_i - \epsilon \Delta \mathbf{w}_{i,m}) S_{i,t})$, we apply M_p shots. Therefore, the number of local shots for each neuron i for the synaptic weights is $2M_p M_{\text{syn}}$. Writing as $\hat{O}_{i,m}^+$ and $\hat{O}_{i,m}^-$ the two estimates obtained with the m -th perturbation, the overall local gradient is estimated as

$$\hat{\nabla} \mathbf{w}_i \log(\text{Tr}(\rho_{i,t}^I S_{i,t})) = \frac{1}{M_{\text{syn}}} \times \sum_{m=1}^{M_{\text{syn}}} \frac{\log \hat{O}_{i,m}^+ - \log \hat{O}_{i,m}^-}{2\epsilon} \Delta \mathbf{w}_{i,m}. \quad (34)$$

2) *Gradient with Respect to the PQC parameter:* To estimate the gradient $\nabla_{\boldsymbol{\theta}_i} \log(\text{Tr}(\rho_{i,t}^I S_{i,t}))$ with respect to the PQC parameters $\boldsymbol{\theta}_i$, we apply the PSR to each individual PQC parameter $\theta_{i,j}$. Given the spiking signals from the parent neurons, neuron i performs a series of local forward passes to evaluate the expected value $\text{Tr}(\rho_{i,t}^I S_{i,t})$ using perturbed parameters $\theta_{i,j} + \pi/2$ and $\theta_{i,j} - \pi/2$. For each angle $\theta_{i,j}$, we generate M_{som} independent evaluations using these shifted values. Each expectation is estimated using M_p measurement shots per forward pass, leading to a total of $2M_p M_{\text{som}}$ local shots per angle $\theta_{i,j}$. Let $\hat{O}_{i,m}^+$ and $\hat{O}_{i,m}^-$ denote the measured estimates of the observable $S_{i,t}$ for the m -th perturbation with shifted parameters $\theta_{i,j} + \pi/2$ and $\theta_{i,j} - \pi/2$, respectively. Using the chain rule and averaging over the M_{som} repetitions, the gradient is estimated as

$$\nabla_{\theta_{i,j}} \log(\text{Tr}(\rho_{i,t}^I S_{i,t})) = \frac{1}{\text{Tr}(\rho_{i,t}^I S_{i,t})} \times \frac{1}{M_{\text{som}}} \sum_{m=1}^{M_{\text{som}}} \frac{\hat{O}_{i,m}^+ - \hat{O}_{i,m}^-}{2}. \quad (35)$$

Algorithm 1: Local learning rule for SQS-based Neural Networks.

Input: Training sequence $\mathbf{o}^T$ extracted from dataset $\mathcal{D}$
Output: Learned model parameters Θ

- 1 **Initialize** parameters Θ
- 2 *Feedforward sampling:*
- 3 Collect M samples of the hidden spiking signals $\mathbf{h}^T(1), \dots, \mathbf{h}^T(M)$ and the corresponding global feedback signals $\ell(1), \dots, \ell(M)$ in (32).
- 4 *Global feedback:*
- 5 Broadcast global learning signals $\ell(1), \dots, \ell(M)$ to all hidden neurons.
- 6 **for** each discrete time $t = 1, \dots, T$ **do**
- 7 *Local gradient estimation:*
- 8 **for** each neuron $i \in \mathcal{S}$ **do**
- 9 Compute the gradient $\nabla_{\boldsymbol{\theta}_i} \log(\text{Tr}(\rho_{i,t}^I S_{i,t}))$ using (34) or (35).
- 10 **end**
- 11 **end**
- 12 *Parameter update:*
- 13 Compute the gradient for each neuron based on (30) or (33) and perform gradient descent.

The proposed local training procedure is summarized in Algorithm 1. The method performs M global forward passes to obtain spike realizations of hidden neurons and scalar feedback signals, followed by $2TM_p|\mathcal{S}|(M_{\text{som}}G + M_{\text{syn}})$ local passes, with G denoting the number of PQC parameters per neuron and $|\mathcal{S}|$ being the number of neurons. This is contrasted with the $2TMM_p|\mathcal{S}|(M_{\text{som}}G + M_{\text{syn}})$ global forward passes required by conventional zero-th order training.

V. EXPERIMENTS

In this section, we evaluate the performance of the proposed SQSNN model on a variety of classification tasks.

A. Tasks

We compare the performance of classical and quantum models – conventional and spiking – on both standard datasets such as MNIST, FMNIST and KMNIST as in [11], and on neuromorphic datasets captured via event-driven cameras [42].

1) *USPS Binary Classification:* We first consider a simple binary classification task using the USPS handwritten digit dataset, focusing on distinguishing between the digits “1” and “7” as in [19]. Each input image is a 16×16 grayscale representation of a digit and is encoded into the spiking domain using rate encoding [43]. Accordingly, for each pixel, a spike train of $T = 10$ time steps is generated by sampling from a Bernoulli distribution, where the spiking probability is proportional to the pixel’s intensity and capped within the range $[0, 0.5]$. The class labels $\{1, 7\}$ are also represented using rate encoding during training. Specifically, there are two output neurons, and the target signal $\mathbf{o}^T$ for the two output neurons are such that the neuron associated with the correct class emits a spike at each time step [19].

2) MNIST, FMNIST, and KMNIST image classification:

Following [11], we also consider image classification tasks based on the standard MNIST, FMNIST, and KMNIST datasets. Each dataset consists of grayscale images of size 28×28 . Unlike [11], in order to highlight the capacity of different models to process event-driven data, we encode the inputs for the MNIST, FMNIST and KMNIST datasets using rate encoding, so that the firing probability for the signal encoding each pixel is proportional to its corresponding pixel value. Accordingly, the experiments on these datasets evaluate the models on spike-encoded temporal representations, rather than on the original dense-valued image inputs. The purpose of the experiments on MNIST, FMNIST, and KMNIST is primarily to validate that the proposed SQSNN remains competitive relative to existing spiking and quantum-spiking models under comparable parameter budgets when operating on spike-encoded temporal data. We set $T = 5$ discrete time steps.

3) *Neuromorphic Data*: Finally, we study classification tasks based on a neuromorphic dataset, namely MNIST-DVS [44], to better showcase the sequence modeling capabilities of SQSNN. The MNIST-DVS dataset consists of labeled spiking signals produced by an event-driven camera with a duration of $T = 80$ time steps [45]. Each data point contains $26 \times 26 = 676$ spike trains, recorded from a dynamic vision sensor (DVS) while it observes handwritten digits moving on a screen. The dataset provides 900 training examples and 100 testing examples per class. For both training and testing as in [46], only the first $T = 20$ time steps are used for classification.

B. Benchmarks

In the following experiments, we benchmark the proposed SQSNN model against the following alternative models:

- *Classical SNN* [43]: Classical feedforward SNN model based on LIF neurons, trained using surrogate gradient descent with backpropagation.
- *QLIF-SNN* [11]: Feedforward neural network implementing QLIF neurons.
- *QNN* [47]: A quantum variational classifier with classical input and read-out layers. For this model, the input is first encoded using a single classical neural layer, and then processed by a PQC using angle embedding and entangling gates. The outputs are measured and passed to a classical fully connected layer.

All models are configured to have approximately the same number of trainable parameters. This is a conventional benchmarking setting in quantum machine learning [31, 32] (see Section I).

C. Architecture

For all tasks and models, we adopt fully-connected feed-forward architectures. The output layers consists of as many neurons as there are classes. For spiking models, we use rate decoding, selecting the class whose corresponding neuron produces the largest number of spikes. In more details, for




Fig. 6. Illustration of the quantum circuit for a SQS neuron with $N = 1$ and $N_{\text{mem}} = 1$, where the first line corresponds to input-output qubit, while the second line is for the memory qubit.

the USPS dataset, we use fully connected models with a single, output, layer of two neurons. For the other datasets, the models also include a hidden layer. For the SQSNN, there are 1000 hidden neurons for the MNIST, FMNIST and KMNIST datasets, and 256 neurons for the MNIST-DVS dataset, while for the SNN and the QLIF-SNN baselines the number of hidden neurons is adjusted to match the total number of trainable parameters in the SQSNN. The QNN model is also set up to have the same number of parameters. The shallow ansatz in Fig. 6 follows the hardware-efficient design principle for near-term devices [48]. The learning rate is set to 5×10^{-4} for all models.

As illustrated in Fig. 6, for each SQS neuron, we configure the quantum circuit with $N = 1$ input-output qubit and $N_{\text{mem}} = 1$ memory qubit. The choice of N and N_{mem} introduces a trade-off between expressivity and complexity. Increasing these parameters enlarges the Hilbert space explored by the SQS neuron, potentially enabling richer temporal dynamics, but also increasing the likelihood of overfitting [49, 50]. Therefore, N and N_{mem} should be treated as architectural hyperparameters and selected through validation procedures that balance accuracy, resource requirements, and generalization performance. In practice, scaling to larger N or N_{mem} would increase the capacity of each SQS neuron, but is currently limited mainly by the exponential cost of classical simulation and by the limited availability of qubits on cloud quantum platforms.

The input currents z_t are encoded into rotation angles ϕ_t^n according to the mapping defined in (12), which determines the input-dependent $R_X(\phi_t^n)$ rotation at each time step. The ansatz $U(z_t, \theta)$ consists of a controlled-RX (CRX) gate followed by single-qubit RX rotations applied independently to each qubit, as in [47]. The CRX gate is parameterized by a rotation angle θ , and is given by the unitary matrix

$$\text{CRX}(\theta) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos(\theta/2) & -i \sin(\theta/2) \\ 0 & 0 & -i \sin(\theta/2) & \cos(\theta/2) \end{bmatrix}.$$

Both the CRX and RX gates contain trainable parameters that are optimized during learning. Therefore, in this setting, each SQS neuron i has three parameters in vector θ_i to be learned.

D. Evaluating the Local Learning Rule

We start by evaluating the properties of the proposed local learning rule, focusing on its per-iteration requirements in




Fig. 7. (a) Test accuracy versus training iterations on a NISQ device without QEM: Test accuracy as a function of training iterations for the QSNM model on the USPS dataset under the proposed local learning rule implemented on IBM’s `ibm_brisbane` without QEM. We vary the number of global passes M and measurement shots for synaptic weights M_{syn} , with $M_{\text{som}} = 1$. (b) Test accuracy versus training iterations on a NISQ device with QEM: Test accuracy as a function of training iterations for the QSNM model on the USPS dataset under the proposed local learning rule, implemented on IBM’s `ibm_brisbane` with zero noise extrapolation algorithm. The ideal local learning refers to the proposed local learning rule with a perfect estimation of spiking probability (15), while the surrogate gradient method is an efficient approximate training approach based on classical simulations that is described in the Appendix.

terms of the number of global forward passes, M , and the numbers of local shots at each neuron, namely $2M_{\text{syn}}$ for the synaptic weights and $2M_{\text{som}}$ for each PQC parameter.

To this end, we implement the local training rule on the IBM quantum computer `ibm_brisbane` using Qiskit [51]. In Fig. 7(a) and Fig. 7(b), we report the test accuracy of the QSNM on the USPS dataset as a function of the training iterations, while varying the parameters M and M_{syn} with $M_{\text{som}} = 1$ and $M_p = 5$. Fig. 7(a) does not assume quantum error mitigation (QEM), while Fig. 7(b) implements QEM using the zero noise extrapolation algorithm provided by the Mitiq library [51].

With $M = 1$ and $M_{\text{syn}} = 5$, the QSNM model is seen to achieve a peak test accuracy of nearly 95%, even in the absence of QEM. In contrast, configurations with lower shot counts for the synaptic weights, e.g., $M_{\text{syn}} = 2$, or a higher numbers of global forward passes, e.g., $M = 5$, exhibit reduced accuracy. Furthermore, increasing the number

of measurement shots from $M_{\text{syn}} = 2$ to $M_{\text{syn}} = 5$ at fixed $M = 5$ leads to a performance drop. We attribute this behavior to the non-stationarity of NISQ hardware. In fact, increasing M or M_{syn} enlarges the number of circuit executions that are run sequentially to form a single gradient estimate, thereby extending the wall-clock window over which the corresponding measurements are aggregated. Over such windows, the relaxation and dephasing times, as well as the gate and readout error rates of superconducting qubits are known to drift, even within a single calibration cycle [52, 53]. Furthermore, the implicit reset inserted between successive circuit executions is imperfect, leaving residual excited-state population that propagates into the following execution as a state-preparation error [54–56]. As a result of these effects, acquiring more samples need not reduce the overall estimation error when this drift is left uncompensated. Overall, the strong performance of the configuration with $(M = 1, M_{\text{syn}} = 5, M_{\text{som}} = 1)$ highlights the potential of local learning in hybrid quantum-classical systems, even with current NISQ devices without QEM.

Using the same hardware, Fig. 7(b) plots the test accuracy obtained by the QSNM model on the USPS dataset as a function of training iterations with QEM. We set $M_{\text{som}} = 1$, and report results for the configurations $(M = 1, M_{\text{syn}} = 2)$, $(M = 1, M_{\text{syn}} = 5)$ and $(M = 5, M_{\text{syn}} = 5)$. For reference, we include the performance of an ideal local training algorithm in which the local gradients (34)–(35) are computed exactly. In practice, this would require either a classical model of the neuron or the use of a large number of shots M_{syn} and M_{som} . We also evaluate a surrogate gradient-based method inspired by methods for classical SNNs [43], in which the spiking probability in (15) is replaced with a suitable sigmoid function, enabling the use of backpropagation. This scheme requires a classical simulator, providing a more efficient alternative for implementation on classical processor such as GPUs. We provide a brief description of this approach in Appendix.

We observe from Fig. 7(b) that increasing either the number of global forward passes M or the number of measurement shots M_{syn} leads to improved test accuracy. However, it is generally necessary to increase both numbers in order to obtain the best performance. In particular, with $M = 5$ and $M_{\text{syn}} = 5$, QSNM obtains a test accuracy around 98% at around 40 iterations, which matches that of the corresponding ideal local learning rule. Overall, the use of QEM is concluded to further improve the performance of QSNMs, while also stabilizing its performance as a function of the number of global forward passes and shots.

Finally, the surrogate gradient-based strategy—which relies on an efficient classical simulation of the system—provides a close approximation of the performance of the ideal local learning rule at convergence. While the proposed local learning rule is inherently hardware-compatible and intended for on-device adaptation, its direct evaluation requires a quantum processor of sufficient scale to handle benchmarks of practical interest. Therefore, consistent with prior works on quantum spiking neural networks [11], we henceforth employ classical simulations using a surrogate gradient-based approach to approximate the learning dynamics of the proposed rule (see

Appendix). This enables scalable evaluation on larger tasks, while preserving the relevance of the proposed local rule as an efficient learning mechanism for NISQ systems.




Fig. 8. Test accuracy and average number of spike per time of SNN, QNN, QLIF-SNN and SQSNN models on the MNIST, FMNIST, and KMNIST datasets: (a) Box plots of test accuracy achieved by each model on the three datasets. (b) Box plot for the average number of spikes per time step during inference.

E. Natural Image Tasks

We now present a comparative evaluation of the four models, SNN, QNN, QLIF-SNN and SQSNN, on the MNIST, FMNIST, and KMNIST image classification tasks. As shown in Fig. 8(a), the SQSNN model consistently achieves the highest test accuracy across all three datasets, demonstrating its superior capability in extracting relevant features from static input. The SNN also performs well, but lags behind the SQSNN, while the QNN and QLIF-SNN show comparatively lower accuracy than SQSNN, particularly on the more challenging KMNIST dataset.

Following common practice in neuromorphic engineering, we also report the average number of spikes per time step during inference. This quantity represents the aggregate spike count across the entire network, rather than per-neuron spiking activity. We adopt this measure to fairly compare models that differ in neuron count, but are constrained to a similar

total parameter budget. As discussed earlier, the number of spikes provides a useful measure of the classical and quantum processing complexity of the SQSNN. As seen in Fig. 8(b), the QLIF-SNN model exhibits the lowest number of spikes, but at the cost of a substantially lower accuracy. Notably, the SQSNN model achieves the most favorable performance trade-off, maintaining both high accuracy and moderate spiking activity. This balance highlights the potential of the SQSNN for energy-efficient neuromorphic computing, especially in hardware-constrained environment.

F. Neuromorphic Vision Task

We now evaluate all spiking models, SNN, QLIF-SNN, and SQSNN on a MNIST-DVS dataset. In order to explore the trade-off between accuracy and number of spikes, we add to the training objective (26) the regularization term

$$\frac{\lambda}{T} \sum_{t=1}^T \frac{(\sum_i u_{i,t})^2}{\sum_i |u_{i,t}|^2}, \quad (36)$$

where $u_{i,t}$ is a measure related to the spiking rate of neuron i at time t , $\lambda > 0$ is a hyperparameter, and the sum over index i is evaluated separately for each layer. Specifically, for the SQSNN, the quantity $u_{i,t}$ denotes the spiking probability $p(s_{i,t} = 1)$; for the QLIF-SNN, it represents the spiking probability given by multi-shot measurements [11]; and for the classical SNN, it corresponds to the membrane potential [43]. The regularizer (36) promotes the sparsity of the vector with entries $\{u_{i,t}\}_i$, encouraging lower spiking rates.




Fig. 9. Accuracy versus average number of spikes per time step for SNN, QLIF-SNN and SQSNN on the MNIST-DVS dataset under varying regularization strengths λ in (36).

In Fig. 9, we compare the test accuracy versus the average number of spikes per time step for the SNN, QLIF-SNN, and SQSNN models on the MNIST-DVS dataset under different regularization strengths λ . The results demonstrate that the SQSNN consistently outperforms both baseline models, achieving the highest accuracy across a broad range of spiking sparsity levels. Notably, SQSNN maintains high performance even with relatively few spikes, indicating its superior efficiency in processing spatio-temporal event-based data. In contrast, QLIF-SNN exhibits significantly lower accuracy at similar or higher spike counts. This highlights the effectiveness

of the quantum-enhanced architecture of SQSNN for extracting discriminative features from neuromorphic inputs.

G. Effect of Multiple Qubits in SQS neurons

To further examine the role of the multi-qubit SQS neuron design beyond the minimal configuration $N = N_{\text{mem}} = 1$, we now revisit the performance on the FMNIST dataset with enlarged quantum dimensions, namely $(N = 2, N_{\text{mem}} = 1)$ and $(N = 2, N_{\text{mem}} = 2)$.

Fig. 10 shows that enlarging the quantum neuron dimension consistently improves classification accuracy while maintaining sparse spike activity. The minimal SQS configuration $(N = 1, N_{\text{mem}} = 1)$ already outperforms both QLIF-SNN and classical SNN baselines with a lower spike rate. Increasing the number of input qubits to $N = 2$ further improves accuracy, and adding an additional memory qubit ($N_{\text{mem}} = 2$) achieves the best overall performance. Importantly, these gains are not accompanied by proportional increases in spike density, indicating improved temporal information integration rather than simple activity amplification.




Fig. 10. Test accuracy and average per-time per-neuron number of spikes for QLIF-SNN, SNN and SQSNN models with configurations $(N = 1, N_{\text{mem}} = 1)$, $(N = 2, N_{\text{mem}} = 1)$, $(N = 2, N_{\text{mem}} = 2)$ on the rate-encoded FMNIST dataset: (a) Box plots of test accuracy achieved by each model. (b) Box plot for the average per-time number of spikes during inference.

H. Neuromorphic Integrated Sensing and Communications

To further demonstrate the practical relevance of the proposed quantum spiking neuron beyond classification bench-

marks, we extend the evaluation to a neuromorphic integrated sensing and communications (N-ISAC) scenario [57]. In this setting, a shared waveform is used simultaneously for digital communication and radar sensing, and the neuromorphic receiver performs joint inference in an event-driven manner.

Specifically, an impulse-radio transmitter employs pulse-position modulation (PPM) to convey a binary sequence $\{x_l\}_{l=1}^L$. Bit 0 is represented by a pulse located at the first chip of a slot, whereas bit 1 is represented by a pulse located at the middle chip within a slot consisting of $2L_b$ chips, where L_b denotes the bandwidth expansion factor. The transmitted waveform propagates through a multipath fading channel that may include a radar target with presence indicator $v_l \in \{0, 1\}$ at a known delay cell, together with multiple clutter paths. The receiver samples the resulting complex baseband signal at the chip rate.

For each slot l , the receiver collects the $2L_b$ complex samples into a vector $\mathbf{y}_l$ and forms a real-valued representation $\bar{\mathbf{y}}_l \in \mathbb{R}^{4L_b}$ by stacking real and imaginary components. The temporal sequence $\{\bar{\mathbf{y}}_l\}_{l=1}^L$ is then processed by a neuromorphic receiver consisting of spiking neurons with two readout heads. These readouts generate a communication decoding decision $\hat{x}_l \in \{0, 1\}$ and a sensing decision $\hat{v}_l \in \{0, 1\}$ for each slot. Training is performed using a weighted sum of cross-entropy losses corresponding to communication and sensing objectives, where a weighting coefficient $\alpha \in [0, 1]$ controls the trade-off between the two objectives. In the experiments, we set $\alpha = 0.3$. Following the N-ISAC formulation, the decoding accuracy and sensing accuracy are defined as $1/L \sum_{l=1}^L \mathbb{1}(\hat{x}_l = x_l)$ and $1/L \sum_{l=1}^L \mathbb{1}(\hat{v}_l = v_l)$ respectively. We refer to [57] for further details.

Fig. 11 reports the performance obtained when adopting SQSNN and SNN at the receiver, respectively. QLIF-SNN is omitted for visual clarity since it consistently achieves the lowest performance among the compared schemes. Fig. 11(a) shows the decoding accuracy as a function of the number of training batches, where each training batch contains 128 samples. As the amount of training data increases, both receivers exhibit improved decoding performance. However, SQSNN consistently achieves higher decoding accuracy across all training regimes, with the advantage being more pronounced when the number of training batches is small. This suggests that SQSNN exhibits stronger representational efficiency in data-constrained settings.

Fig. 11(b) presents the sensing accuracy under the same training conditions. Sensing performance improves with increasing training batches for both models, while SQSNN maintains a performance advantage in the low training regimes. Fig. 11(c) reports the average number of spikes per time step. While SNN maintains a relatively high spike rate, SQSNN consistently operates with fewer spikes per time step. This demonstrates that the proposed quantum spiking architecture achieves superior sensing and decoding performance with reduced neural activity, reflecting improved event-driven efficiency.




Fig. 11. Test decoding accuracy, sensing accuracy, and average number of spikes per time step of the N-ISAC system using SNN and SQSNN, where each training batch contains 128 samples: (a) Decoding accuracy versus the number of training batches. (b) Sensing accuracy versus the number of training batches. (c) Average number of spikes per time step versus the number of training batches.

VI. CONCLUSIONS

In this work, we have introduced the stochastic quantum spiking (SQS) neuron, a novel quantum computational model that leverages the internal state space of multi-qubit quantum circuits to realize a spiking unit with inherent quantum memory. Unlike prior quantum spiking neuron models that require repeated measurements and classical memory mechanisms, the SQS neuron supports single-shot, event-driven spike generation, making it more aligned with the principles of neuromorphic computing.

We further studied SQS neural networks (SQSNNs) built from SQS neurons, and demonstrated that they can be trained using a local, hardware-efficient learning rule. This rule only requires the application of perturbation-based gradient estimates within each neuron, while steering learning via global feedback signals from the output to the hidden neurons.

By combining the sparse temporal processing efficiency of neuromorphic systems with the rich representational capacity of quantum computing, the proposed SQSNN architecture offers a promising direction for hybrid quantum neuromorphic intelligence. Our results open the door to the development of practical, trainable quantum spiking neural networks that can be implemented on future quantum hardware.

Future work may investigate the implementation on suitable quantum computers such as neutral-atom platforms (see Section I), the integration with quantum error mitigation, and the implementation on modular multi-core architectures with constraints on the bandwidth between cores as in classical neuromorphic chips [6].

APPENDIX A: SURROGATE GRADIENT LEARNING

This appendix introduces a surrogate gradient-based learning technique enabling the use of backpropagation for a classical simulator of an SQSNN. The approach is inspired by well-established methods for classical SNNs [43].

To elaborate, we consider a classification task with $|\mathcal{O}|$ classes, where $\mathcal{O}$ is the set of output neurons. Given an input at each time t , the output neurons produce a vector of spikes $\hat{o}_t = [\hat{o}_{1,t}, \dots, \hat{o}_{|\mathcal{O}|,t}]$, with $\hat{o}_{i,t} \in \{0, 1\}^N$. The spike rate for each output neuron $i \in \mathcal{O}$ over the full sequence of length T is evaluated as $r_i = \frac{1}{T} \sum_{t=1}^T \hat{o}_{i,t}$, providing an estimate of

the spiking probability (26). The loss in (26) is approximated as the cross-entropy between the spike rates r_i and the target outputs $\mathbf{o}_i^T$

$$\mathcal{L}(\Theta) = \frac{1}{|\mathcal{B}||\mathcal{O}|NT} \sum_{\mathbf{o}^T \in \mathcal{B}} \sum_{i=1}^{|\mathcal{O}|} \sum_{n=1}^N \sum_{t=1}^T [-o_{i,t}^n \log(r_i^n) - (1 - o_{i,t}^n) \log(1 - r_i^n)], \quad (37)$$

where $\mathcal{B} \subseteq \mathcal{D}$ is a mini-batch.

During the forward process, the output $\mathbf{s}_{i,t}$ of each neuron i at each time t is sampled from the Bernoulli distribution with probability $p_{\Theta_i}(\mathbf{s}_{i,t} | \mathbf{s}_{\mathcal{P}_i \cup \{i\}}^{t-1})$. Applying the chain rule, we compute the gradient of the loss (37) with respect to model parameters Θ_i of each neuron i as

$$\frac{\partial \mathcal{L}(\Theta)}{\partial \Theta_i} = \frac{\partial \mathcal{L}(\Theta)}{\partial \mathbf{s}_{i,t}} \cdot \frac{\partial \mathbf{s}_{i,t}}{\partial p_{\Theta_i}(\mathbf{s}_{i,t} | \mathbf{s}_{\mathcal{P}_i \cup \{i\}}^{t-1})} \cdot \frac{\partial p_{\Theta_i}(\mathbf{s}_{i,t} | \mathbf{s}_{\mathcal{P}_i \cup \{i\}}^{t-1})}{\partial \Theta_i}. \quad (38)$$

However, the spiking signals $\mathbf{s}_{i,t}$ are random and thus the chain rule (38) cannot be directly used. To address this problem, we replace the N -dimensional binary random signal $\mathbf{s}_{i,t}$ with a deterministic N -dimensional vector taking values in the interval $[0, 1]$. Specifically, denoting as $\sigma(\cdot)$ the element-wise sigmoid function $\sigma(x) = 1/(1 + \exp(-x))$, we replace vector $\mathbf{s}_{i,t}$ with the surrogate $\sigma(p_{\Theta_i}(\mathbf{s}_{i,t} | \mathbf{s}_{\mathcal{P}_i \cup \{i\}}^{t-1}))$, yielding the partial derivative approximation

$$\frac{\partial \mathcal{L}(\Theta)}{\partial \Theta_i} = \frac{\partial \mathcal{L}(\Theta)}{\partial \mathbf{s}_{i,t}} \cdot \frac{\partial \sigma(p_{\Theta_i}(\mathbf{s}_{i,t} | \mathbf{s}_{\mathcal{P}_i \cup \{i\}}^{t-1}))}{\partial p_{\Theta_i}(\mathbf{s}_{i,t} | \mathbf{s}_{\mathcal{P}_i \cup \{i\}}^{t-1})} \cdot \frac{\partial p_{\Theta_i}(\mathbf{s}_{i,t} | \mathbf{s}_{\mathcal{P}_i \cup \{i\}}^{t-1})}{\partial \Theta_i}. \quad (39)$$

This approximation is applied during the backward pass of backpropagation to estimate the gradient $\nabla_{\Theta} \mathcal{L}(\Theta)$.

REFERENCES

- [1] B. Rajendran, O. Simeone, and B. M. Al-Hashimi, "Towards efficient and reliable AI through neuromorphic principles," *Philosophical Transactions A of the Royal Society*, 2025.
- [2] D. Marković and J. Grollier, "Quantum neuromorphic computing," *Applied physics letters*, vol. 117, no. 15, 2020.

- [3] O. Simeone, *Classical and Quantum Uncertainty, Information, and Correlation*. Cambridge University Press, 2026.
- [4] S. Bravyi, D. Gosset, and R. König, "Quantum advantage with shallow circuits," *Science*, vol. 362, no. 6412, pp. 308–311, 2018.
- [5] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*. Cambridge university press, 2010.
- [6] M. Davies, A. Wild, G. Orchard, Y. Sandamirskaya, G. A. F. Guerra, P. Joshi, P. Plank, and S. R. Risbud, "Advancing neuromorphic computing with loihi: A survey of results and outlook," *Proceedings of the IEEE*, vol. 109, no. 5, pp. 911–934, 2021.
- [7] O. Simeone, "Modern neuromorphic AI: From intra-token to inter-token processing," *arXiv preprint arXiv:2601.00245*, 2026.
- [8] B. Rajendran, A. Sebastian, M. Schmuker, N. Srinivasa, and E. Eleftheriou, "Low-power neuromorphic hardware for signal processing applications: A review of architectural and system-level design approaches," *IEEE Signal Processing Magazine*, vol. 36, no. 6, pp. 97–110, 2019.
- [9] A. Mehonic, A. Sebastian, B. Rajendran, O. Simeone, E. Vasilaki, and A. J. Kenyon, "Memristors—from in-memory computing, deep learning acceleration, and spiking neural networks to the future of neuromorphic and bio-inspired computing," *Advanced Intelligent Systems*, vol. 2, no. 11, p. 2000085, 2020.
- [10] N. Skatchkovsky, H. Jang, and O. Simeone, "Bayesian continual learning via spiking neural networks," *Frontiers in Computational Neuroscience*, vol. 16, p. 1037976, 2022.
- [11] D. Brand and F. Petruccione, "A quantum leaky integrate-and-fire spiking neuron and network," *npj Quantum Information*, vol. 10, no. 1, pp. 1–8, 2024.
- [12] J. Dudas, B. Carles, E. Plouet, F. A. Mizrahi, J. Grollier, and D. Marković, "Quantum reservoir computing implementation on coherently coupled quantum oscillators," *npj Quantum Information*, vol. 9, no. 1, p. 64, 2023.
- [13] C. Gyurik, F. Wudarski, E. Philip, A. Sannia, H. Sadeghi, O. Kyriienko, D. Venturelli, and A. A. Gentile, "From quantum feature maps to quantum reservoir computing: perspectives and applications," *arXiv preprint arXiv:2510.01797*, 2025.
- [14] K. Fujii and K. Nakajima, "Harnessing disordered-ensemble quantum dynamics for machine learning," *Physical Review Applied*, vol. 8, no. 2, p. 024030, 2017.
- [15] H. W. Lau, A. Hayashi, A. Sakurai, W. J. Munro, and K. Nemoto, "Modular quantum extreme reservoir computing," *arXiv preprint arXiv:2412.19336*, 2024.
- [16] J. Murauer, R. Krishnakumar, S. Tornow, and M. Geierhos, "Feedback connections in quantum reservoir computing with mid-circuit measurements," *arXiv preprint arXiv:2503.22380*, 2025.
- [17] Y. Takaki, K. Mitarai, M. Negoro, K. Fujii, and M. Kitagawa, "Learning temporal data with a variational quantum recurrent neural network," *Physical Review A*, vol. 103, no. 5, p. 052414, 2021.
- [18] I. Nikoloska, O. Simeone, L. Banchi, and P. Veličković, "Time-warping invariant quantum recurrent neural networks via quantum-classical adaptive gating," *Machine Learning: Science and Technology*, vol. 4, no. 4, p. 045038, 2023.
- [19] H. Jang, O. Simeone, B. Gardner, and A. Gruning, "An introduction to probabilistic spiking neural networks: Probabilistic models, learning rules, and applications," *IEEE Signal Processing Magazine*, vol. 36, no. 6, pp. 64–77, 2019.
- [20] A. Deshpande, M. Hinsche, S. Najafi, K. Sharma, R. Sweke, and C. Zoufal, "Dynamic parameterized quantum circuits: expressive and barren-plateau free," *arXiv preprint arXiv:2411.05760*, 2024.
- [21] A. Gentile, S. Paesani, R. Santagati, J. Wang, N. Wiebe, D. Tew, J. O'Brien, and M. Thompson, "Noise resilience of bayesian quantum phase estimation tested on a si quantum photonic chip," in *European Quantum Electronics Conference*, p. EB_6_2, Optica Publishing Group, 2017.
- [22] K. Rudinger, G. J. Ribeill, L. C. Govia, M. Ware, E. Nielsen, K. Young, T. A. Ohki, R. Blume-Kohout, and T. Proctor, "Characterizing midcircuit measurements on a superconducting qubit using gate set tomography," *Physical Review Applied*, vol. 17, no. 1, p. 014014, 2022.
- [23] M. Cain, Q. Xu, R. King, L. R. Picard, H. Levine, M. Endres, J. Preskill, H.-Y. Huang, and D. Bluvstein, "Shor's algorithm is possible with as few as 10,000 reconfigurable atomic qubits," *arXiv preprint arXiv:2603.28627*, 2026.
- [24] J. Preskill, "Beyond nisq: The megaquop machine," 2025.
- [25] T. M. Graham, L. Phuttitarn, R. Chinnarasu, Y. Song, C. Poole, K. Jooya, J. Scott, A. Scott, P. Eichler, and M. Saffman, "Midcircuit measurements on a single-species neutral alkali atom quantum processor," *Physical Review X*, vol. 13, no. 4, p. 041051, 2023.
- [26] Y. Yu, K. Yan, D. Biswas, V. N. Zhang, B. Harraz, C. Noel, C. Monroe, and A. Kozhanov, "In situ midcircuit qubit measurement and reset in a single-species trapped-ion quantum computing system," *Physical Review Research*, vol. 7, no. 4, p. 043355, 2025.
- [27] A. Cowtan, S. Dilkes, R. Duncan, A. Krajenbrink, W. Simmons, and S. Sivarajah, "On the qubit routing problem," *arXiv preprint arXiv:1902.08091*, 2019.
- [28] D. B. Tan, D. Bluvstein, M. D. Lukin, and J. Cong, "Compiling quantum circuits for dynamically field-programmable neutral atoms array processors," *Quantum*, vol. 8, p. 1281, 2024.
- [29] M. Webber, S. Herbert, S. Weidt, and W. K. Hensinger, "Efficient qubit routing for a globally connected trapped ion quantum computer," *Advanced Quantum Technologies*, vol. 3, no. 8, p. 2000027, 2020.
- [30] T. Xiao, X. Zhai, X. Wu, J. Fan, and G. Zeng, "Practical advantage of quantum machine learning in ghost imaging," *Communications Physics*, vol. 6, no. 1, p. 171, 2023.
- [31] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, "The power of quantum neural networks," *Nature Computational Science*, vol. 1, no. 6, pp. 403–409, 2021.
- [32] T. Fellner, D. Kreplin, S. Tovey, and C. Holm, "Quantum vs. classical: A comprehensive benchmark study for predicting time series with variational quantum machine learning," *arXiv preprint arXiv:2504.12416*, 2025.
- [33] J. Chen, S. Park, and O. Simeone, "Knowing when to stop: Delay-adaptive spiking neural network classifiers with reliability guarantees," *IEEE Journal of Selected Topics in Signal Processing*, vol. 19, no. 1, pp. 88–102, 2025.
- [34] S. Karilanova, S. Dey, and A. Özçelikkale, "State-space model inspired multiple-input multiple-output spiking neurons," *arXiv preprint arXiv:2504.02591*, 2025.
- [35] C. Frenkel, D. Bol, and G. Indiveri, "Bottom-up and top-down approaches for the design of neuromorphic processing systems: Tradeoffs and synergies between natural and artificial intelligence," *Proceedings of the IEEE*, vol. 111, no. 6, pp. 623–652, 2023.
- [36] O. Simeone *et al.*, "An introduction to quantum machine learning for engineers," *Foundations and Trends® in Signal Processing*, vol. 16, no. 1-2, pp. 1–223, 2022.
- [37] H. Jang, N. Skatchkovsky, and O. Simeone, "VOWEL: A local online learning rule for recurrent networks of probabilistic spiking winner-take-all circuits," in *2020 25th International Conference on Pattern Recognition (ICPR)*, pp. 4597–4604, IEEE, 2021.
- [38] A. Sannia, G. L. Giorgi, and R. Zambrini, "Exponential concentration and symmetries in quantum reservoir computing," *arXiv preprint arXiv:2505.10062*, 2025.
- [39] J. C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE transactions on automatic control*, vol. 37, no. 3, pp. 332–341, 1992.
- [40] M. Schuld and F. Petruccione, *Machine learning with quantum computers*, vol. 676. Springer, 2021.
- [41] J. C. Spall, "An overview of the simultaneous perturbation

- method for efficient optimization,” *Johns Hopkins apl technical digest*, vol. 19, no. 4, pp. 482–492, 1998.
- [42] G. Gallego *et al.*, “Event-based vision: A survey,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 1, pp. 154–180, 2020.
 - [43] J. K. Eshraghian, M. Ward, E. O. Neftci, X. Wang, G. Lenz, G. Dwivedi, M. Bennamoun, D. S. Jeong, and W. D. Lu, “Training spiking neural networks using lessons from deep learning,” *Proceedings of the IEEE*, vol. 111, no. 9, pp. 1016–1054, 2023.
 - [44] T. Serrano-Gotarredona and B. Linares-Barranco, “Poker-dvs and mnist-dvs. their history, how they were made, and other details,” *Frontiers in neuroscience*, vol. 9, p. 481, 2015.
 - [45] J. Chen, N. Skatchkovsky, and O. Simeone, “Neuromorphic wireless cognition: Event-driven semantic communications for remote inference,” *IEEE Transactions on Cognitive Communications and Networking*, vol. 9, no. 2, pp. 252–265, 2023.
 - [46] Z. Wu, H. Zhang, Y. Lin, G. Li, M. Wang, and Y. Tang, “Liaf-net: Leaky integrate and analog fire network for lightweight and efficient spatiotemporal information processing,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 11, pp. 6249–6262, 2021.
 - [47] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, “Circuit-centric quantum classifiers,” *Physical Review A*, vol. 101, no. 3, p. 032308, 2020.
 - [48] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, *et al.*, “Variational quantum algorithms,” *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, 2021.
 - [49] L. Bianchi, J. Pereira, and S. Pirandola, “Generalization in quantum machine learning: A quantum information standpoint,” *PRX Quantum*, vol. 2, no. 4, p. 040321, 2021.
 - [50] M. C. Caro, H.-Y. Huang, N. Ezzell, J. Gibbs, A. T. Sornborger, L. Cincio, P. J. Coles, and Z. Holmes, “Out-of-distribution generalization for learning quantum dynamics,” *Nature Communications*, vol. 14, no. 1, p. 3751, 2023.
 - [51] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon, S. Martiel, P. D. Nation, L. S. Bishop, A. W. Cross, *et al.*, “Quantum computing with qiskit,” *arXiv preprint arXiv:2405.08810*, 2024.
 - [52] J. Etxezarreta Martinez, P. Fuentes, A. deMarti iOlius, J. Garcia-Frias, J. R. Fonollosa, and P. M. Crespo, “Multiqubit time-varying quantum channels for nisq-era superconducting quantum processors,” *Physical Review Research*, vol. 5, no. 3, p. 033055, 2023.
 - [53] H. Wang, P. Liu, Y. Liu, J. Gu, J. Baker, F. T. Chong, and S. Han, “Dgr: Tackling drifted and correlated noise in quantum error correction via decoding graph re-weighting,” *arXiv preprint arXiv:2311.16214*, 2023.
 - [54] I. Q. Documentation, “IBM quantum computer / Qubit initialization.” [Online]. <https://quantum.cloud.ibm.com/docs/en/guides/repetition-rate-execution>.
 - [55] B. Rost, L. Del Re, N. Earnest, A. Kemper, B. Jones, and J. Freericks, “Long-time error-mitigating simulation of open quantum systems on near term quantum computers, npj quantum inf,” 2025.
 - [56] J. Tan, C. Xu, T. Trochatos, and J. Szefer, “Extending and defending attacks on reset operations in quantum computers,” in *2024 International Symposium on Secure and Private Execution Environment Design (SEED)*, pp. 73–83, 2024.
 - [57] J. Chen, N. Skatchkovsky, and O. Simeone, “Neuromorphic integrated sensing and communications,” *IEEE Wireless Communications Letters*, vol. 12, no. 3, pp. 476–480, 2023.