

# Improving Quantum Machine Learning via Heat-Bath Algorithmic Cooling

Nayeli A. Rodríguez-Briones<sup>1,2,\*</sup> and Daniel K. Park<sup>3,4,5,†</sup>

<sup>1</sup>Atominstitut, *Technische Universität Wien*, Stadionallee 2, 1020 Vienna, Austria

<sup>2</sup>Miller Institute for Basic Research in Science, *University of California Berkeley*, Berkeley, California 94720, USA

<sup>3</sup>Department of Applied Statistics, *Yonsei University*, Seoul 03722, Republic of Korea

<sup>4</sup>Department of Statistics and Data Science, *Yonsei University*, Seoul 03722, Republic of Korea

<sup>5</sup>Department of Quantum Information, *Yonsei University*, Seoul 03722, Republic of Korea

 (Received 10 January 2025; revised 3 December 2025; accepted 4 February 2026; published 12 March 2026)

This work introduces an approach rooted in quantum thermodynamics to enhance sampling efficiency in quantum machine learning (QML). We propose conceptualizing quantum supervised learning as a thermodynamic cooling process. Building on this concept, we develop a quantum refrigerator protocol that enhances sample efficiency during training and prediction without the need for Grover iterations or quantum phase estimation. Inspired by heat-bath algorithmic cooling protocols, our method alternates entropy compression and thermalization steps to decrease the entropy of qubits, increasing polarization toward the dominant bias. This technique minimizes the computational overhead associated with estimating classification scores and gradients, presenting a practical and efficient solution for QML algorithms compatible with noisy intermediate-scale quantum devices.

DOI: [10.1103/24n3-tskj](https://doi.org/10.1103/24n3-tskj)

## I. INTRODUCTION

Quantum machine learning (QML) stands at the intersection of quantum information processing (QIP) and data science, exploring the fundamental limits of physical systems' ability to learn from data and generalize. By operating on radically different principles, QML holds the potential of surpassing its classical counterparts in analyzing complex data distributions and identifying intricate patterns [1–5]. However, quantum mechanics introduces unique challenges that are absent in classical ML approaches [6]. A critical issue arises from the probabilistic nature of quantum measurements. In QML, both training and inference rely on information extracted from probability distributions associated with the measurements used in the protocol, such as the expectation value of an observable [7–10]. This process inherently leads to finite sampling errors in QML algorithms.

In theory, adapting Quantum Amplitude Estimation (QAE) [11,12] to a QML algorithm can quadratically

reduce sampling errors. However, QAE requires multiple rounds of Grover-like operations [13,14], which significantly limits its feasibility for Noisy Intermediate-Scale Quantum (NISQ) computing [15,16]. Furthermore, what QAE achieves is often excessive for ML tasks. For instance, in classification problems, it is sufficient to determine the sign of a measured statistic, such as an expectation value, while its magnitude can remain undetermined. This scenario prompts a reevaluation of more practical sampling reduction techniques and motivates investigation into the possibilities of surpassing quadratic improvements in QML.

In response, this work introduces an approach rooted in quantum thermodynamics, reframing quantum supervised learning as a *thermodynamic cooling process*. In this framework, the quantum state plays the role of the data representation, and the direction of its population bias (i.e., the sign of its polarization) encodes the model's output. From this viewpoint, we show that reducing the entropy of the quantum states, while preserving their bias direction, mitigates finite sampling errors and improves the quality of statistical estimates. Building on this concept, we develop a quantum refrigerator protocol that enhances sample efficiency in both the training and prediction phases of QML, without relying on complex operations like Grover iterations and quantum phase estimation. Our technique draws inspiration from algorithmic cooling protocols [17–26], where entropy is systematically extracted through alternating steps of unitary operations and thermalization.

\*Contact author: [nayeli.briones@tuwien.ac.at](mailto:nayeli.briones@tuwien.ac.at)

†Contact author: [dkd.park@yonsei.ac.kr](mailto:dkd.park@yonsei.ac.kr)

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

A central innovation of this work is the introduction of the bidirectional cooling concept. This extends conventional cooling approaches by allowing quantum states with either positive or negative initial polarization to be cooled while preserving the sign of their bias. Remarkably, the same cooling operation applies to both positive and negative inputs and requires no prior knowledge of the initial bias direction (i.e., the correct model output). We introduce the notion of bidirectional cooling at two complementary levels. First, we present a unitary form of bidirectional cooling, in which global unitary operations increase the magnitude of the polarization while preserving the initially unknown sign of a target subsystem (a target qubit). Second, we extend this idea to a cyclic setting by constructing explicit multi-round protocols that alternate these bidirectional cooling unitary operations with thermalization, thereby establishing a novel Bidirectional Quantum Refrigerator (BQR) scheme.

We demonstrate through both theoretical analysis and numerical simulations that our proposed protocol significantly reduces the number of measurements required to estimate classification scores and gradients, leading to a substantial decrease in the overall computational cost of QML. Furthermore, we test the protocols under realistic NISQ noise conditions, showing that it establishes a physically grounded and noise-resilient pathway to enhanced learning performance.

In summary, this work introduces a thermodynamic framework for QML, where entropy-reduction processes act as operational primitives to improve learning efficiency. Specifically, we (i) develop Bidirectional Quantum Protocols that can be embedded within QML circuits to reduce the entropy of data-encoding qubits while preserving the sign of their polarization; (ii) show through theoretical analysis and numerical simulations that the protocols markedly decrease the number of measurements required to estimate classification scores and gradients, thereby reducing the overall computational cost of QML across a broad range of initial polarizations; (iii) evaluate the classification performance of BQR-enhanced quantum classifiers using different datasets and simulation setups; and (iv) test the protocol under noise models, demonstrating its robustness and suitability for current NISQ devices.

The paper is organized as follows. Section II introduces the necessary background and formalizes the classification problem. In Sec. II A, we present the framework for variational quantum classifiers, formulating the hypothesis function and classification rule using single-qubit measurements. In Sec. III, we discuss the challenge of finite sampling errors, presenting bounds on error probabilities for prediction and training. This section also highlights the need for a bidirectional cooling method to enhance the magnitude of polarization of the classification score. In Sec. IV, we provide a brief review of Heat-Bath

Algorithmic Cooling. In Secs. V and VI, we present the core contributions of this work: the Bidirectional Cooling Protocols. Section V A begins with a comprehensive analysis of the unitary case of bidirectional cooling. Next, in Sec. VI A, we introduce the Bidirectional Quantum Refrigerator (BQR) scheme, demonstrating its ability to significantly reduce the number of measurements required to estimate classification scores and gradients. In Sec. VI B, we present a practical adaptation that employs  $k$ -local unitaries to enhance its implementability. In Sec. VI C, we provide an overview of the protocols introduced in this work. In Sec. VII, we show the classification performance evaluation. Finally, Sec. VIII concludes the paper with a discussion of our findings and outlines promising directions for future research.

## II. BACKGROUND ON THE CLASSIFICATION FRAMEWORK

Classification is a cornerstone of data analysis, underpinning a wide range of applications [27]. Approaches to this task span from traditional statistical methods and signal processing to advanced techniques in both classical and QML. Supervised learning for classification begins with a sample dataset  $\mathcal{D} = \{(x_1, y_1), \dots, (x_s, y_s)\} \subseteq \mathbb{R}^d \times \mathbb{Z}_2$ , where  $x_i$  denotes the feature vector and  $y_i$  represents its binary label (or class). Typically, each sample  $(x_i, y_i)$  is assumed to be drawn from an unknown joint probability distribution  $P_{X,Y}$ , i.e.,  $(x_i, y_i) \sim P_{X,Y}$ . Then, the objective is to construct a hypothesis  $f(x)$  such that, for a given loss function  $L$ , which quantifies the difference between  $f(x)$  and the true label  $y$ , the expected loss  $\mathbb{E}_{(x,y) \sim P_{X,Y}}(L(f(x), y))$  is minimized. In practice, the search is often confined to a family of functions parameterized by  $\theta$ , expressed as  $f(x, \theta)$ . Moreover, since the classifier can only be constructed based on the  $s$  available samples, the task is typically addressed by minimizing the empirical risk,  $\sum_{i=1}^s L(f(x_i, \theta), y_i)/s$  with respect to  $\theta$ . Note that multiclass classifiers with  $l > 2$  labels can be constructed from binary classifiers using strategies such as One-vs-One or One-vs-Rest comparisons [27–29].

### A. Quantum classifier with single-qubit measurement

The classification problem can be solved on a quantum computer by training a variational quantum circuit (VQC). In this framework, the canonical form of the hypothesis function is expressed as  $f(x, \theta) = \text{sign}(\langle x | U^\dagger(\theta) M U(\theta) | x \rangle)$ , where  $|x\rangle$  is an  $n$ -qubit representation of the data,  $M$  is a Hermitian matrix and  $\text{sign}(\cdot) \in \{-1, +1\}$  denotes the sign function.

Without loss of generality, we can assume that  $M$  is an  $n$ -qubit Pauli operator, i.e.,  $M \in \{I, X, Y, Z\}^{\otimes n}$ . This assumption is well-suited for binary classification, as all Pauli operators have eigenvalues of  $+1$  and  $-1$ . Consequently, the outcome of any variational quantum

binary classifier (VQBC) can be written as  $q(x, \theta) = \langle x | U^\dagger(\theta) U_c^\dagger Z_1 U_c U(\theta) | x \rangle$ , where  $Z_1$  is the single-qubit Pauli-Z operator acting nontrivially on a single qubit and  $U_c$  is an  $n$ -qubit Clifford operator. The term  $q(x, \theta)$  is referred to as the classification score. The VQBC then assigns a label to  $x$  based on the sign of this score, i.e.,  $\text{sign}(q(x, \theta))$ .

An important observation is that any quantum binary classifier can be constructed by training a VQC with a single-qubit Pauli measurement. Moreover, single-qubit measurements are particularly advantageous for mitigating the barren plateaus phenomenon [8,30], which is one of the pressing challenges in variational quantum algorithms. This effectiveness is demonstrated through various techniques, including the use of hierarchical circuit architecture [31–35], or local observables [36]. In these approaches,  $M$  can be selected as  $Z_1$ , eliminating the need for Clifford transformations.

Now, we rewrite the classification score as

$$q(x, \theta) = \text{Tr}(Z \rho_1(x, \theta)), \quad (1)$$

where  $\rho_1(x, \theta) = \text{Tr}_{n-1}(U_c U(\theta) | x \rangle \langle x | U^\dagger(\theta) U_c^\dagger)$  is a single-qubit density operator and  $\text{Tr}_{n-1}(\cdot)$  denotes the partial trace over the  $n - 1$  qubits that are not being measured. Without loss of generality, the Clifford operator  $U_c$  can be absorbed into the parameterized unitary. By expressing the single-qubit density matrix as

$$\rho_1(x, \theta) = \frac{I + \alpha(x, \theta)Z + \beta(x, \theta)X + \gamma(x, \theta)Y}{2}, \quad (2)$$

we see that the classification score is given by  $q(x, \theta) = \alpha(x, \theta)$ . The classification rule for an unseen data point  $\tilde{x}$  can then be written as

$$\tilde{y} = \text{sign}(\alpha(\tilde{x}, \theta^*)), \quad (3)$$

where  $\theta^*$  is the optimal set of parameters obtained by training the VQC. When  $\alpha(\tilde{x}, \theta^*) = 0$ , the decision can be made at uniformly random.

Since only the  $Z$  component of the density matrix affects classification, we can assume that the final state of the VQC is

$$\rho_z(x, \theta) = \frac{I + \alpha(x, \theta)Z}{2}. \quad (4)$$

### III. REDUCTION OF FINITE SAMPLING ERROR

We aim to minimize the sampling error in VQBC, which arises from the finite number of repetitions (or shots) used to estimate the classification score in Eq. (1). The relationship between the number of repetitions and the estimation

error is described by the Chebyshev inequality

$$\Pr[|\mu - \langle M \rangle| \geq \epsilon] \leq \sigma^2 / (k\epsilon^2) \quad (5)$$

where  $\mu$  is an average value obtained from  $k$  repetitions,  $\langle M \rangle$  is the expectation value to be estimated, and  $\sigma^2 = \langle M^2 \rangle - \langle M \rangle^2$  is the variance.

When predicting the label of unseen data  $\tilde{x}$ , the classifier assigns the label based on the sign of the estimated expectation value  $\mu$ , which is obtained from a finite number of measurements. To ensure that the estimated classification outcome matches the model's ideal prediction in the infinite-shot limit, the estimation error must satisfy  $|\mu - \langle M \rangle| < |\langle M \rangle|$ . Otherwise, the estimate  $\mu = \langle M \rangle \pm |\mu - \langle M \rangle|$  may have a different sign than  $\langle M \rangle$ . This condition defines the minimum requirement for the prediction to remain robust against sampling noise.

In our case,  $\langle M \rangle = \alpha(\tilde{x}, \theta^*)$  and  $\sigma^2 = \langle M^2 \rangle - \langle M \rangle^2 = 1 - \alpha^2(\tilde{x}, \theta^*)$ . Therefore, the error probability is bounded from above as

$$\Pr[\text{error}] = \Pr[|\mu - \langle M \rangle| \geq |\langle M \rangle|] \leq \frac{1 - \alpha^2(\tilde{x}, \theta^*)}{k\alpha^2(\tilde{x}, \theta^*)}. \quad (6)$$

In contrast, the training process aims to find the set of parameters  $\theta$  such that, for all samples  $(x_j, y_j) \in \mathcal{D}$ , the following conditions hold:

$$\begin{aligned} q(x_j, \theta) &> b \text{ if } y_j = +1, \\ q(x_j, \theta) &< -b \text{ if } y_j = -1, \end{aligned} \quad (7)$$

where  $0 \leq b < 1$  defines the margin. These conditions can be expressed as a single inequality that the training samples should satisfy:  $y_j q(x_j, \theta) - b > 0 \forall j = 1, \dots, s$ . Thus, the training process seeks to minimize the hinge loss, defined as  $l(\theta | x_j, y_j) = \max(0, b - y_j q(x_j, \theta))$ . The gradient of the hinge loss is given by

$$\frac{dl}{d\theta} = \begin{cases} -y_j \frac{\partial q(x_j, \theta)}{\partial \theta} & \text{if } y_j q(x_j, \theta) < b \\ 0 & \text{otherwise} \end{cases}. \quad (8)$$

Equation (8) highlights that implementing gradient-based optimization requires estimating  $q(x_j, \theta)$  with an error smaller than  $|q(x_j, \theta) - b|$  to accurately assess whether  $y_j q(x_j, \theta) < b$ . Consequently, the probability of making an incorrect decision at this step is bounded by

$$\Pr[\text{error}] \leq \frac{1 - \alpha^2(x_j, \theta)}{k(\alpha(x_j, \theta) - b)^2}. \quad (9)$$

Additionally, the estimation error must be sufficiently small to ensure effective parameter updates. The gradient's sign indicates the direction of the parameter update, while

its magnitude determines the step size in that direction. Notably, correctly estimating the sign of the gradient is more critical than its magnitude, as the latter can be scaled by an adaptively adjusted learning rate. This implies that  $\partial q(x_j, \theta)/\partial \theta$  must be estimated within the error margin smaller than  $|\partial q(x_j, \theta)/\partial \theta|$ . This requirement results in an upper bound on the error probability similar to Eq. (6), as shown in Appendix C.

The error probabilities in Eqs. (6) and (9) highlight the importance of having a large magnitude for  $\alpha(x, \theta)$  to minimize the number of repetitions required. This observation motivates the development of a protocol that increases  $|\alpha(\tilde{x}, \theta)|$ . Such a protocol cannot be a locally unitary process on the single qubit, as it must alter the purity of its state. Various algorithmic cooling techniques, based on engineering the system-bath interaction, exist that can increase the population of a chosen basis state [17–21, 23–25, 37–40]. However, it is crucial to note that the sign of  $\alpha(x, \theta)$  is not known a priori in the QML context. Thus, conventional algorithmic cooling techniques, which solely increase the population of a predetermined basis state, cannot achieve our goal. Instead, we require a bidirectional protocol capable of dynamically transforming the single-qubit density matrix as follows:

$$\frac{I + \alpha Z}{2} \rightarrow \frac{I + \alpha' Z}{2} \text{ with } \begin{cases} \alpha' > \alpha & \text{if } \alpha > 0 \\ \alpha' < \alpha & \text{if } \alpha < 0 \end{cases} \quad (10)$$

Bidirectional cooling can also be motivated from the perspective of quantum state discrimination. To illustrate this, consider splitting the sample dataset as  $\mathcal{D} = \mathcal{D}_+ \cup \mathcal{D}_-$  where  $\mathcal{D}_\pm = \{(x_i^\pm, \pm 1)\}_{i=1}^{s_\pm}$  denotes the subset containing only data labeled as  $\pm 1$ , and  $s_\pm$  denotes the size of the  $\pm 1$ -class dataset. As noted in Ref. [41], the optimization process can be interpreted as finding  $\theta$  that maximizes the discrimination between the density matrices  $\rho_+ = \sum_{i=1}^{s_+} \rho_z(x_i^+, \theta)/s_+$  and  $\rho_- = \sum_{i=1}^{s_-} \rho_z(x_i^-, \theta)/s_-$ , which represent the ensemble of data points in each class, with a high success probability. In this picture, the empirical risk (i.e., training loss) is bounded from below by the error probability in distinguishing the two ensembles. These ensembles can also be expressed as  $\rho_\pm = (I + \bar{\alpha}_\pm Z)/2$ , where  $\bar{\alpha}_\pm$  is the average polarization of  $\rho_z(x_i^\pm, \theta)$ . Assuming  $s_+ = s_-$ , the error probability of discriminating  $\rho_+$  and  $\rho_-$  is

$$\begin{aligned} \text{Pr}[\text{error}] &= \frac{1}{2} \left( 1 - \left\| \frac{s_+}{s} \rho_+ - \frac{s_-}{s} \rho_- \right\|_1 \right) \\ &= \frac{1}{2} - \frac{1}{4} |\bar{\alpha}_+ - \bar{\alpha}_-|. \end{aligned} \quad (11)$$

Therefore, increasing the magnitude of the polarizations via bidirectional cooling reduces the lower bound of the empirical risk.

While the main results of this work are demonstrated through VQBC, the transformation outlined in Eq. (10) is applicable to any binary classifier established based on the classification score given in Eq. (1). An example of this broader applicability is provided in Appendix D.

The observation that reducing the entropy of the measured qubit is beneficial is not only central to our framework, but also consistent with previous studies on entanglement-induced barren plateaus. These works have shown that strong entanglement between the visible (measured) qubit and hidden qubits hinders efficient training of the model [42, 43]. To address this, several strategies have been proposed [43], all of which involve either modifying the machine learning model, such as altering the structure or parameters of the variational circuit, or changing the training strategy.

In contrast, we introduce an approach rooted in quantum thermodynamics as an independent mechanism for reducing finite sampling error, regardless of whether the system exhibits a barren plateau, without altering either the classifier structure or training procedure. Specifically, this approach achieves the transformation in Eq. (10) without requiring any prior information about  $\alpha$ .

Before introducing our main contribution—the Bidirectional Quantum Refrigerator protocols—we begin with a brief review of Heat-Bath Algorithmic Cooling and its essential components, which serve as the conceptual foundation for our approach.

## IV. HEAT-BATH ALGORITHMIC COOLING

### A. A brief review

Advances in quantum information science have enabled the development of cooling techniques for quantum systems based on controlled entropy manipulation through sequences of quantum operations. In particular, Heat-Bath Algorithmic Cooling (HBAC) protocols cool a quantum system by repeatedly applying a two-step process: (i) a unitary operation that redistributes entropy within the system, removing entropy from a target subsystem while concentrating it in the complementary subsystem; and (ii) a dissipative process, such as thermalization or a reset operation, that expels the accumulated entropy into a heat bath. [17–23, 25, 37, 39].

The standard HBAC setting (though not restricted to this specific configuration) consists of a register of  $n$  qubits: one target qubit to be cooled;  $m$  reset qubits that can be refreshed, e.g., by being brought into thermal contact with a heat bath; and  $n_A = n - m - 1$  auxiliary qubits that assist in redistributing entropy during the unitary operation.

The standard HBAC protocol operates in repeated rounds, each round consisting of two steps:

**Step 1. Unitary Step: Entropy Compression.** A global unitary operation  $U(\rho)$  is applied to the target, auxiliary,

and reset qubits. This operation is designed to redistribute entropy across the entire register, extracting entropy from the target + auxiliary subsystems and concentrating it in the reset qubits. When the entropy compression is designed to maximize the ground-state population of the target qubit, the target becomes colder while the reset qubits become correspondingly warmer.

A single-shot unitary step is often referred to simply as Algorithmic Cooling (AC), corresponding to HBAC without the heat-bath reset. In this fully unitary setting, the optimal entropy-compression operation is a global unitary that diagonalizes the state in the energy eigenbasis and reorders the populations so that the largest diagonal entries are placed in the first half of the spectrum (with the target qubit corresponding to the first subsystem of the  $n$ -qubit register). This population permutation maximizes the probability that the target qubit occupies its ground state through unitary operations [22,26].

**Step 2. Dissipative Step: Thermal or Reset Operations.** The reset qubits are refreshed by coupling them to a heat bath (or by applying an appropriate dephasing thermal map), thereby expelling the accumulated entropy from the reset subsystem without external energy input. When full thermalization is implemented, this step is effectively equivalent to replacing the reset qubits with fresh qubits drawn from a thermal reservoir.

Iterating these two steps gradually increases the polarization of the target subsystem, driving it toward an asymptotic cooling limit determined by the bath temperature, the system size ( $n, m$ ), and the specific thermal operation used [19,25]. These steps are illustrated in Fig. 1.

Various HBAC protocols have been proposed and experimentally realized on diverse platforms, underscoring their versatility and practical relevance [17,19–21,23,37–39,44–50]. In particular, this work is inspired by the Partner-Pairing Algorithm (PPA-HBAC) protocol [18,19,25]. In PPA-HBAC, the unitary in Step 1 performs a descending sort of the diagonal elements of the system’s density matrix. In the subsequent thermal reset step, the  $m$  reset qubits are fully thermalized via contact with the bath. Note that although other engineered thermalization processes can yield lower cooling limits, full thermalization is convenient for the purposes in this work, as it is operationally equivalent to replacing the reset qubits with fresh qubits drawn from the same sample. The maximum polarization achievable on a target qubit by the PPA-HBAC [19,25] is given by

$$\alpha_{\text{HBAC}}^{\infty} = \frac{(1 + \alpha_0)^{md} - (1 - \alpha_0)^{md}}{(1 + \alpha_0)^{md} + (1 - \alpha_0)^{md}} \quad (12)$$

$$= \tanh [md * \operatorname{arctanh}(\alpha_0)], \quad (13)$$

where  $\alpha_0$  is the polarization of the thermal bath (or ensemble) qubits,  $m$  is the number of reset qubits used per round,

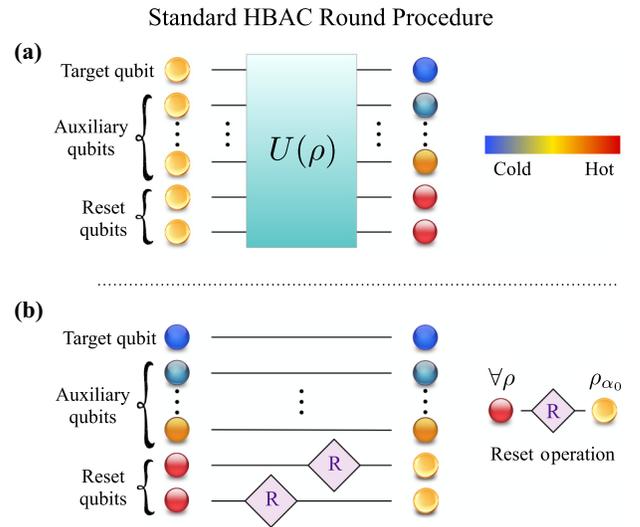


FIG. 1. A round of the standard heat-bath algorithmic cooling protocol. (a) Step 1. Unitary Step: Entropy compression. A global unitary operation  $U(\rho)$  acting on the target, auxiliary, and reset qubits coherently redistributes entropy across the entire register. This operation extracts entropy from the target and auxiliary qubits and concentrates it in the reset subsystem, thereby increasing the population of the target qubit toward its ground state (colder, blue states). (b) Step 2. Dissipative Step: Thermalization or Reset Operations. The reset qubits are refreshed through interaction with a heat bath, removing the accumulated entropy. The illustration corresponds to the case of full thermalization, which is operationally equivalent to replacing the reset qubits with fresh qubits prepared in the bath state  $\rho_{\alpha_0}$ , illustrated here for the case of  $m = 2$  reset qubits. Repeated iteration of these two steps constitutes the standard heat-bath algorithmic cooling protocol. The color gradient from blue to red indicates the polarization of the qubits, from colder (blue) to hotter (red) states.

and  $d$  is the dimension of the auxiliary system ( $d = 2^{n_A} = 2^{n-m-1}$ ).

## B. HBAC for quantum classification: A new perspective

While HBAC has traditionally been developed to enhance ground-state populations in quantum systems, its applicability to classification tasks—where the sign of an observable encodes the output label—has not been previously explored. HBAC methods have demonstrated the ability to increase the ground-state population  $p_0$ , through alternating rounds of entropy compression unitaries and thermalization steps. However, these protocols are inherently unidirectional, as they are specifically designed to increase  $p_0$ , resulting in a monotonic change in  $\alpha(x, \theta)$ . Moreover, constructing optimal unitaries in these protocols requires prior knowledge of the sign of  $\alpha(x, \theta)$ . This dependence makes previous HBAC protocols unsuitable for classification problems, where the sign

of  $\alpha(x, \theta)$ —which determines the gradient direction or predicts the label—must remain unknown a priori.

In this study, we extend the foundational elements of HBAC to classification problems for the first time. We introduce a quantum refrigerator designed to reliably increase polarization in the bias direction, regardless of the unknown parameter  $\alpha(x, \theta)$ . Specifically, we design and utilize entropy compression gates that obey the transformation given by Eq. (10) for unknown polarization, functioning effectively in both bias directions. Based on these findings, we introduce a family of cooling protocols called Bidirectional Quantum Refrigerators (BQRs).

Moreover, unlike conventional HBAC methods, our protocols not only perform bidirectional cooling but also operate cyclically: after extracting the enhanced qubits, the remaining subsystem—in its updated state—is recycled to generate multiple subsequent enhanced targets. This cyclic operation substantially reduces the required qubit resources while maintaining effective cooling performance.

## V. BIDIRECTIONAL COOLING

### A. Unitary Bidirectional Cooling (UBC)

A key step of heat-bath algorithmic cooling is entropy compression via unitary operations—a process that reduces the entropy of a target subsystem by coherently redistributing it across the entire system. Here, we extend this scheme by introducing a concept that we call **Unitary Bidirectional Cooling (UBC)**: a unitary compression method that amplifies the polarization magnitude of a target qubit irrespective of the sign of its initial bias, as described in Eq. (10). The same unitary operation simultaneously enhances both positive and negative polarizations, preserving their sign while increasing their magnitude. This single-shot bidirectional compression forms the core of the Bidirectional Quantum Refrigerator (BQR) developed in the following sections, where successive bidirectional compression steps are alternated with reset operations. Unlike conventional HBAC approaches, our BQR method recycles the working body of the refrigerator, generating multiple purified samples with reduced entropy and thereby lowering the overall resource overhead.

**Setup and UBC Protocol**—The system consists of  $n$  qubits: one target qubit, whose polarization magnitude is to be enhanced, and  $n - 1$  auxiliary qubits that facilitate entropy compression. In unitary bidirectional cooling, the goal is twofold: to increase the population of the target qubit’s  $|0\rangle$  state when the initial polarization is positive and to increase the population of its  $|1\rangle$  state when the polarization is negative. Crucially, the same unitary operation must accomplish both behaviors without prior knowledge of the polarization sign; that is, it must operate correctly regardless of whether the initial polarization is positive or negative. The following theorem characterizes

the optimal unitary bidirectional cooling protocol on  $n$  identical qubits, followed by a corollary that provides an explicit arrangement of the populations.

*Theorem 1 (Optimal Unitary Bidirectional Cooling on  $n$  Identical Qubits).* Consider a system of  $n$  identical qubits, each prepared in the state

$$\rho = \begin{pmatrix} p & 0 \\ 0 & 1-p \end{pmatrix},$$

with  $p = (1 + \alpha)/2$ , where  $\alpha \in [-1, 1]$  denotes the unknown polarization parameter. The joint state  $\rho^{\otimes n}$  is diagonal in the computational basis, with entries of the form  $p^k(1-p)^{n-k}$  corresponding to basis states  $|x \in \{0, 1\}^n\rangle$ , where  $k$  is the number of zeros in the bit string  $x$ , and  $n - k$  is the number of ones. Optimal bidirectional cooling of the target qubit (the first qubit) through unitary operations is achieved by a permutation that assigns the  $2^{n-1}$  populations with the largest exponents  $k$  to the first half of the diagonal; the order within this block is irrelevant. This permutation simultaneously increases the polarization  $\alpha'$  of the target qubit when  $\alpha > 0$  and decreases it when  $\alpha < 0$ , thereby increasing  $|\alpha'|$  in both cases while preserving its sign.

*Proof.* We begin by noting the *monotone ordering of the diagonal weights*. The spectrum of  $\rho^{\otimes n}$  consists of the entries  $\lambda_k = p^k(1-p)^{n-k}$  with multiplicity  $\binom{n}{k}$  for  $k = 0, \dots, n$ .

When  $\alpha > 0$  (i.e.,  $p > 1 - p$ ), the sequence  $\lambda_k$  is strictly *decreasing* as  $k$  decreases,

$$p^n > p^{n-1}(1-p) > p^{n-2}(1-p)^2 > \dots > (1-p)^n, \\ \text{if } \alpha > 0.$$

Conversely, when  $\alpha < 0$  (i.e.,  $p < 1 - p$ ), the same sequence in the same order is strictly *increasing*

$$p^n < p^{n-1}(1-p) < p^{n-2}(1-p)^2 < \dots < (1-p)^n, \\ \text{if } \alpha < 0.$$

After applying a global unitary  $U$ , the state of the  $n$  qubits becomes  $\rho'_T := U\rho^{\otimes n}U^\dagger$ . The updated population of the target qubit’s  $|0\rangle$  level, denoted  $p'$  ( $\rho'_T$ ), is obtained by summing all diagonal entries of  $\rho'_T$  in which the target qubit is in state  $|0\rangle$ ,

$$p'(\rho'_T) = \sum_{y \in \{0,1\}^{n-1}} \langle 0y | \rho'_T | 0y \rangle.$$

The corresponding updated polarization of the target qubit is then given by  $\alpha' = 2p'(\rho'_T) - 1$ .

By the Schur-Horn theorem, to maximize or minimize  $p'(\rho'_T)$  over all global unitaries  $U$  acting on the diagonal

input state  $\rho^{\otimes n}$ , it suffices to consider permutation unitaries that reorder its diagonal populations. Namely, unitaries that generate coherences by rotating the input diagonal state cannot improve either extremum of the sum defining  $p'(\rho'_T)$  that is achievable from its eigenvalues (see Appendix A 1 for details).

Therefore, to *maximize*  $\alpha'$  when  $\alpha > 0$ , we must maximize the sum  $p'(\rho'_T)$ . By the above argument, this is achieved by assigning the  $2^{n-1}$  largest eigenvalues  $\lambda_k$  to the  $|0\rangle$  block, i.e., to the first half of the diagonal of the joint state (the ordering within that block is irrelevant). Conversely, when  $\alpha < 0$ , we must *minimize* the same sum, which is attained by assigning the  $2^{n-1}$  smallest eigenvalues  $\lambda_k$  to the same  $|0\rangle$  block.

Finally, as noted at the beginning, the eigenvalues vary monotonically with  $k$  in opposite directions for the two sign regimes: when  $\alpha > 0$  they increase with  $k$ , whereas when  $\alpha < 0$  they decrease with  $k$ . Hence, the single rule “assign the  $2^{n-1}$  populations with the largest values of  $k$  to the  $|0\rangle$  block” simultaneously maximizes  $p'$  when  $\alpha > 0$  (producing a larger positive  $\alpha'$ ) and minimizes  $p'$  when  $\alpha < 0$  (producing a more negative  $\alpha'$ ). In both cases, this increases  $|\alpha'|$  while preserving the sign of the polarization. This completes the proof of optimal unitary bidirectional compression. ■

**Explicit Permutation Rule.** Theorem 1 establishes the optimality condition for bidirectional entropy compression. The permutations required to enhance  $|\alpha'|$  are obtained by swapping entries between the first and second blocks whenever a state in the first block, of the form  $|0y\rangle$  with  $y \in \{0, 1\}^{n-1}$ , has fewer zeros in its full bit string than some state in the second block, of the form  $|1y'\rangle$  with  $y' \in \{0, 1\}^{n-1}$ .

Note that for the case of  $\alpha > 0$ , the results coincide with those of conventional unidirectional cooling, which aims to enhance the ground-state population [19,23]. There is no unique permutation that achieves the maximum value of  $\alpha'$ . A convenient way to implement an optimal unitary is the so-called Mirror Protocol [26,51], which prescribes swapping each state  $|0y\rangle$  with its complementary bitstring  $|1\tilde{y}\rangle$  whenever the number of zeros  $k$  in the binary string  $0y$  satisfies  $k < n - k$ . The complement  $\tilde{y}$  is obtained by flipping every bit of  $y$  (i.e.,  $0 \leftrightarrow 1$  at each position). For example, in the case of three qubits, one such permutation swaps the populations of the states  $|011\rangle$  and  $|100\rangle$ . The explicit circuit implementation of this operation on three qubits, denoted  $U_{C_3}$ , is shown in Fig. 4.

The following corollary provides an alternative explicit population ordering that also achieves the optimal increase in  $|\alpha'|$ , as follows directly from Theorem 1. Note that this optimal increase in  $|\alpha'|$  can be realized by different permutations, so the unitary achieving the optimal value is not

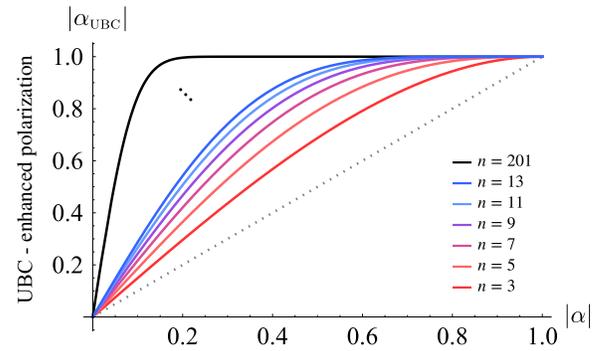


FIG. 2. Enhanced polarization magnitude  $|\alpha_{\text{UBC}}|$  resulting from optimal bidirectional entropy compression unitary, plotted as a function of the initial  $|\alpha|$  for different numbers of qubits. The black dotted line represents the baseline, initial polarization. For even numbers of  $n = 2j$ , the result is the same as for the previous odd number  $2j - 1$ .

unique; it may be chosen according to additional optimization criteria, such as energy cost or the number of required swaps.

**Corollary 1 (Unitary Bidirectional Cooling via Nonincreasing Binomial Weights on  $n$  Identical Qubits).** The permutation unitary that reorders the diagonal elements of the joint state  $\rho^{\otimes n}$ , defined in Theorem 1, according to the sequence  $p^n, p^{n-1}(1-p), p^{n-2}(1-p)^2, \dots, (1-p)^n$  implements an optimal unitary for bidirectional cooling of the target qubit. Each value  $p^k(1-p)^{n-k}$  appears with multiplicity  $\binom{n}{k}$ .

The ordering given in Corollary 1 represents one possible arrangement that achieves the optimal value of  $|\alpha'|$  for bidirectional cooling of the target qubit. Any permutation of the populations within the first half—or within the second half—of the diagonal of the joint state yields the same value of  $|\alpha'|$  for the target qubit.

The optimal enhanced polarization of a target qubit in a string of  $n$  qubits, achieved via unitary bidirectional cooling, with initial polarization  $\alpha$ , is given by

$$\alpha \longrightarrow \alpha'_{\text{UBC}} = 2 \left[ \sum_{i=0}^{\lfloor (n-1)/2 \rfloor} \binom{n}{i} p^{n-i} (1-p)^i \right] - 1, \quad (14)$$

where  $p = (1 + \alpha)/2$ , and  $\lfloor \cdot \rfloor$  denotes the floor function. One finds that  $\alpha'_{\text{UBC}}$  takes the same value for  $n = 2j - 1$  and  $n = 2j$ , for all  $j = 1, 2, 3, \dots$  Figure 2 illustrates the enhanced polarization magnitude  $|\alpha'_{\text{UBC}}|$  as a function of the initial polarization  $|\alpha|$  for different values of  $n$ . Note that when  $\alpha > 0$ , the optimal enhanced polarization coincides with the standard AC (unitary cooling) limit for  $n$  identical qubits  $\alpha'_{\text{AC}}$  [22,23,26].

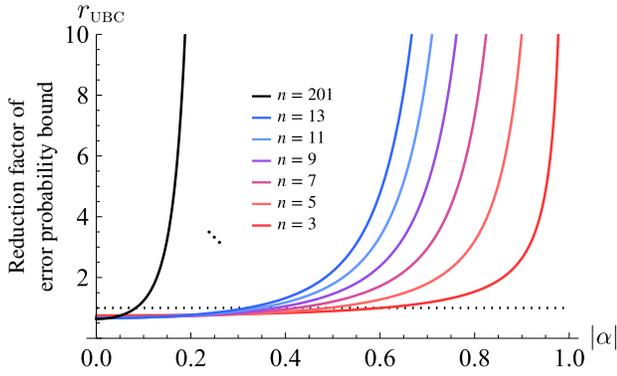


FIG. 3. Reduction factor of the sampling error-probability bound  $r_{\text{UBC}}$  achieved by enhancing the polarization via unitary bidirectional cooling on  $n$  qubits, plotted as a function of the initial polarization  $|\alpha|$ , using the error precision  $\epsilon(\alpha) = |\alpha|$ . Curves are shown for different values of  $n$ .

The enhanced polarization can be approximated as

$$\alpha'_{\text{UBC}} \simeq \text{erf}(\xi) := \frac{2}{\sqrt{\pi}} \int_0^\xi e^{-t^2} dt, \quad (15)$$

where  $\xi = n\alpha/\sqrt{2n(1-\alpha^2)}$  [22].

In the context of classification tasks, after enhancing the polarization via unitary bidirectional cooling on  $n$  qubits, the sampling error probability bound in Eq. (6) is reduced by the following factor:

$$r_{\text{UBC}} = \frac{1-\alpha^2}{1-\alpha'^2_{\text{UBC}}} \left( \frac{\epsilon(\alpha'_{\text{UBC}})}{\epsilon(\alpha)} \right)^2 \frac{1}{n}, \quad (16)$$

where  $\epsilon(\alpha)$  represents the requested precision. This comparison assumes equal total qubit resources; in other words, it contrasts  $k$  independent repetitions using the original qubits with  $k/n$  repetitions using the polarization-enhanced qubits. As illustrated in Fig. 3, the reduction factor can be substantial, with significant improvements that grow systematically as the number of qubits  $n$  increases. This demonstrates a substantial performance gain enabled by unitary bidirectional cooling across most of the initial polarization range.

The behavior of  $r_{\text{UBC}}$  in the different polarization regimes is as follows [using the error precision associated with the bound in Eq. (6), where  $\epsilon(\alpha) = |\alpha|$ ; further details are provided in Appendix A 2]: **(1) In the high-polarization regime**, the UBC protocol yields a very large improvement, and in the limit where the polarization approaches unity, the improvement factor diverges. **(2) In the intermediate-polarization regime**, the UBC protocol shows a clear advantage ( $r_{\text{UBC}} > 1$ ). In this region, one can employ a suitable number  $n$  of qubits for which UBC yields an improvement. Moreover, the enhancement factor

grows approximately exponentially, details are provided in Appendix A 2. **(3) In the low-polarization regime**, the enhanced polarization satisfies  $\alpha'_{\text{UBC}} \approx \alpha\sqrt{2n/\pi}$ , which yields  $r_{\text{UBC}} \approx 2(1-\alpha^2)/(\pi-2n\alpha^2)$ . In the limit  $|\alpha| \approx 0$ , this reduces to  $r_{\text{UBC}} \approx 2/\pi < 1$ , indicating that UBC does not provide an improvement in the extreme low-polarization regime.

The area under the  $r_{\text{UBC}}$  curves over the full range of initial polarization  $\alpha$  diverges for all  $n \geq 3$ . This indicates that, although the UBC unitary strategy does not provide an advantage in the very low-polarization regime, the sampling error probability bound is improved on average when the polarization is uniformly distributed, for all  $n \geq 3$ , with greater improvement as  $n$  increases. Moreover, the size of the low-polarization region with no advantage decreases as  $n$  grows. To go beyond unitary bidirectional cooling, we introduce a cyclic protocol inspired by HBAC, described in the next section.

## VI. BIDIRECTIONAL QUANTUM REFRIGERATOR

We define the *Bidirectional Quantum Refrigerator (BQR)* as a cyclic protocol that acts on an  $n$ -qubit register to increase the magnitude  $\alpha$  in a target qubit while preserving its sign without any prior knowledge of that sign. The BQR consists of multiple rounds of (i) unitary bidirectional cooling operations and (ii) qubit-reset steps. After the target qubit is enhanced and removed from the system, the remaining qubits in their updated state are recycled to prepare subsequent target qubits, serving as the working body of the refrigerator that can operate in a steady state.

**Setup and BQR Protocol**—The system setup for each cycle consists of a string of  $n$  qubits: one target qubit to be improved,  $m$  reset qubits that can be replaced with fresh qubits from the sample, and  $n-m-1$  auxiliary qubits that assist in the entropy compressions. The BQR protocol purifies a target qubit through  $N_{\text{rounds}}$  rounds, each consisting of two steps: first, a sequence of unitaries that increase the dominant bias of the target qubit, followed by the replacement of the  $m$  reset qubits, which pumps the excess entropy out of the system. After the  $N_{\text{rounds}}$ , the enhanced target qubit is removed and is ready to be measured for a classification task, while the remaining  $n-1$  qubits are reused to prepare the next target qubit. A new input qubit, in initial state  $\rho_\alpha$ , is introduced at the end of the chain to complete the  $n$ -qubit register and repeat the process. This new input qubit becomes the last element of the chain, and the qubit that was previously second—already prepared in the earlier cycle—now serves as the new target qubit, as depicted in Fig. 4. Recycling the  $n-1$  qubits reduces both the required qubit resources and the number of steps needed to prepare subsequent targets. Once in their steady state  $\tilde{\rho}_{\text{QR}}$ , these qubits effectively form the working body of a

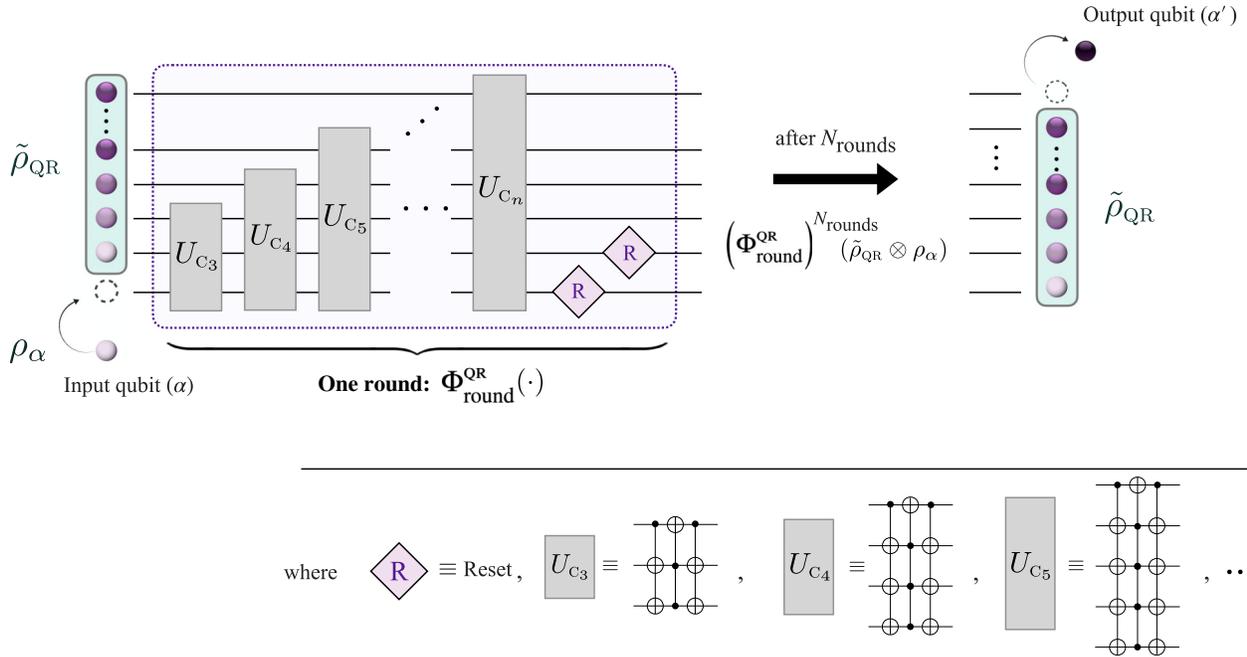


FIG. 4. Bidirectional quantum refrigerator. Circuit diagram of the **progressive boundary entropy compression-bidirectional quantum refrigerator** protocol acting on an  $n$ -qubit register with  $m$  reset qubits ( $m = 2$  shown). The unitary stage consists of a stairlike sequence of unitaries  $U_{C_3}, \dots, U_{C_n}$ , where each  $U_{C_j}$  performs a boundary entropy-compression step by swapping the states  $|0\rangle|1\rangle^{\otimes(j-1)}$  and  $|1\rangle|0\rangle^{\otimes(j-1)}$  on the last  $j$  qubits. After the unitary stage, the  $m$  reset qubits are refreshed, completing one round of the protocol. The protocol runs for  $N_{\text{rounds}}$  to prepare the target qubit; once prepared, the target qubit is extracted ready for the classification task, and the process restarts by recycling the remaining  $n - 1$  qubits to prepare subsequent target qubits.

refrigerator that enables repeated initialization, as detailed in the following sections and appendices.

In this work, we introduce a cooling protocol within the BQR framework: the Progressive Boundary Entropy Compression BQR, defined as follows.

### A. The Progressive Boundary Entropy Compression BQR

The **Progressive Boundary Entropy Compression BQR** (PBEC-BQR) protocol is a bidirectional quantum refrigerator running with identical rounds, whose unitary step consists of a sequence of unitaries  $U_{C_3}, \dots, U_{C_n}$ . Each  $U_{C_j}$  swaps the states  $|0\rangle|1\rangle^{\otimes(j-1)}$  and  $|1\rangle|0\rangle^{\otimes(j-1)}$ , and the sequence is implemented in a stairlike manner: starting with  $U_{C_3}$  acting on the last three qubits of the register, followed by  $U_{C_4}$  on the last four qubits, and so on, up to  $U_{C_n}$ , as depicted in Fig. 4. Then,  $m$  qubits are fully reset by replacing them with qubits from the sample in the state  $\rho_\alpha$ .

Note that each  $U_{C_j}$  performs an entropy-compression step on the  $j$ -qubit subsystem by swapping a single pair of states—those located at the boundary between the middle energy levels of the  $j$ -qubit spectrum. This motivates the name *boundary entropy compression on  $j$  qubits*.

In each round, the global effect of the set of unitaries is given by

$$U_{\text{QR}}(n) = U_{C_n}(\mathbb{1}_2 \otimes U_{C_{(n-1)}})(\mathbb{1}_2 \otimes U_{C_{(n-2)}}) \cdots (\mathbb{1}_2 \otimes U_{C_3}).$$

This can be represented as the following block matrix:

$$U_{\text{QR}}(n) = \begin{bmatrix} \mathbb{1}_3 & 0 & 0 \\ 0 & \sigma_{\text{QR}}(n) & 0 \\ 0 & 0 & \mathbb{1}_3 \end{bmatrix}, \quad (17)$$

where  $\mathbb{1}_3$  is the identity matrix with a diagonal of size three,  $\sigma_{\text{QR}}(3) = \sigma_x$ , and

$$\sigma_{\text{QR}}(n) := \begin{bmatrix} \mathbb{1}_2 \otimes \begin{bmatrix} \sigma_{\text{QR}}(n-1) & 0 \\ 0 & \mathbb{1}_2 \end{bmatrix} & 0 \\ 0 & \sigma_x \end{bmatrix}, \quad \text{for } n > 3. \quad (18)$$

The effect of a single round on a system of  $n$  qubits in the total state  $\rho_T$  is given by

$$\rho_T \rightarrow \Phi_{\text{round}}^{\text{QR}}(\rho_T) := \text{Tr}_m \left[ U_{\text{QR}}(n) \rho_T U_{\text{QR}}^\dagger(n) \right] \otimes \rho_\alpha^{\otimes m}. \quad (19)$$

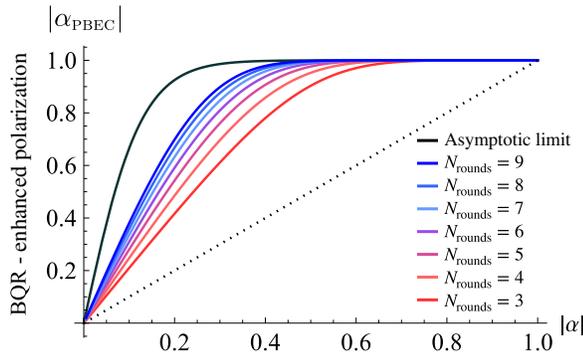


FIG. 5. Enhanced polarization  $\alpha_{\text{PBEC}}$  achieved by the progressive boundary entropy compression-bidirectional quantum refrigerator protocol for  $n = 5$  qubits, shown as a function of the initial polarization magnitude  $|\alpha|$  for different numbers of rounds. The black dotted line indicates the initial polarization baseline, while the solid black line marks the asymptotic polarization.

The BQR reuses the output system for subsequent preparations. Namely, after preparing an enhanced qubit, the remaining  $n - 1$  qubits are recycled as the updated system for the quantum refrigerator, along with a new fresh qubit to complete the  $n$  qubits needed for the circuit to run again. The recycled set of qubits reaches an output steady-state  $\tilde{\rho}_{\text{QR}}(n, \alpha, N_{\text{rounds}}, m)$ , which reduces the number of resources required to prepare enhanced qubits ( $\sim mN_{\text{rounds}} + 1$ , instead of  $n + mN_{\text{rounds}}$ ), and improves the convergence rate.

The enhanced polarization of the target qubit,  $\alpha_{\text{PBEC}}$ , depends on the refrigerator configuration  $(n, N_{\text{rounds}}, m)$  and on the initial polarization of the sample qubits  $\alpha$ . Figure 5 shows this enhancement for  $n = 5$  and various values of  $N_{\text{rounds}}$  as a function of  $|\alpha|$ . Additional details are provided in Appendix B. In the cooling limit, the (PBEC)–BQR protocol can asymptotically achieve—albeit at a different rate and for both signs of  $\alpha$ —the maximum achievable polarization magnitude of the PPA-HBAC protocol, given in Eq. (12) (details are provided in Appendix B).

It is important to emphasize that, for our purposes in classification tasks, operating near the cooling limit is not necessary. Instead, the optimal configuration for classification tasks involves using only a small number of rounds, striking a balance between achieving significant polarization enhancement and using qubit resources efficiently to maximize the reduction in sampling error. Notably, the first few rounds provide the largest gains in polarization, whereas subsequent rounds yield diminishing returns.

The sampling-error probability-reduction ratio, based on the bound in Eq. (6) and the enhanced polarization achieved by the PBEC-QR protocol, is given by

$$r_{\text{PBEC}} = \frac{\alpha^{-2} - 1}{\alpha_{\text{PBEC}}^{-2}(n, m, N_{\text{rounds}}, \alpha) - 1} \cdot \frac{k_{\text{BQR}}}{k_{\text{C}}}, \quad (20)$$

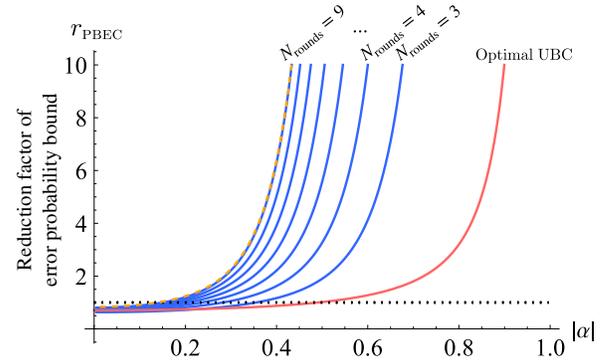


FIG. 6. Reduction factor of the error-probability bound for the progressive boundary entropy compression-bidirectional quantum refrigerator with  $n = 5$  qubits, shown in blue for different numbers of rounds as a function of the initial  $|\alpha|$ . The pink line shows the performance of unitary bidirectional cooling, and the yellow dashed line corresponds to an adaptive bidirectional quantum refrigerator using optimal per-round compressions (shown here for  $N_{\text{rounds}} = 9$ ). The progressive boundary entropy compression-bidirectional quantum refrigerator, despite using identical rounds, closely matches the adaptive scheme for  $n = 5$ ,  $m = 2$ , and  $N_{\text{rounds}} = 9$ —whereas deviations from this number of rounds lead to a visible performance gap (not shown).

where  $\alpha_{\text{PBEC}}(n, m, N_{\text{rounds}}, \alpha)$  denotes the PBEC-BQR’s enhanced polarization,  $k_{\text{C}}$  is the number of shots in the baseline method without BQR enhancement, and  $k_{\text{BQR}}$  is the number of measurement shots used in the BQR scheme, assuming both approaches have the same qubit resources.

Figure 6 illustrates the reduction factor of the error-probability bound for the PBEC-BQR operating with  $n = 5$  qubits as a function of the initial  $|\alpha|$ , shown in blue for various numbers of rounds. The pink line represents the improvement achieved through an optimal single-shot entropy compression on the same register, while the yellow dashed line shows an example for  $N_{\text{rounds}} = 9$  using adaptive compressions in a BQR scheme (where each round applies its own optimal entropy-compression unitary), obtained from numerical simulations. In this configuration ( $n = 5$ ,  $N_{\text{rounds}} = 9$ ,  $m = 2$ ), the performance of the PBEC-BQR closely approaches the adaptive BQR behavior, exhibiting an almost exact agreement. For other configurations with fewer rounds, the gap becomes more noticeable (not depicted in the figure), although still small, and our protocol remains more practical and experimentally feasible due to its use of identical rounds.

The PBEC-BQR provides a substantial enhancement of the error-probability reduction factor compared to the Unitary Bidirectional Cooling (UBC) protocol, offering consistently superior performance across the full range of initial polarizations. Moreover, achieving the optimal configuration does not require operating near the cooling limit or using an excessively large number of rounds, as

shown in Eq. (20). Instead, it requires balancing significant polarization enhancement with the efficient use of qubit resources to minimize sampling error.

The PBEC-BQR significantly improves the sampling-error probability bound, both on average over all initial polarizations for a uniformly distributed  $\alpha$  and locally within the medium- and high-polarization regimes. Although the protocol does not provide an advantage in the very low-polarization regime, the region in which it is advantageous can be expanded by optimizing the number of qubits of the register and the number of rounds.

### B. BQR with $k$ -local compressions

In this section, we introduce a more experimentally feasible BQR protocol, in which the unitary step is carried out using only  $k$ -local bidirectional operations acting on  $k$  neighboring qubits at a time. We refer to this protocol as the  $k$ -local Bidirectional Quantum Refrigerator.

The  $k$ -local *Bidirectional Quantum Refrigerator* ( $k$ -local BQR) protocol is a bidirectional quantum refrigerator that operates in identical rounds. Its unitary step consists of a sequence of local compression unitaries  $U_{C_k}$ , which swap the states  $|0\rangle|1\rangle^{\otimes(k-1)}$  and  $|1\rangle|0\rangle^{\otimes(k-1)}$  of  $k$  neighboring qubits. These unitaries are applied in a staircaselike pattern across the register, starting from the last qubits and progressing upward, as illustrated in Fig. 7 for the case  $k = 3$ . By restricting the compression to  $k$ -local neighborhoods, this scheme becomes substantially easier to implement on hardware while still preserving the core performance advantages of the full BQR protocol. After the unitary step,  $m$  qubits are fully reset by replacing them with fresh qubits prepared in the state  $\rho_\alpha$ .

Given its practical relevance, we focus in this section on the  $k = 3$  implementation, employing  $U_{C_3}$  as the three-local compression unitary on an  $n$ -qubit register with  $m = 2$  reset qubits, as depicted in Fig. 7. In the three-local BQR,

the global unitary corresponding to one round is given by

$$U_{\text{QR}_{3\text{local}}} = (U_{C_3} \otimes \mathbb{1}_{2^{n-3}})(\mathbb{1}_2 \otimes U_{C_3} \otimes \mathbb{1}_{2^{n-4}}) \cdots (\mathbb{1}_{2^{n-3}} \otimes U_{C_3}).$$

In the cooling limit, the three-local BQR protocol working on a  $n$  qubit register has an asymptotic target polarization of

$$\alpha_{\text{QR}_{3\text{local}}}^\infty = \frac{2p_\alpha^{F_n}}{\left(\sum_{i=0}^{F_n-1} (-1)^i \binom{F_n}{i} p_\alpha^i\right) - 1},$$

where  $F_n$  is the  $n$ th Fibonacci number ( $F_n = F_{n-1} + F_{n-2}$ , for  $n \geq 3$ , with  $F_1 = F_2 = 1$ ; see details in Appendix B).

Importantly, as mentioned in the previous section, it is not necessary for the quantum refrigerator to operate near the cooling limit. Instead, the optimal configuration involves only a small number of rounds, balancing substantial polarization enhancement with the efficient use of qubit resources to maximize the reduction in sampling error.

The reduction factor of the error-probability bound obtained using the  $k$ -local BQR protocol, denoted  $r_{\text{QR}_{k\text{-local}}}$ , has the same structure as Eq. (20). Using the enhanced polarization  $\alpha_{\text{QR}_{k\text{-local}}}$  corresponding to  $k$ -local BQR. Figure 8 shows the reduction factor of the error-probability bound for the three-local BQR operating with  $n = 5$  qubits for different numbers of rounds, as a function of the initial  $|\alpha|$ . The pink line represents the improvement achieved through UBC on  $n = 5$ , while the yellow dashed line indicates an upper bound for  $N_{\text{rounds}} = 9$  obtained from simulations with an adaptive-round BQR. Note that although the PBEC-BQR protocol introduced in the previous subsection achieves a greater improvement (see Fig. 6 for results under the same system setup), the  $k$ -local BQR protocol still provides a substantial reduction in the error-estimation bound, both on average for uniformly distributed polarizations and outside the low-polarization regime.

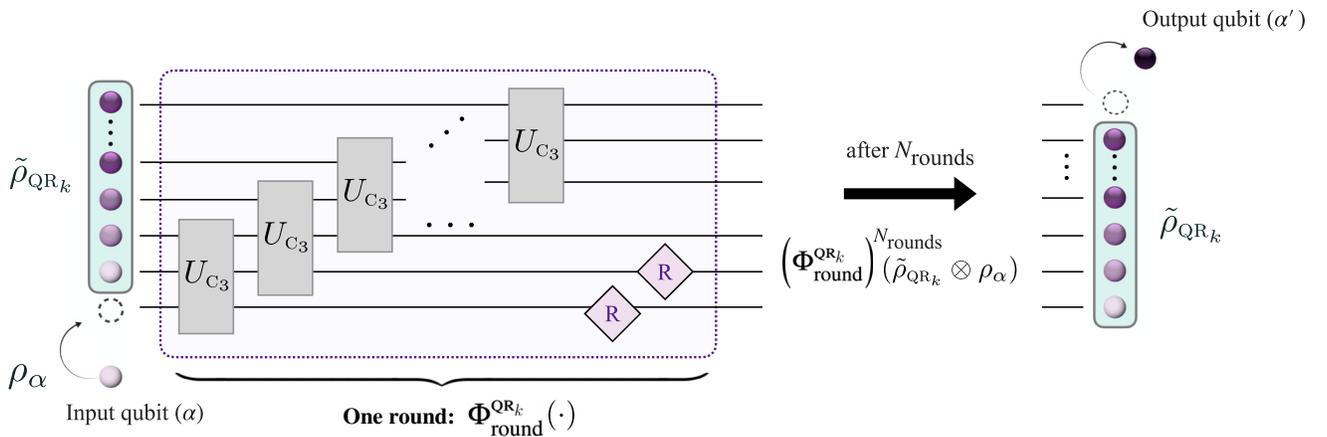


FIG. 7.  $k$ -local Bidirectional Quantum Refrigerator ( $k$ -local BQR). The image illustrates the example for  $k = 3$ .

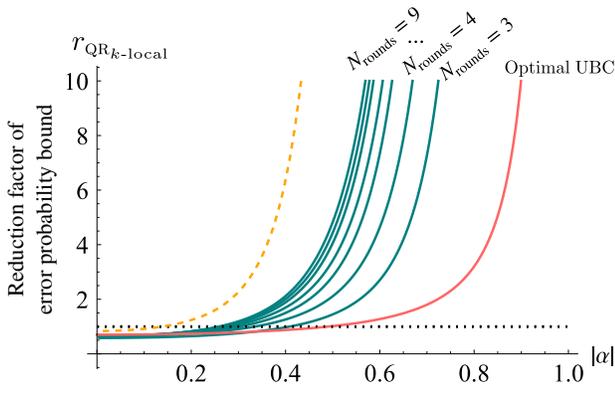


FIG. 8. Reduction factor of the error-probability bound for the three-local bidirectional quantum refrigerator with  $n = 5$ , shown in green for different numbers of rounds as a function of the initial  $|\alpha|$ . The pink line shows the improvement achieved through single-shot unitary bidirectional cooling on an  $n = 5$  register, while the yellow dashed line indicates the upper bound obtained from simulations of the adaptive-round bidirectional quantum refrigerator with  $N_{\text{rounds}} = 9$ . Although the performance of the three-local bidirectional quantum refrigerator shows a noticeable gap relative to this upper bound—reflecting its reduced optimality—it remains significantly more practical to implement while still achieving substantial reductions in the error-probability bound.

### C. Overview of bidirectional cooling methods

Here, we summarize the hierarchy of the bidirectional cooling mechanisms introduced in this work. These schemes differ in their operational structure, resource requirements, and practicality, yet all share the key feature of sign-preserving, bidirectional cooling. A schematic overview of the protocols is shown in Fig. 9.

**(a) Unitary Bidirectional Cooling (UBC).** UBC is a single-shot entropy-compression operation applied to a register of  $n$  qubits to increase the magnitude of the polarization  $|\alpha|$  of a target qubit while preserving its unknown sign. It serves as the fundamental primitive underlying all protocols developed in this work.

**(b) Bidirectional Quantum Refrigerator (BQR).** The BQR extends UBC into a multiround, cyclic cooling protocol. Each round consists of one UBC step acting on an  $n$ -qubit register, followed by the reset of  $m$  qubits. After  $N_{\text{rounds}}$ , the enhanced target qubit is extracted, and the remaining  $n - 1$  qubits are recycled as the working body of the refrigerator, enabling the preparation of multiple purified outputs. We introduce three explicit variants of the BQR protocol, which differ in the particular implementation of the unitary step:

**(b.1) Progressive Boundary Entropy Compression BQR (PBEC-BQR).** PBEC-BQR consists of repeated identical rounds in which the unitary cooling step is implemented by applying a sequence of unitaries  $U_{C_3}, \dots, U_{C_n}$  in a staircaselike manner: starting with  $U_{C_3}$  acting on the last three qubits of the register, followed by  $U_{C_4}$  on the

last four, and so on up to  $U_{C_n}$ , as depicted in Fig. 4. Each  $U_{C_j}$  swaps the states at the boundary between the first and second blocks of the spectrum of the  $j$ -qubit register, i.e.,  $|0\rangle|1\rangle^{\otimes(j-1)}$  and  $|1\rangle|0\rangle^{\otimes(j-1)}$ .

**(b.2) Local Bidirectional Quantum Refrigerator ( $k$ -local BQR).** The  $k$ -local BQR is a more experimentally feasible variant in which the bidirectional compression stage is implemented using  $k$ -local unitaries acting on neighboring qubits in a staircaselike pattern. This protocol operates in identical rounds. Its unitary step consists of a sequence of local compression unitaries  $U_{C_k}$ , each swapping  $|0\rangle|1\rangle^{\otimes(k-1)}$  and  $|1\rangle|0\rangle^{\otimes(k-1)}$ , on  $k$  neighboring qubits. These unitaries are applied sequentially across the register, starting from the last qubits and moving upward, as illustrated in Fig. 7 for  $k = 3$ .

**(b.3) Adaptive BQR.** The adaptive BQR is a BQR protocol in which the rounds are not identical, instead, the unitary step may vary from one round to the next depending on the state of the system. In this work, we introduce the concept and study numerical examples to compare our explicit protocols against the optimal adaptive BQR obtained numerically for the scenarios presented here.

*Comparison of Resource Requirements and Cooling limits.* In terms of resource usage, UBC requires only an  $n$ -qubit register to produce a single enhanced target qubit. In contrast, the BQR schemes require additional qubits to sustain the cyclic rounds of cooling. To generate  $k_{\text{BQR}}$  enhanced qubits, a BQR scheme uses  $k_{\text{BQR}}(mN_{\text{rounds}} + 1) + n$  qubits. Regarding cooling limits, the strongest asymptotic cooling is achieved by the PBEC-BQR and the optimal adaptive BQR, followed by the  $k$ -local BQR, and finally the single-shot UBC. (We note, however, that if UBC were implemented as a single global unitary on all the qubits used in the BQR simultaneously, it would achieve the best polarization enhancement, but such a global operation is not practical to implement.)

#### 1. Comparison of error-probability-bound reduction

All schemes provide substantial improvements in the reduction factor of the error-probability bound in both the high- and intermediate-polarization regimes, with the BQR family outperforming the UBC. In the low-polarization regime, however, the enhanced polarization does not yield a meaningful improvement. This performance region can be extended by appropriately adjusting protocol parameters, although the extreme low-polarization limit remains challenging to cover. In addition, the sampling-error probability bound is improved on average when the polarization is uniformly distributed.

#### 2. Comparison of feasibility

Among the BQR variants, the  $k$ -local BQR offers the most practical implementation. By restricting compression operations to  $k$ -local neighborhoods, it becomes

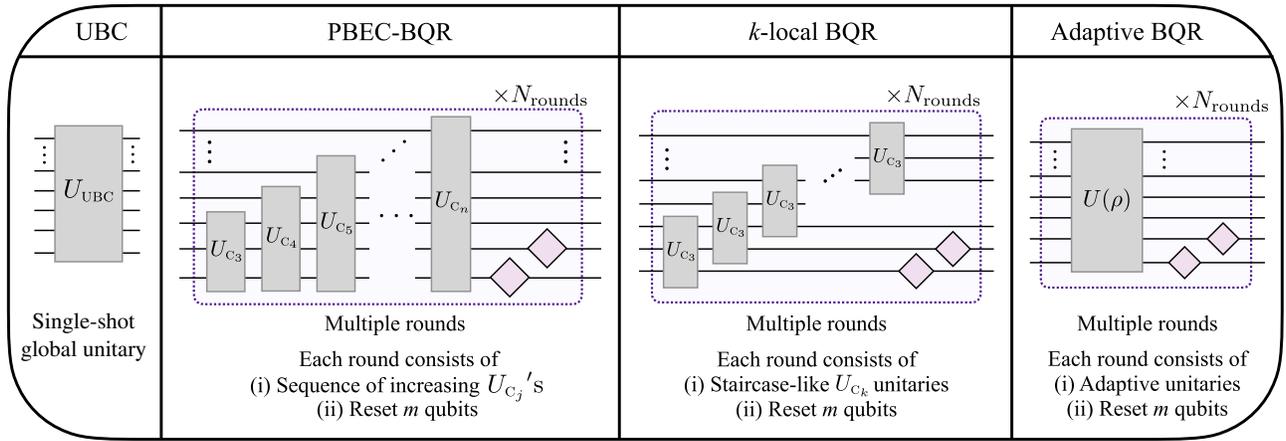


FIG. 9. Schematic overview of the bidirectional cooling methods. From left to right, the figure illustrates: (i) the unitary bidirectional cooling protocol (Sec. VA), implemented as a single-shot global unitary; (ii) the progressive boundary entropy compression-bidirectional quantum refrigerator (Sec. VIA), whose unitary step consists of a sequence of increasing  $U_{C_j}$  operations (complete circuit shown in Fig. 4); (iii) the  $k$ -local bidirectional quantum refrigerator protocol (Sec. VIB), implemented via staircaselike  $k$ -local unitaries  $U_{C_k}$  acting on neighboring qubits (full circuit in Fig. 7); and (iv) the adaptive bidirectional quantum refrigerator, in which the unitary step may vary from one round to the next depending on the state of the system. In this work, this protocol is used only numerically to obtain optimal benchmarks for comparison with the explicit constructions above. All bidirectional quantum refrigerator variants operate in multiple rounds, each consisting of a unitary compression stage followed by a reset of  $m$  qubits.

significantly more hardware-friendly while still preserving the essential performance advantages of the more general BQR protocols. Moreover, for certain system configurations, the  $k$ -local BQR can even yield a better reduction of the error-probability bound than the PBEC-BQR.

*Comparison with Previous Algorithmic Cooling Techniques.* While certain unitaries appearing in AC/HBAC constructions may incidentally exhibit bidirectional behavior under specific circumstances, this concept has neither been identified nor formalized, and it has not been used to design explicit sign-preserving cooling protocols. In this work, we introduce this notion explicitly and develop the first systematic family of bidirectional cooling methods. Previous HBAC protocols were conceived as schemes that increase the population of  $|0\rangle$ , with optimal constructions typically requiring prior knowledge of the sign of  $\alpha$ . In contrast, UBC, BQR, and  $k$ -local BQR provide explicit bidirectional, sign-preserving cooling without requiring such knowledge. This distinction is essential in supervised QML settings, where the sign of  $\alpha(x, \theta)$  encodes the label or the gradient direction and must remain unknown. The BQR protocols introduced here thus constitute entirely new classes of algorithmic cooling schemes, tailored specifically to QML applications.

## VII. CLASSIFICATION PERFORMANCE EVALUATION

To evaluate the effectiveness of the proposed BQR in practical QML scenarios, we conducted a series of simulated experiments utilizing Qiskit to compare classification

performance with and without BQR. These experiments were designed to assess whether enhancing the polarization of a measurement qubit through BQR translates to tangible improvements in classification accuracy.

We focus on the three-local BQR and compare its ability to classify the target data accurately in the finite measurement shots against the conventional sampling method. This is a conservative choice: if the three-local version outperforms the conventional baseline, then the full BQR protocol—which provides stronger cooling—is expected to perform even better. Specifically, we used BQR with the following parameters: five system qubits ( $n = 5$ ), two reset qubits ( $m = 2$ ), and two rounds ( $N_{\text{rounds}} = 2$ ).

To isolate the impact of finite sampling error, which BQR is intended to reduce, from the expressivity or trainability of QML models themselves, we simplified the problem setup. Specifically, we assumed access to the final quantum state in the form of a single-qubit reduced density matrix whose  $Z$ -polarization encodes the classification signal, as described in Sec. II A.

The number of measurement shots with BQR was intentionally limited (e.g., to 10 or 100) to reflect realistic resource constraints in QML, where repeated sampling is an overhead not present in classical machine learning. In this context, the shot count can be viewed as an additional computational cost required to extract reliable output from a quantum classifier. Our goal is to evaluate how effectively BQR mitigates this cost by improving the reliability of predictions from fewer measurements. To ensure a fair comparison, we allocated proportionally more measurement shots to the conventional (non-BQR)

TABLE I. Classification accuracy (%) with and without BQR for  $n = 5$ ,  $m = 2$ , and  $N_{\text{rounds}} = 2$ , using the three-local BQR. The BQR-based method used  $k_{\text{BQR}} \in \{10, 100\}$  measurement shots, while the conventional baseline used  $k_c = k_{\text{BQR}} \times m \times (N_{\text{rounds}} - 1) + n$  shots, namely, 25 and 205, respectively. Reported values are the mean and standard deviation over 100 random trials.  $p$ -values are computed using Welch’s  $t$ -test to assess the statistical significance of accuracy differences between the BQR and baseline methods.

Dataset	BQR	Classification accuracy (%)	
		$k_{\text{BQR}} = 10$	$k_{\text{BQR}} = 100$
Uniform	With	$95.8 \pm 1.8$	$99.3 \pm 0.7$
	Without	$93.1 \pm 2.3$	$97.7 \pm 1.3$
	$p$ -value	$6.8 \times 10^{-17}$	$4.6 \times 10^{-20}$
Gaussian	With	$96.7 \pm 1.2$	$95.6 \pm 0.5$
	Without	$94.3 \pm 1.9$	$94.9 \pm 0.8$
	$p$ -value	$2.0 \times 10^{-20}$	$2.7 \times 10^{-10}$
Iris (setosa vs virginica)	With	$99.7 \pm 0.6$	$100.0 \pm 0.0$
	Without	$98.5 \pm 1.0$	$99.8 \pm 0.4$
	$p$ -value	$6.1 \times 10^{-18}$	$3.4 \times 10^{-5}$
Wine (Barolo vs Grignolino)	With BQR	$85.2 \pm 2.6$	$87.5 \pm 1.4$
	Without	$83.6 \pm 2.6$	$86.3 \pm 1.8$
	$p$ -value	$3.7 \times 10^{-5}$	$1.2 \times 10^{-6}$
Handwritten digit (2 vs 5)	With BQR	$99.2 \pm 0.9$	$100.0 \pm 0.2$
	Without	$98.1 \pm 1.3$	$99.6 \pm 0.6$
	$p$ -value	$7.9 \times 10^{-12}$	$5.1 \times 10^{-8}$
Sonar	With BQR	$87.1 \pm 3.4$	$90.3 \pm 2.1$
	Without	$83.4 \pm 3.5$	$89.1 \pm 2.3$
	$p$ -value	$1.0 \times 10^{-12}$	$2.6 \times 10^{-4}$
Diabetes	With BQR	$73.0 \pm 4.3$	$75.8 \pm 3.6$
	Without	$70.7 \pm 4.7$	$74.7 \pm 3.6$
	$p$ -value	$3.9 \times 10^{-4}$	$2.9 \times 10^{-2}$

baseline to compensate for the additional resources used in BQR, including system and reset qubits and the number of cooling rounds. Specifically, the number of shots in the baseline method was scaled as:  $k_c = k_{\text{BQR}} \times m \times (N_{\text{rounds}} - 1) + n$ , where  $k_{\text{BQR}}$  is the number of measurement shots used in the BQR-based schemes.

To evaluate performance, we constructed binary classification tasks from several real-world and synthetic datasets. For each dataset, we sampled 50 data points from each class to create balanced tasks. We repeated this sampling and evaluation procedure 100 times, generating 100 independent classification tasks per dataset. This repetition serves as a randomized statistical ensemble, allowing us to perform hypothesis testing on the accuracy results to assess the statistical significance and robustness of observed improvements.

The simulation results are summarized in Table I. Across all datasets tested, the classifier with BQR consistently outperformed the conventional sampling baseline. Welch’s  $t$ -tests [52] confirmed that the accuracy improvements were statistically significant ( $p$ -value  $< 0.05$  in all

instances). Full details on the datasets and simulation setups are provided in Appendix F.

## VIII. CONCLUSIONS AND DISCUSSION

This study integrates quantum thermodynamics with QML to enhance sampling efficiency through a bidirectional quantum refrigerator technique. By conceptualizing quantum supervised learning as a thermodynamic cooling process, the proposed method significantly reduces the finite sampling error or the number of repetitions required for accurate classification, without the need for Grover-like operations. This practical approach applies to various QML models, including VQCs and quantum kernel methods, demonstrating its versatility and broad impact. Note that this connection goes beyond a formal analogy: the proposed cooling process is operationally implemented within a quantum circuit that physically reduces the entropy of the qubits. The technique is particularly suited for NISQ devices, addressing the current limitations of quantum hardware and making QML algorithms more practical and scalable. The interdisciplinary nature of this work, bridging QIP, thermodynamics, and data science, not only advances the field of QML but also opens new research avenues in quantum thermodynamics, particularly for advancing algorithmic cooling techniques.

An interesting conceptual connection can be drawn between our method and the data reuploading technique, a well-established strategy in QML for building nonlinear models. In classification, the distinguishability between two density matrices representing the ensemble of quantum states with different labels, which is quantified by the trace distance, critically impacts both the empirical and expected risks [53]. Since trace distance is contractive under completely positive and trace-preserving (CPTP) maps, the performance of a VQC with fixed state preparation followed by parameterized gates is fundamentally limited by the initial data encoding. Through the lens of Ref. [54], we observe that data reuploading can potentially alleviate this limitation. Specifically, the work shows that any data reuploading circuit can be mapped to an explicit form consisting of a fixed data encoding stage followed by a parameterized quantum circuit, where ancillary qubits unfold the repeated encodings. This structure effectively reshapes the data embedding in a higher-dimensional Hilbert space. Our protocol can also be viewed as this unfolding process that effectively encodes data in an extended Hilbert space. However, it is more directly designed to increase the distinguishability of label-conditioned states by engineering a system-bath interaction with a data-dependent thermal reservoir. Moreover, our method does not require access to the data encoding circuit and operates directly on quantum states, including those with unknown or fixed preparation procedures.

This connection raises intriguing avenues for future exploration. For instance, it may be possible to construct a data reuploading version of the BQR circuit, allowing for a trade-off analysis between circuit depth and qubit overhead. We also note that, for Pauli measurements, the final state of a data reuploading circuit can be expressed in the form of Eq. (2), implying that our method can be directly applied to multiple copies of such circuits to reduce finite sampling error without loss of generality.

A remaining challenge is determining whether the proposed PBEC-BQR cooling protocol is optimal for reducing finite sampling errors. If optimality cannot be established, further refinement of the protocol to enhance its performance presents a compelling direction for future research. Additionally, investigating how coherence and nonclassical correlations within the system and bath qubits can be harnessed to improve cooling efficiency offers an intriguing avenue. A detailed quantitative analysis of how our method mitigates the barren plateau effect also deserves further investigation. Note that our technique can be employed as an independent mechanism in conjunction with existing strategies for barren plateau mitigation [43,55]. Another open question is whether these cooling techniques can be adapted to reduce finite sampling errors in the estimation of quantum kernels. This application presents unique challenges, as the quantities to be estimated involve the kernel matrix elements, rather than simple binary outcomes. Consequently, the naive application of the proposed method, which primarily aids in sign estimation, would not directly suffice. Addressing this challenge will broaden the applicability of insights from quantum thermodynamics to QML.

## ACKNOWLEDGMENTS

This work was supported by Institute of Information & communications Technology Planning & evaluation (IITP) grant funded by the Korea government (No. 2019-0-00003, Research and Development of Core Technologies for Programming, Running, Implementing and Validating of Fault-Tolerant Quantum Computing System), the Yonsei University Research Fund of 2024 (No. 2024-22-0147), and the National Research Foundation of Korea (No. RS-2025-02309510), and the Ministry of Trade, Industry, and Energy (MOTIE), Korea, under the Industrial Innovation Infrastructure Development Project (No. RS-2024-00466693). N.A.R.-B. acknowledges funding from the European Research Council (Consolidator Grant ‘‘Cocoquest’’ No. 101043705); the European Union’s Marie Skłodowska-Curie Postdoctoral Fellowship (Grant Agreement No. 101204616, QTF-Ataulfo); and the Austrian Research Promotion Agency (FFG) through Project No. FO999914030 (MUSIQ), funded by the European Union—NextGenerationEU. N.A.R.-B. also acknowledges support from the Miller Institute for Basic

Research in Science at the University of California, Berkeley, during the initial months of the project. We thank Hyukjoon Kwon for helpful discussions.

## DATA AVAILABILITY

The data that support the findings of this article are openly available [82].

## APPENDIX A: BIDIRECTIONAL ENTROPY COMPRESSION

### 1. Schur-Horn theorem on the setting of Theorem 1

Let  $\lambda$  be the eigenvalue vector of  $\rho^{\otimes n}$  arranged in non-increasing order and let  $\mathbf{d}(U) = \text{diag}(U\rho^{\otimes n}U^\dagger)$  be the diagonal vector in the computational basis after any unitary  $U$ . By the Schur-Horn theorem,  $\mathbf{d}(U)$  is majorized by  $\lambda$ . Then, the sum of the largest  $m$  diagonal entries is bounded by the sum of the largest  $m$  eigenvalues:  $\sum_{j=1}^m d_{(j)}(U) \leq \sum_{j=1}^m \lambda_{(j)}$ .

Taking  $m = 2^{n-1}$  shows that, among all unitaries, the maximum possible value of the sum of *any*  $2^{n-1}$  diagonal entries equals the sum of the  $2^{n-1}$  largest eigenvalues. Partition the computational basis into two blocks: the block where the target qubit is in state  $|0\rangle$  (the ‘‘first half’’) and the block where it is in state  $|1\rangle$  (the ‘‘second half’’). Then

$$p'(U\rho^{\otimes n}U^\dagger) = \sum_{j \in \mathcal{I}_0} d_j(U),$$

where  $\mathcal{I}_0$  indexes the  $2^{n-1}$  diagonal positions of the  $|0\rangle$ -block. Hence, to *maximize*  $\alpha'$  when  $\alpha > 0$ , we must maximize this sum, which by the previous paragraph is achieved by assigning the  $2^{n-1}$  largest eigenvalues to the  $|0\rangle$ -block (the order within that block is irrelevant). Conversely, when  $\alpha < 0$ , the magnitude  $|\alpha'|$  increases as  $p'$  *decreases*; thus we must *minimize* the same sum, which is attained by assigning the  $2^{n-1}$  *smallest* eigenvalues to the  $|0\rangle$ -block.

Finally, observe that the eigenvalues of  $\rho^{\otimes n}$  for a register of  $n$  identical qubits are monotonic functions of  $k$ , but with opposite behavior in the two sign regimes: when  $\alpha > 0$  they increase with  $k$ , whereas when  $\alpha < 0$  they decrease with  $k$ . Therefore, the single rule ‘‘assign the  $2^{n-1}$  populations with the largest  $k$  to the  $|0\rangle$ -block’’ maximizes  $p'$  when  $\alpha > 0$  (yielding a larger positive  $\alpha'$ ) and minimizes  $p'$  when  $\alpha < 0$  (yielding a more negative  $\alpha'$ ), in both cases increasing  $|\alpha'|$  while preserving the sign. This proves the stated optimal bidirectional compression.

### 2. Approximation of the reduction of finite sampling error

In this subsection, we derive approximations for the reduction in the sampling-error probability bound after

applying an optimal entropy-compression operation on  $n$  qubits in different polarization regimes. The enhancement factor  $r_{\text{UBC}}(n, \alpha, \alpha_{\text{UBC}})$  defined in Eq. (16), together with the optimal enhanced polarization  $\alpha_{\text{UBC}}$  of Eq. (14) and the estimation error associated with the correct sign of the expectation value ( $\epsilon(\alpha) = |\alpha|$ ), can be rewritten as

$$r_{\text{UBC}}(\alpha, \alpha_{\text{UBC}}, n) = \frac{1}{n} \cdot \frac{1 - \alpha^{-2}}{1 - \alpha_{\text{UBC}}^{-2}}. \quad (\text{A1})$$

To approximate that expression in the different approximation regimes, we have the following considerations:  $\alpha_{\text{UBC}} = 2 \left[ \sum_{i=0}^{\lfloor (n-1)/2 \rfloor} \binom{n}{i} p^{n-i} (1-p)^i \right] - 1$ , where  $p = (1 + \alpha)/2$  and  $\lfloor \cdot \rfloor$  denotes the floor function. The enhanced polarization can be approximated to  $\alpha_{\text{UBC}} \simeq \text{erf}(\xi) := 2/\sqrt{\pi} \int_0^\xi e^{-t^2} dt$ , where  $\xi = n\alpha/\sqrt{2n(1-\alpha^2)}$ . The reduction factor of the probability error bound can be rewritten as

$$r_{\text{UBC}} \simeq \frac{1}{n} \cdot \frac{1 - \alpha^{-2}}{1 - \text{erf}\left(\frac{n\alpha}{\sqrt{2n(1-\alpha^2)}}\right)^{-2}}. \quad (\text{A2})$$

**(1) Small-polarization regime ( $\alpha \ll 1$ ).** In this regime, the argument of the error function satisfies  $\xi = (n\alpha/\sqrt{2n(1-\alpha^2)}) \simeq \sqrt{n/2}\alpha$ . Using the small- $\xi$  expansion of the error function,

$$\text{erf}(\xi) = \frac{2}{\sqrt{\pi}} \sum_{j=0}^{\infty} \frac{(-1)^j \xi^{2j+1}}{2j+1 j!} \simeq \frac{2}{\sqrt{\pi}} \xi = \frac{\sqrt{2n}\alpha}{\sqrt{\pi}}, \quad (\text{A3})$$

and substituting this into Eq. (A1), the enhancement factor becomes

$$r_{\text{UBC}} \simeq \frac{2}{\pi} \cdot \frac{1 - \alpha^2}{1 - \frac{2n\alpha^2}{\pi}} \sim \frac{2}{\pi} < 1, \quad (\alpha \rightarrow 0). \quad (\text{A4})$$

**(2) Intermediate-polarization regime.** For the intermediate polarization regime, we use the large-argument asymptotic expansion of the error function,

$$\text{erf}(\xi) = 1 - \frac{e^{-\xi^2}}{\xi\sqrt{\pi}} \sum_{j=0}^{\infty} (-1)^j \frac{(2j-1)!!}{(2\xi^2)^j}, \quad (\text{A5})$$

and, for  $\xi \gg 1$  (i.e.,  $e^{-\xi^2}/(\xi\sqrt{\pi}) \ll 1$ ), we approximate

$$\text{erf}(\xi) \simeq 1 - \frac{e^{-\xi^2}}{\xi\sqrt{\pi}}. \quad (\text{A6})$$

Writing  $\text{erf}(\xi) \approx 1 - \delta$  with  $\delta = e^{-\xi^2}/(\xi\sqrt{\pi})$ , we have

$$\text{erf}(\xi)^{-2} \approx (1 - \delta)^{-2} \approx 1 + 2\delta, \quad (\text{A7})$$

so the denominator of the reduction factor  $r_{\text{UBC}}$  simplifies to

$$n(1 - \text{erf}(\xi)^{-2}) \simeq n(1 - (1 + 2\delta)) = -\frac{2ne^{-\xi^2}}{\xi\sqrt{\pi}}. \quad (\text{A8})$$

Hence,  $r_{\text{UBC}} \simeq -(\xi\sqrt{\pi}/2n) e^{\xi^2} (1 - \alpha^{-2})$ , so that the behavior of  $r_{\text{UBC}}$  in this regime is dominated by an exponentially growing factor,  $r_{\text{UBC}} \propto e^{\xi^2}$ , where  $\xi^2 = n\alpha^2/[2(1-\alpha^2)]$ .

**(3) Large-polarization regime ( $\alpha \rightarrow 1$ ).** When the initial polarization approaches unity, the denominator of  $\xi = n\alpha/\sqrt{2n(1-\alpha^2)}$  tends to zero, causing the argument  $\xi$  to diverge. In this limit, the error function saturates,  $\text{erf}(\xi) \xrightarrow[\xi \rightarrow \infty]{} 1$ . Substituting  $\text{erf}(\xi) \simeq 1$  into Eq. (A1), and noting that  $\alpha_{\text{UBC}} \rightarrow 1$  as well, we obtain  $r_{\text{UBC}} \xrightarrow[\alpha \rightarrow 1]{} \infty$ .

Thus, in the large-polarization regime, the reduction factor diverges: optimal entropy compression produces an exponentially enhanced improvement in the sampling-error probability bound.

## APPENDIX B: EVOLUTION UNDER THE PBEC BIDIRECTIONAL QUANTUM REFRIGERATOR

The global effect of the set of unitaries in the PBEC-BQR protocol is given as follows:

$$U_{\text{QR}}(n) = U_{C_n}(\mathbb{1}_2 \otimes U_{C_{(n-1)}})(\mathbb{1}_{2^2} \otimes U_{C_{(n-2)}}) \cdots (\mathbb{1}_{2^{(n-3)}} \otimes U_{C_3}),$$

where each  $U_{C_j}$  has the effect of swapping the states  $|0\rangle|1\rangle^{\otimes(j-1)}$  and  $|1\rangle|0\rangle^{\otimes(j-1)}$ . Specifically, each unitary  $U_{C_j}$  has the following block matrix representation:

$$U_{C_j} = \begin{bmatrix} \mathbb{1}_{2^{(n-1)-1}} & 0 & 0 \\ 0 & \sigma_x & 0 \\ 0 & 0 & \mathbb{1}_{2^{(n-1)-1}} \end{bmatrix}. \quad (\text{B1})$$

The general expression of  $U_{\text{QR}}(n)$  can be obtained by replacing Eq. (B1) as follows:

For  $n = 3$ ,

$$U_{\text{QR}}(3) = U_{C_3} = \begin{bmatrix} \mathbb{1}_3 & 0 & 0 \\ 0 & \sigma_x & 0 \\ 0 & 0 & \mathbb{1}_3 \end{bmatrix}. \quad (\text{B2})$$

For  $n = 4$ ,

$$U_{\text{QR}}(4) = U_{C_4}(\mathbb{1}_2 \otimes U_{\text{QR}}(3)) = U_{C_4}(\mathbb{1}_2 \otimes U_{C_3}) \quad (\text{B3})$$

$$= \begin{bmatrix} \mathbb{1}_7 & 0 & 0 \\ 0 & \sigma_x & 0 \\ 0 & 0 & \mathbb{1}_7 \end{bmatrix} \left( \mathbb{1}_2 \otimes \begin{bmatrix} \mathbb{1}_3 & 0 & 0 \\ 0 & \sigma_x & 0 \\ 0 & 0 & \mathbb{1}_3 \end{bmatrix} \right) \quad (\text{B4})$$

$$= \begin{bmatrix} \mathbb{1}_3 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_x & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \mathbb{1}_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma_x & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbb{1}_2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_x & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \mathbb{1}_3 \end{bmatrix} = \quad (\text{B5})$$

$$= \begin{bmatrix} \mathbb{1}_3 & & 0 & & 0 & & 0 \\ 0 & \mathbb{1}_2 \otimes \begin{bmatrix} \sigma_x & 0 \\ 0 & \mathbb{1}_2 \end{bmatrix} & & & 0 & & 0 \\ 0 & & 0 & & & & \sigma_x \\ 0 & & 0 & & & & 0 \\ & & & & & & 0 \\ & & & & & & \mathbb{1}_3 \end{bmatrix}. \quad (\text{B6})$$

Similarly, for general  $n > 3$ ,

$$U_{\text{QR}}(n) = U_{C_n}(\mathbb{1}_2 \otimes U_{\text{QR}}(n-1)). \quad (\text{B7})$$

Thus,  $U_{\text{QR}}(n)$  can be represented in matrix form as the following block matrix:

$$U_{\text{QR}}(n) = \begin{bmatrix} \mathbb{1}_3 & 0 & 0 \\ 0 & \sigma_{\text{QR}}(n) & 0 \\ 0 & 0 & \mathbb{1}_3 \end{bmatrix}, \quad (\text{B8})$$

where  $\sigma_{\text{QR}}(3) = \sigma_x$ , and  $\sigma_{\text{QR}}(n > 3)$  is defined as

$$\sigma_{\text{QR}}(n) := \begin{bmatrix} \mathbb{1}_2 \otimes \begin{bmatrix} \sigma_{\text{QR}}(n-1) & 0 \\ 0 & \mathbb{1}_2 \end{bmatrix} & 0 \\ 0 & \sigma_x \end{bmatrix}, \quad \text{for } n > 3, \quad (\text{B9})$$

or equivalently,

$$\sigma_{\text{QR}}(n) := \begin{bmatrix} \mathbb{1}_{(2^{n-2}-2)} \otimes \begin{bmatrix} \sigma_x & 0 \\ 0 & \mathbb{1}_2 \end{bmatrix} & 0 \\ 0 & \sigma_x \end{bmatrix}, \quad \text{for } n > 3.$$

### 1. Evolution under PBEC-BQR

The effect of a single round on a system of  $n$  qubits in the state  $\rho$ , with  $m$  reset qubits, is given by

$$\rho \rightarrow \Phi_{\text{round}}^{\text{QR}}(\rho) := \text{Tr}_m \left[ U_{\text{QR}}(n) \rho U_{\text{QR}}^\dagger(n) \right] \otimes \rho_\alpha^{\otimes m}. \quad (\text{B10})$$

Then, the effect of the refrigerator operating with  $N_{\text{rounds}}$  is as follows:

$$\rho \rightarrow \left( \Phi_{\text{round}}^{\text{QR}} \right)^{N_{\text{rounds}}}(\rho). \quad (\text{B11})$$

The polarization of the target qubit is obtained as  $\alpha' = \text{Tr}(Z \rho_{\text{target}}^{\text{output}})$ , where  $\rho_{\text{target}}^{\text{output}} = \text{Tr}_{\text{target}} \left( \Phi_{\text{round}}^{\text{QR}} \right)^{N_{\text{rounds}}}(\rho)$  is

the enhanced state of the target qubit after the implementation of the PBEC-BQR.

Since the quantum refrigerator recycles the  $n - 1$  qubits remaining after extracting the target qubit and incorporates a fresh qubit from the sample, the updated input system is expressed as

$$\tilde{\rho} = \text{Tr}_{\text{target}} \left( \left( \Phi_{\text{round}}^{\text{QR}} \right)^{N_{\text{rounds}}}(\rho) \right) \otimes \rho_\alpha. \quad (\text{B12})$$

After multiple iterations of the refrigerator, the state  $\tilde{\rho}$  converges to a fixed point determined by the number of rounds,  $N_{\text{rounds}}$ , and the number of reset qubits,  $m$ . In the asymptotic cooling limit, as  $N_{\text{rounds}}$  increases, the refrigerator reaches the maximum polarization achievable by the optimal PPA Heat-Bath Algorithmic Cooling protocol, as demonstrated in the next subsection.

For the classification problem, however, it is not necessary to operate in the asymptotic cooling limit. Instead, the refrigerator will function at the fixed point corresponding to a smaller number of rounds  $\tilde{\rho}_{\text{QR}}(n, N_{\text{rounds}}, m)$ , as detailed in the section following the asymptotic discussion.

#### a. PBEC-BQR cooling limit

After each round of the PBEC-BQR, since  $m$  qubits are reset, the state of the system takes the form

$$\rho \rightarrow \rho_{\text{comp}} \otimes \rho_\alpha^{\otimes m}, \quad (\text{B13})$$

where  $\rho_{\text{comp}}$  represents the state of the qubits after removing the reset qubits. Without loss of generality, the vector corresponding to the diagonal of  $\rho_{\text{comp}}$  can be expressed as

$$\text{diag}(\rho_{\text{comp}}) = \begin{bmatrix} A_1 \\ A_2 \\ \dots \\ A_{2^{n-m}} \end{bmatrix}. \quad (\text{B14})$$

In particular, for the case  $m = 2$ , the form of the state is given as follows:

$$\text{diag}(\rho) \rightarrow \begin{bmatrix} A_1 \\ A_2 \\ \dots \\ A_{2^{n-2}} \end{bmatrix} \otimes \begin{bmatrix} p_\alpha^2 \\ p_\alpha(1-p_\alpha) \\ p_\alpha(1-p_\alpha) \\ (1-p_\alpha)^2 \end{bmatrix}, \quad (\text{B15})$$

where  $p_\alpha = (1 + \alpha)/2$ . Note that the elements of the density matrix for the reset qubits are already sorted in decreasing order when  $\alpha > 0$  or in increasing order when  $\alpha < 0$ .

The full form of Eq. (B15) is given as

$$\text{diag}(\rho) \rightarrow \begin{bmatrix} A_1 p_\alpha^2 \\ A_1 p_\alpha (1 - p_\alpha) \\ A_1 p_\alpha (1 - p_\alpha) \\ A_1 (1 - p_\alpha)^2 \\ A_2 p_\alpha^2 \\ A_2 p_\alpha (1 - p_\alpha) \\ A_2 p_\alpha (1 - p_\alpha) \\ A_2 (1 - p_\alpha)^2 \\ A_3 p_\alpha^2 \\ A_3 p_\alpha (1 - p_\alpha) \\ A_3 p_\alpha (1 - p_\alpha) \\ A_3 (1 - p_\alpha)^2 \\ \vdots \\ \vdots \\ A_{2^{n-2}} p_\alpha^2 \\ A_{2^{n-2}} p_\alpha (1 - p_\alpha) \\ A_{2^{n-2}} p_\alpha (1 - p_\alpha) \\ A_{2^{n-2}} (1 - p_\alpha)^2 \end{bmatrix}. \quad (\text{B16})$$

After a new round of the PBEC-BQR, the state is transformed in the following way:

$$\begin{bmatrix} A_1 p_\alpha^2 \\ A_1 p_\alpha (1 - p_\alpha) \\ A_1 p_\alpha (1 - p_\alpha) \\ A_1 (1 - p_\alpha)^2 \\ A_2 p_\alpha^2 \\ A_2 p_\alpha (1 - p_\alpha) \\ A_2 p_\alpha (1 - p_\alpha) \\ A_2 (1 - p_\alpha)^2 \\ A_3 p_\alpha^2 \\ A_3 p_\alpha (1 - p_\alpha) \\ A_3 p_\alpha (1 - p_\alpha) \\ A_3 (1 - p_\alpha)^2 \\ A_4 p_\alpha^2 \\ \vdots \\ \vdots \\ A_{2^{n-2-1}} (1 - p_\alpha)^2 \\ A_{2^{n-2}} p_\alpha^2 \\ A_{2^{n-2}} p_\alpha (1 - p_\alpha) \\ A_{2^{n-2}} p_\alpha (1 - p_\alpha) \\ A_{2^{n-2}} (1 - p_\alpha)^2 \end{bmatrix} \xrightarrow{\Phi_{\text{round}}^{\text{QR}}} \begin{bmatrix} A_1 p_\alpha^2 \\ A_1 p_\alpha (1 - p_\alpha) \\ A_1 p_\alpha (1 - p_\alpha) \\ A_2 p_\alpha^2 \\ A_1 (1 - p_\alpha)^2 \\ A_2 p_\alpha (1 - p_\alpha) \\ A_2 p_\alpha (1 - p_\alpha) \\ A_2 (1 - p_\alpha)^2 \\ A_3 p_\alpha^2 \\ A_2 (1 - p_\alpha)^2 \\ A_3 p_\alpha (1 - p_\alpha) \\ A_3 p_\alpha (1 - p_\alpha) \\ A_4 p_\alpha^2 \\ A_3 (1 - p_\alpha)^2 \\ \vdots \\ \vdots \\ A_{2^{n-2}} p_\alpha^2 \\ A_{2^{n-2-1}} (1 - p_\alpha)^2 \\ A_{2^{n-2}} p_\alpha (1 - p_\alpha) \\ A_{2^{n-2}} p_\alpha (1 - p_\alpha) \\ A_{2^{n-2}} (1 - p_\alpha)^2 \end{bmatrix}. \quad (\text{B17})$$

The elements in the diagonal that are permuted are the following:

$$A_{(i+1)} p_\alpha^2 \leftrightarrow A_i (1 - p_\alpha)^2, \quad \text{for } i = 1, 2, \dots, 2^{n-2} - 1,$$

as shown in color in the transformation.

In the asymptotic limit, as the number of rounds increases, the state of the elements converges to one that is invariant under the  $\Phi_{\text{round}}^{\text{QR}}$  operation. Consequently, in this asymptotic limit, the following condition must be satisfied:

$$A_{(i+1)} p_\alpha^2 = A_i (1 - p_\alpha)^2, \quad \text{for } i = 1, 2, \dots, 2^{n-2} - 1. \quad (\text{B18})$$

This condition has the form of the one given by the conventional HBAC using 2 reset qubits, see Eq. (S29) of the supplemental material in Ref. [25]. From here, asymptotic polarization for the target qubit as the  $N_{\text{rounds}}$  grows is the maximum achievable for the HBAC with  $m = 2$ ,

$$\begin{aligned} \alpha_{\text{PBEC}}^\infty(n, m = 2) &= \frac{(1 + \alpha)^{2^{n-2}} - (1 - \alpha)^{2^{n-2}}}{(1 + \alpha)^{2^{n-2}} + (1 - \alpha)^{2^{n-2}}} \\ &= \tanh [2^{n-2} \text{arctanh}(\alpha)]. \end{aligned}$$

Similarly, it is now straightforward to see that by following the same derivations, for any  $m > 2$ , the condition that holds in the asymptotic cooling limit is given by

$$A_{(i+1)} p_\alpha^m = A_i (1 - p_\alpha)^m, \quad \text{for } i = 1, 2, \dots, 2^{n-m} - 1, \quad (\text{B19})$$

which gives the general condition for the asymptotic cooling limit, see Eq. (S29) of the supplemental material in Ref. [25],

$$\alpha_{\text{PBEC}}^\infty = \alpha_{\text{HBAC}}^\infty = \frac{(1 + \alpha)^{m 2^{n-m-1}} - (1 - \alpha)^{m 2^{n-m-1}}}{(1 + \alpha)^{m 2^{n-m-1}} + (1 - \alpha)^{m 2^{n-m-1}}} \quad (\text{B20})$$

$$= \tanh [m 2^{n-m-1} \text{arctanh}(\alpha)]. \quad (\text{B21})$$

### b. Enhanced polarization of the PBEC-BQR for finite $N_{\text{rounds}}$

The enhanced polarization of the target qubit,  $\alpha_{\text{PBEC}}$ , depends on the refrigerator configuration parameters ( $n$ ,  $N_{\text{rounds}}$ , and  $m$ ), which determine the state of the BQR after  $N_{\text{rounds}}$ , denoted  $\tilde{\rho}_{\text{QR}}(n, N_{\text{rounds}}, \alpha)$ . In this subsection, we present the general evolution of the PBEC-BQR under  $N_{\text{rounds}}$ , the corresponding  $N_{\text{rounds}}$ -dependent state of the refrigerator, and the resulting enhanced polarization for an arbitrary configuration. We then illustrate explicit examples of the enhancement for configurations with  $n = 4$  and  $n = 5$  qubits, both using  $m = 2$  reset qubits, showing the behavior as a function of the number of rounds.

In general, the state of the  $n$  system immediately after the  $j$ th round, takes the following form:

$$\text{diag}(\rho^{(j)}) = \begin{bmatrix} A_1^{(j)} \\ A_2^{(j)} \\ \dots \\ A_{2^{n-m}}^{(j)} \end{bmatrix} \otimes \text{diag}(\rho_\alpha^m). \quad (\text{B22})$$

Let us denote by  $\vec{A}^{(j)}$  the diagonal vector of the state of the first  $n - m$  qubits. By implementing one round of the BQR, following the transformation given in Eq. (B10), the diagonal vector is transformed using a stochastic matrix  $M$  as follows:

$$\vec{A}^{(j)} \xrightarrow{\Phi_{\text{round}}^{\text{QR}}} \vec{A}^{(j+1)} = M \cdot \vec{A}^{(j)}. \quad (\text{B23})$$

After  $N_{\text{rounds}}$ , that diagonal vector is given by

$$\vec{A}^{N_{\text{rounds}}} = M^{N_{\text{rounds}}} \cdot \vec{A}^{(0)}, \quad (\text{B24})$$

where  $\vec{A}^{(0)}$  represents to the initial state when the refrigerator begins operation. Since this state is recycled from the previous cycle of the refrigerator, and  $M^{N_{\text{round}}}$  has attractive fixed points, the refrigerator will eventually operate in a steady state. In this regime, the input state to the refrigerator matches the output state after removing the enhanced target qubit. This condition implies that the following equation must be solved to determine the state  $\tilde{\rho}_{\text{QR}}$ :

$$\tilde{\rho}_{\text{QR}} = \text{Tr}_{\text{target}} \left[ \left( \Phi_{\text{round}}^{\text{QR}} \right)^{N_{\text{rounds}}} (\tilde{\rho}_{\text{QR}} \otimes \rho_\alpha) \right]. \quad (\text{B25})$$

Expressed in terms of the diagonal vector of the fixed point,  $\vec{A}_{\text{QR}}$ , this condition becomes

$$\vec{A}_{\text{QR}} = \text{Tr}_{\text{target}} \left[ M^{N_{\text{rounds}}} \cdot \vec{A}_{\text{QR}} \right] \otimes \text{diag}(\rho_\alpha), \quad (\text{B26})$$

which must be solved for a given  $N_{\text{rounds}}$ . The updated polarization  $\alpha'$  for the target qubit, after the  $N_{\text{rounds}}$  is given by

$$\alpha' = \text{Tr} [Z \rho_{\text{target}}], \quad (\text{B27})$$

$$\text{where } \rho_{\text{target}} = \text{Tr}_{\text{target}} \left[ M^{N_{\text{rounds}}} \cdot \vec{A}_{\text{QR}} \right]. \quad (\text{B28})$$

(a) Example for  $n = 4$ , with  $m = 2$  for  $N_{\text{rounds}} = 3$ :

The state of the  $n = 4$  system immediately after a round that resets  $m = 2$  qubits takes the form

$$\text{diag}(\rho_4^{(j)}) = \begin{bmatrix} A_1^{(j)} \\ A_2^{(j)} \\ A_3^{(j)} \\ A_4^{(j)} \end{bmatrix} \otimes \text{diag}(\rho_\alpha^2). \quad (\text{B29})$$

By implementing one round of the BQR, following the transformation given in Eq. (B10), the stochastic matrix  $M$  that describes that transformation is

$$\vec{A}^{(j)} \xrightarrow{\Phi_{\text{round}}^{\text{QR}}} \vec{A}^{(j+1)} = M \cdot \vec{A}^{(j)}, \quad (\text{B30})$$

where

$$M_4 = \begin{bmatrix} p_\alpha(2-p_\alpha) & p_\alpha^2 & 0 & 0 \\ (1-p_\alpha)^2 & 2p_\alpha(1-p_\alpha) & p_\alpha^2 & 0 \\ 0 & (1-p_\alpha)^2 & 2p_\alpha(1-p_\alpha) & p_\alpha^2 \\ 0 & 0 & (1-p_\alpha)^2 & 1-p_\alpha^2 \end{bmatrix}.$$

This is derived as in Eq. (B17), followed by tracing out the two reset qubits to obtain the updated diagonal vector  $\vec{A}^{(j+1)}$ .

Note that for  $n = 4$  and  $m = 2$ , the vector  $\vec{A}_{\text{QR}}$  is given in the form  $\vec{A}_{\text{QR}} = \{p_{\text{QR}}, 1 - p_{\text{QR}}\} \otimes \{p_\alpha, 1 - p_\alpha\}$ , since two qubits are reset and the target qubit is removed. Thus, we only need to solve the parameter  $p_{\text{QR}}$ . Then, the condition given in Eq. (B26) to obtain the state of the refrigerator, for this case when the refrigerator is operating with  $N_{\text{rounds}} = 3$  is expressed as

$$p_{\text{QR}} = p_\alpha^2 \cdot \frac{1 + 2p_\alpha + 16p_\alpha^2 - 70p_\alpha^3 + 86p_\alpha^4 - 34p_\alpha^5}{1 - 20p_\alpha^2 + 106p_\alpha^3 - 219p_\alpha^4 + 198p_\alpha^5 - 66p_\alpha^6}.$$

Substituting this result in Eq. (B28) with the updated  $\vec{A}_{\text{QR}}$ , the enhanced polarization of the target qubits is given by

$$\frac{\alpha'}{\alpha} = \frac{943 - 1608\alpha^2 + 2711\alpha^4 - 2248\alpha^6 + 925\alpha^8 - 240\alpha^{10} + 29\alpha^{12}}{16(25 - 19\alpha^2 + 59\alpha^4 - 33\alpha^6)}.$$

Note that the plots in the main manuscript were obtained through a direct simulation of the circuit for the bidirectional refrigerator.

(b) Example for  $n = 5$  with  $m = 2$ :

$$M_5 = \begin{bmatrix} p_\alpha(2-p_\alpha) & p_\alpha^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ (1-p_\alpha)^2 & 2p_\alpha(1-p_\alpha) & p_\alpha^2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & (1-p_\alpha)^2 & 2p_\alpha(1-p_\alpha) & p_\alpha^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & (1-p_\alpha)^2 & 2p_\alpha(1-p_\alpha) & p_\alpha^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & (1-p_\alpha)^2 & 2p_\alpha(1-p_\alpha) & p_\alpha^2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & (1-p_\alpha)^2 & 2p_\alpha(1-p_\alpha) & p_\alpha^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & (1-p_\alpha)^2 & 2p_\alpha(1-p_\alpha) & p_\alpha^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & (1-p_\alpha)^2 & 1-p_\alpha^2 \end{bmatrix}, \quad (\text{B31})$$

and similarly for  $n > 5$ , with  $m = 2$ .

As a concrete example, solving for the case of  $N_{\text{rounds}} = 3$  with the refrigerator running in the steady state yields target qubits with the following enhanced polarization:  $\alpha' = \alpha(1444856 - 3513097\alpha^2 + 7571347\alpha^4 - 7884719\alpha^6 + 5496725\alpha^8 - 1352250\alpha^{10} + 681150\alpha^{12} - 838878\alpha^{14} + 272130\alpha^{16} - 39085\alpha^{18} - 4929\alpha^{20} + 1661\alpha^{22} + 97\alpha^{24}) / (32(21456 - 47629\alpha^2 + 134063\alpha^4 - 129925\alpha^6 + 76115\alpha^8 + 11917\alpha^{10} + 5277\alpha^{12} - 14023\alpha^{14} + 81\alpha^{16} + 12\alpha^{18}))$ . This expression provides the exact analytical form of the red line in Fig. 5, which corresponds to  $\alpha_{\text{QR}}$  for  $N_{\text{rounds}} = 3$ . Substituting this result into Eq. (20) generates the corresponding line in Fig. 6. The figures were reproduced numerically by simulating the circuit. These sections are intended to illustrate exact analytical results.

## 2. Evolution under the $k$ -local BQR

For the  $k$ -local BQR, the evolution retains the same structure as in Eq. (B10), but with a different unitary step. In the example considered in the main text, where  $k = 3$ ,

$$M_{4(3\text{-local})} = \begin{bmatrix} p_\alpha(2-p_\alpha) & p_\alpha^2 & 0 & 0 \\ (1-p_\alpha)^2 & p_\alpha(1-p_\alpha) & p_\alpha & 0 \\ 0 & 1-p_\alpha & p_\alpha(1-p_\alpha) & p_\alpha^2 \\ 0 & 0 & (1-p_\alpha)^2 & 1-p_\alpha^2 \end{bmatrix}.$$

In the asymptotic limit, as the number of rounds increases, the vector  $\vec{A}_\infty$  is invariant under the effect of one

The diagonal state of the  $n = 5$  system immediately after a round that resets  $m = 2$  qubits takes the form  $\text{diag}(\rho^{(j)}) = A_1^{(j)}, A_2^{(j)}, \dots, A_8^{(j)} \otimes \text{diag}(\rho_\alpha)$ . The corresponding stochastic matrix that transforms the diagonal vector  $\vec{A}$ , following the effect of a round  $\Phi_{\text{round}}^{\text{QR}}(\cdot)$ , is given by

the unitary applied in each round is given by

$$U_{\text{QR}3\text{local}} = (U_{C_3} \otimes \mathbb{1}_{2^{n-3}})(\mathbb{1}_2 \otimes U_{C_3} \otimes \mathbb{1}_{2^{n-4}}) \cdots (\mathbb{1}_{2^{n-3}} \otimes U_{C_3}).$$

Thus, the effect of a single round on a system of  $n$  qubits in the state  $\rho$  is given by

$$\Phi_{\text{round}}^{\text{QR}(k=3)}(\rho) := \text{Tr}_m \left[ U_{\text{QR}3\text{local}} \rho U_{\text{QR}3\text{local}}^\dagger \right] \otimes \rho_\alpha^{\otimes m}. \quad (\text{B32})$$

An analogous procedure determines the steady state of the refrigerator and the corresponding enhanced polarization in the  $k$ -local setting. In this case, the steady state is generated by iterating the  $k$ -local round unitary  $U_{\text{QR}3\text{local}}$ , replacing the global unitary  $U_{\text{QR}}$  used in the PBEC-BQR.

### a. Three-local BQR cooling limit

(c) Asymptotic polarization for  $n = 4$  and  $m = 2$ :

The stochastic matrix that describes the evolution in the form of Eq. (B23) is given by

more round,

$$\vec{A}_\infty = M_{4(3\text{-local})} \cdot \vec{A}_\infty. \quad (\text{B33})$$

Solving this condition, we found that, in the asymptotic limit, the  $n$  qubits are in a product state, with each qubit having a different ground-state population. Considering the qubits in order, starting from the end of the string and moving toward the target qubit, the populations are given by

$$\vec{p}^\infty(n=4) = \left\{ p_\alpha, p_\alpha, \frac{p_\alpha^2}{1-2p_\alpha+2p_\alpha^2}, \frac{p_\alpha^3}{1-3p_\alpha+3p_\alpha^2} \right\}.$$

The two qubits at the end of the string retain their initial populations,  $p_\alpha$ , as they are reset after each round. In

contrast, the other qubits in the string exhibit improved populations.

As the size of the string increases, the steady state follows a specific pattern, which will be described in the following subsections.

(d) Asymptotic polarization for  $n = 5$  and  $m = 2$ :

The stochastic matrix that describes the evolution in the form of Eq. (B23) is given by

$$M_{5(3\text{-local})} = \begin{bmatrix} p_\alpha(2-p_\alpha) & p_\alpha^2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ (1-p_\alpha)^2 & p_\alpha(1-p_\alpha) & p_\alpha & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1-p_\alpha & p_\alpha(1-p_\alpha) & p_\alpha^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & p_\alpha(2-p_\alpha) & p_\alpha^2 & 0 & 0 & 0 \\ 0 & 0 & (1-p_\alpha)^2 & 1-p_\alpha^2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & (1-p_\alpha)^2 & p_\alpha(1-p_\alpha) & p_\alpha & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1-p_\alpha & p_\alpha(1-p_\alpha) & p_\alpha^2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1-p_\alpha & p_\alpha(1-p_\alpha) & p_\alpha^2 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & (1-p_\alpha)^2 & 1-p_\alpha^2 \end{bmatrix}.$$

Solving the condition  $\vec{A}_\infty = M_{5(3\text{-local})} \cdot \vec{A}_\infty$ , the corresponding vector of enhanced ground-state populations of each qubit is given as

$$\vec{p}_\infty(n) = \left\{ p_\alpha, p_\alpha, \frac{p_\alpha^2}{1-2p_\alpha+2p_\alpha^2}, \frac{p_\alpha^3}{1-3p_\alpha+3p_\alpha^2}, \frac{p_\alpha^5}{1-5p_\alpha+10p_\alpha^2-10p_\alpha^3+5p_\alpha^5} \right\}.$$

(e) Asymptotic polarization for  $n$  and  $m = 2$ :

By iterating the described derivations, in the asymptotic limit of the three-local BQR will be in a product state, with the following list of the local ground-state populations for the qubits, counting from the end of the string of qubits to the target qubit

$$\vec{p}_\infty(n) = \left\{ p_\alpha, p_\alpha, \frac{p_\alpha^2}{1-2p_\alpha+2p_\alpha^2}, \frac{p_\alpha^{F_4}}{\sum_{i=0}^{F_4-1} (-1)^i \binom{F_4}{i} p_\alpha^i}, \frac{p_\alpha^{F_5}}{\sum_{i=0}^{F_5-1} (-1)^i \binom{F_5}{i} p_\alpha^i}, \dots, \frac{p_\alpha^{F_n}}{\sum_{i=0}^{F_n-1} (-1)^i \binom{F_n}{i} p_\alpha^i} \right\}$$

for  $n \geq 3$ ,

where  $F_j$  is the  $j$ th Fibonacci number (i.e.,  $F_j = F_{j-1} + F_{j-2}$  with  $F_1 = 1, F_2 = 1$ ).

However, as mentioned in the main text, it is not necessary for the quantum refrigerator to operate near the asymptotic cooling limit. Instead, the optimal configuration requires only a small number of rounds, striking a balance between significant polarization enhancement and efficient use of qubit resources to maximize the reduction in sampling error, as illustrated in Fig. 8.

### APPENDIX C: ESTIMATING THE GRADIENT OF THE CLASSIFICATION SCORE

Writing the classification score as  $q(x, \theta) = \langle x | \exp(i\theta G) M \exp(-i\theta G) | x \rangle$  with a Hermitian operator  $G$ , the gradient can be expressed as

$$\frac{\partial q(x, \theta)}{\partial \theta} = \langle x | \exp(i\theta G) i[G, M] \exp(-i\theta G) | x \rangle. \quad (C1)$$

Since  $i[G, M]$  is Hermitian, the above equation represents measuring the expectation value of  $\tilde{M} := i[G, M]$  on the same VQC that is used to compute  $q(x, \theta)$ .

In our binary classification setting,  $M$  is assumed to be a Pauli observable. Therefore, if  $G$  is a Pauli operator, then  $\tilde{M}$  is also a Pauli operator. In this case, the Clifford transformation technique allows us to compute the gradient as

$$\frac{\partial q(x, \theta)}{\partial \theta} = \text{Tr}(Z\sigma_1(x, \theta)), \quad (C2)$$

where  $\sigma_1(x, \theta) = \text{Tr}_{n-1} \left( V_c \exp(-i\theta G) |x\rangle \langle x| \exp(i\theta G) V_c^\dagger \right)$  and  $V_c$  represents a Clifford gate. This expression is analogous to Eq. (1) of the main text, except for the difference between  $U_c$  and  $V_c$ , as  $U_c$  transforms  $M$  to  $Z_1$  and  $V_c$  transforms  $\tilde{M}$  to  $Z_1$ . Thus, estimating the gradient with an additive error less than its magnitude results in a similar error probability bound as described in Eq. (6) of the main text.

#### APPENDIX D: KERNEL-BASED QUANTUM BINARY CLASSIFIERS

Although the main text focuses primarily on variational quantum binary classifiers (VQBCs), the transformation in Eq. (10) can be applied more broadly to any binary classifier established based on the classification score in Eq. (1). For instance, quantum kernelized binary classifiers (QKBC) utilizing Hadamard test or swap test circuits [56–62] also compute the classification score as expressed in Eq. (1). In such instances, the quantum circuits are designed so that the single-qubit Pauli-Z measurement yields

$$\langle Z \rangle = \frac{1}{s} \sum_{i=1}^s y_i k(x_i, \tilde{x}), \quad (\text{D1})$$

where  $k(x_i, \tilde{x})$  represents the kernel function quantifying the similarity between a sample data point  $x_i$  and the test data  $\tilde{x}$ ,  $y_i \in \{+1, -1\}$  denotes the class label for  $x_i$ , and  $s$  is the total number of samples. The QKBC adheres to the same classification rule as described in Eq. (3) for the VQBC. Another notable example is the quantum support vector machine proposed in Ref. [63]. Therefore, the importance of developing an efficient technique to achieve Eq. (10) extends beyond VQBC, encompassing broader applications in QML.

#### APPENDIX E: NUMERICAL ANALYSIS OF NOISE EFFECTS

To assess the robustness of the quantum refrigerator protocols under experimentally realistic conditions, we performed numerical simulations that incorporate the dominant noise mechanisms observed in current NISQ devices. We modeled the effects of energy relaxation and control errors through generalized amplitude damping (GAD) and depolarizing channels, respectively. These channels capture the primary sources of decoherence and gate imperfections present across all major NISQ hardware platforms, including superconducting [64–67], trapped-ion [68,69], and neutral-atom [70–72] systems.

It is worth emphasizing that the quantum refrigerator operates entirely within the subspace of diagonal states: all operations in the protocol—resets, permutations, and conditional swaps—preserve diagonality at every step. As a

consequence, the protocol is inherently insensitive to any noise processes that act only on quantum coherences, such as dephasing, phase-damping, or control-phase fluctuations [15,73,74]. These noise channels, while detrimental in general quantum algorithms, leave the populations of diagonal states unaffected and therefore do not influence the outcome of the refrigerator. In this sense, the protocol is intrinsically robust against coherence-degrading noise, and our numerical simulations incorporate this property by focusing on the population-changing channels (GAD and depolarizing) that directly affect the diagonal elements of the density matrix.

##### a. Noise model

The generalized amplitude-damping channel describes energy relaxation to a thermal bath with equilibrium excitation probability  $p_{\text{GAD}}$  and damping probability  $\gamma = 1 - e^{-\tau/T_1}$ , where  $T_1$  is the longitudinal relaxation time and  $\tau$  the gate duration. The corresponding Kraus operators are as follows:

$$\begin{aligned} E_0 &= \sqrt{p_{\text{GAD}}} \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{pmatrix}, \quad E_1 = \sqrt{p_{\text{GAD}}} \begin{pmatrix} 0 & \sqrt{\gamma} \\ 0 & 0 \end{pmatrix}, \\ E_2 &= \sqrt{1-p_{\text{GAD}}} \begin{pmatrix} \sqrt{1-\gamma} & 0 \\ 0 & 1 \end{pmatrix}, \\ E_3 &= \sqrt{1-p_{\text{GAD}}} \begin{pmatrix} 0 & 0 \\ \sqrt{\gamma} & 0 \end{pmatrix}. \end{aligned} \quad (\text{E1})$$

This channel reduces to standard amplitude damping when  $p_{\text{GAD}} = 0$ , and in diagonal state it affects the populations of the computational basis states without introducing coherences.

The depolarizing channel models stochastic gate errors that randomize the qubit state with probability  $p_{\text{depol}}$ ,

$$\mathcal{D}_{p_{\text{depol}}}(\rho) = (1 - p_{\text{depol}})\rho + \frac{p_{\text{depol}}}{d} \mathbb{I}, \quad (\text{E2})$$

where  $d$  is the Hilbert-space dimension of the affected subsystem ( $d = 2$  for single-qubit gates and  $d = 4$  for two-qubit gates). Since the working states of the refrigerator remain diagonal, dephasing-type errors have no effect beyond these channels.

##### b. Typical and worst-case noise regimes

We considered two representative noise regimes. First, the *typical NISQ regime* corresponds to a set of parameters chosen to lie within the ranges reported for present-day superconducting, trapped-ion, and neutral-atom devices [64–72]. Note that, this regime does not reflect the calibration of any particular hardware platform, but rather encapsulates the characteristic magnitude of gate and relaxation

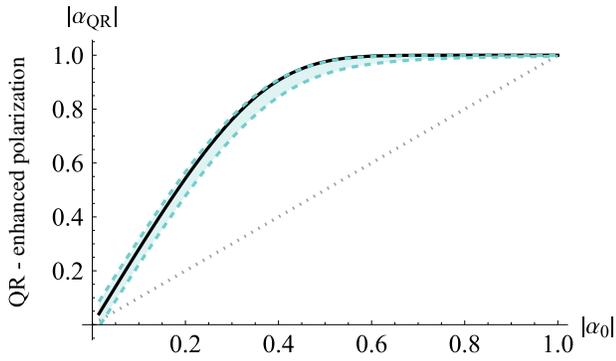


FIG. 10. Enhanced polarization  $\alpha_{QR}$  of the quantum refrigerator, with the progressive boundary entropy compression-bidirectional quantum refrigerator protocol operating with  $n = 5$  qubits and  $N_{\text{rounds}} = 5$  as a function of the initial polarization  $|\alpha_0|$ , under typical noise parameters representative of current noisy intermediate-scale quantum devices. The dotted straight line indicates the baseline initial polarization, while the solid black curve corresponds to the ideal noise-free enhancement. The green zone shows the expected polarization in the presence of typical noise channels—generalized amplitude damping and depolarizing noise.

errors that are broadly representative of current NISQ technologies. By avoiding platform-specific assumptions, these values define a common benchmark that realistically captures the overall performance envelope of existing quantum processors. For the *typical NISQ regime*, the noise parameters are  $p_{\text{GAD}} = 10^{-5}$ ,  $\gamma_{1q} = 5 \times 10^{-4}$ ,  $\gamma_{2q} = 2.5 \times 10^{-3}$ ,  $p_{\text{depol}}^{(1q)} = 10^{-4}$ , and  $p_{\text{depol}}^{(2q)} = 5 \times 10^{-3}$ . The *worst-case regime* extends beyond the upper end of the noise range currently reported for NISQ hardware and serves as a conservative stress test for the protocol's robustness. It employs larger noise strengths to probe the limits of stability and performance:  $p_{\text{GAD}} = 0.9$ ,  $\gamma_{1q} = 2.7 \times 10^{-3}$ ,  $\gamma_{2q} = 1.3 \times 10^{-2}$ ,  $p_{\text{depol}}^{(1q)} = 10^{-3}$ , and  $p_{\text{depol}}^{(2q)} = 2 \times 10^{-2}$ .

Simulations were performed for the five-qubit quantum refrigerator configuration, with the PBEC-BQR protocol operating with two reset qubits per round and  $N_{\text{rounds}} = 5$ . Figures 10 and 11 summarize the resulting performance under the noise conditions defined above. Figure 10 shows the enhancement of the target-qubit polarization relative to the initial polarization baseline (dotted line) and to the ideal noise-free case (solid black line). In the presence of moderate noise (green curve), the achievable polarization is slightly reduced, yet the system achieves a steady state. Figure 11 presents the corresponding reduction factor of the error-probability bound, which quantifies the improvement in the effective error probability achieved through the cooling process.

Across the full range of initial polarizations, and under typical NISQ-level noise, the system reliably converges toward a steady state that remains close to the ideal fixed

point. This robustness arises from two complementary features of the protocol. First, the dynamics of the quantum refrigerator exhibit attractive fixed points that guarantee convergence toward a steady state upon iteration. Second, the diagonal structure of the protocol enhances stability: since all operations—resets, swaps, and basis permutations—preserve diagonality, the population dynamics are inherently resilient to small stochastic fluctuations in the reset and permutation steps. Consequently, noise processes that act solely on quantum coherences, such as dephasing, phase-damping, or control-phase fluctuations, have no measurable influence on the outcome, as they leave diagonal states invariant throughout the evolution.

Even under typical NISQ-level noise, the cooling process continues to reduce the effective error probability over nearly the same range of initial polarizations as observed in the ideal, noise-free case. Note that, the generalized amplitude-damping channel introduces an asymmetry in the estimation-error bounds for positive and negative initial polarizations. Because amplitude damping drives the population toward one polarization direction, the resulting magnitude of  $|\alpha'|$  in Eq. (6) differs for the two labels. This effect is visible in Fig. 11, where the upper dotted green curve corresponds to the case less affected by damping, and the lower dotted curve represents the configuration in which  $|\alpha'|$  is more strongly degraded by noise. The overall reduction factor of the error-probability bound is therefore computed as

$$r_{QR} = \frac{2r_{QR}^+ r_{QR}^-}{r_{QR}^+ + r_{QR}^-},$$

where  $r_{QR}^+$  and  $r_{QR}^-$  denote the reduction factors for states with positive and negative polarization, respectively. In Fig. 11, the solid green line represents the global  $r_{QR}$

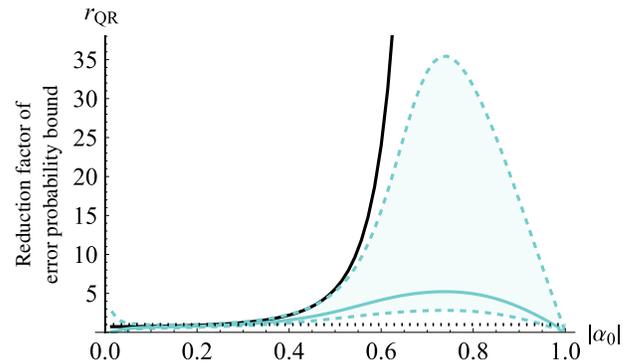


FIG. 11. Reduction factor of the error-probability bound obtained for the same protocol and noise parameters. The reduction factor quantifies the improvement in the effective error probability achieved through the cooling process, illustrating the robustness of the protocol in the typical noisy intermediate-scale quantum regime.

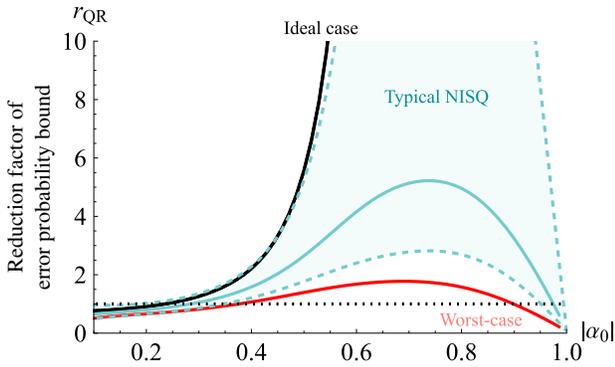


FIG. 12. Reduction factor of the error-probability bound for the five-qubit progressive boundary entropy compression-bidirectional quantum refrigerator with  $N_{\text{rounds}} = 5$ , evaluated under both the typical noisy intermediate-scale quantum noise regime (green curve) and the worst-case regime (solid red line). Even under the deliberately exaggerated noise strengths of the worst-case scenario, the protocol continues to yield an enhancement over nearly the same range of initial polarizations as in the ideal case, demonstrating its robustness against realistic and strongly dissipative noise conditions.

obtained for the case where both polarization labels occur with equal probability in the classification task.

To further assess the robustness of the protocol, we simulated the same five-qubit configuration under the worst-case noise regime defined above. Figure 12 compares the reduction factor of the error-probability bound for the typical NISQ parameters (green curve) with that obtained under the worst-case noise conditions (solid red line). Remarkably, even in this highly dissipative regime—corresponding to relaxation times on the order of tens of microseconds and two-qubit gate errors approaching a few percent—the protocol continues to yield a reduction in the effective error probability over nearly the same range of initial polarizations as in the ideal case. The results approach the boundary beyond which no enhancement is obtained, yet the persistence of a reduction factor larger than unity across most of the region demonstrates that the quantum refrigerator retains its operational advantage and steady-state convergence well beyond typical NISQ noise conditions, underscoring its intrinsic robustness.

## APPENDIX F: DETAILS OF CLASSIFICATION EXPERIMENTS

This appendix provides additional details on the datasets, preprocessing procedures, and simulation setup used for evaluating the classification performance of BQR-enhanced quantum classifiers.

All classification simulations were implemented in Python using Qiskit [75] and scikit-learn [76]. We used a

combination of synthetic and real-world binary classification datasets, as listed below.

- (a) Synthetic datasets: Two synthetic datasets were created using polarization values  $\alpha \in [-1, 1]$ . For the uniform case, positive and negative class values were sampled from  $(0, 1]$  and  $[-1, 0)$ , respectively. For the Gaussian case, we used  $\mathcal{N}(+0.6, 0.3^2)$  and  $\mathcal{N}(-0.6, 0.3^2)$  distributions and clipped the samples to lie within  $[-1, 1]$ .
- (b) Iris: We used the standard Iris dataset [77] and formed two binary classification tasks: (i) setosa vs versicolor and (ii) setosa vs virginica. For each task, we applied principal component analysis (PCA) to project the original four-dimensional features to one dimension, and rescaled the values to  $[-1, 1]$ .
- (c) Wine: We selected class 0 (Barolo) and class 1 (Grignolino) from the Wine dataset [78]. PCA was used for dimensionality reduction, and the resulting values were rescaled to  $[-1, 1]$ .
- (d) Handwritten digits: We used the digits dataset [79] provided by scikit-learn [76], which contains  $8 \times 8$  grayscale images of handwritten digits from 0 to 9. For binary classification, we selected samples labeled as 2 and 5. PCA was applied to project the 64-dimensional image vectors to one dimension, followed by rescaling to the  $[-1, 1]$  range.
- (e) Sonar: The UCI Sonar dataset contains sonar returns labeled as “Rock” or “Mine” [80]. We encoded these labels as 0 and 1, applied LDA to reduce the dimensionality to 1, and scaled the values to  $[-1, 1]$ .
- (f) Diabetes: We used the Pima Indians Diabetes dataset [81], which contains diagnostic measurements for predicting whether a patient has diabetes. Linear discriminant analysis (LDA) was used to reduce the original eight features to one dimension, followed by rescaling to  $[-1, 1]$ .

For each dataset, we randomly sampled 50 data points from each class to form a balanced binary classification task. This process was repeated 100 times to generate statistically independent test scenarios. Classification was based on the sign of the sample mean from the computational basis (i.e.,  $Z$ -basis) measurements on the target qubit.

We used the three-local BQR with five system qubits ( $n = 5$ ), two reset qubits ( $m = 2$ ), and two rounds ( $N_{\text{rounds}} = 2$ ). The number of measurement shots was set to  $k_{\text{BQR}} \in \{10, 100\}$ . For fair comparison, the conventional baseline was allocated proportionally more shots using the expression  $k_c = k_{\text{BQR}} \times m \times (N_{\text{rounds}} - 1) + n$ , which corresponds to 25 and 205 shots for  $k_{\text{BQR}} = 10$  and 100, respectively.

The classification accuracy was averaged over 100 randomized trials, and Welch’s  $t$ -test was used to assess

statistical significance between the BQR-enhanced classifier and the baseline classifier without BQR. The full implementation code is publicly available at [82].

- 
- [1] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd, Quantum machine learning, *Nature* **549**, 195 (2017).
- [2] Vojtech Havlicek, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature* **567**, 209 (2019).
- [3] Yunchao Liu, Srinivasan Arunachalam, and Kristan Temme, A rigorous and robust quantum speed-up in supervised machine learning, *Nat. Phys.* **17**, 1013 (2021).
- [4] Jennifer R. Glick, Tanvi P. Gujarati, Antonio D. Córcoles, Youngseok Kim, Abhinav Kandala, Jay M. Gambetta, and Kristan Temme, Covariant quantum kernels for data with group structure, *Nat. Phys.* **20**, 479 (2024).
- [5] Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, and Jarrod R. McClean, Quantum advantage in learning from experiments, *Science* **376**, 1182 (2022).
- [6] M. Cerezo, Guillaume Verdon, Hsin-Yuan Huang, Lukasz Cincio, and Patrick J. Coles, Challenges and opportunities in quantum machine learning, *Nat. Comput. Sci.* **2**, 567 (2022).
- [7] Edward Farhi and Hartmut Neven, Classification with quantum neural networks on near term processors, [arXiv:1802.06002](https://arxiv.org/abs/1802.06002).
- [8] Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [9] Iris Cong, Soonwon Choi, and Mikhail D. Lukin, Quantum convolutional neural networks, *Nat. Phys.* **15**, 1273 (2019).
- [10] M. Cerezo, Andrew Arrasmith, Ryan Babbush, Simon C. Benjamin, Suguru Endo, Keisuke Fujii, Jarrod R. McClean, Kosuke Mitarai, Xiao Yuan, Lukasz Cincio, and Patrick J. Coles, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [11] Gilles Brassard, Peter Hoyer, Michele Mosca, and Alain Tapp, Quantum amplitude amplification and estimation, *Contemp. Math.* **305**, 53 (2002).
- [12] Ashley Montanaro, Quantum speedup of Monte Carlo methods, *Proc. R. Soc. A: Math. Phys. Eng. Sci.* **471**, 20150301 (2015).
- [13] Yohichi Suzuki, Shumpei Uno, Rudy Raymond, Tomoki Tanaka, Tamiya Onodera, and Naoki Yamamoto, Amplitude estimation without phase estimation, *Quantum Inf. Process.* **19**, 1 (2020).
- [14] Dmitry Grinko, Julien Gacon, Christa Zoufal, and Stefan Woerner, Iterative quantum amplitude estimation, *npj Quantum Inf.* **7**, 52 (2021).
- [15] John Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [16] Kishor Bharti, Alba Cervera-Lierta, Thi Ha Kyaw, Tobias Haug, Sumner Alperin-Lea, Abhinav Anand, Matthias Degroote, Hermanni Heimonen, Jakob S. Kottmann, Tim Menke, Wai-Keong Mok, Sukin Sim, Leong-Chuan Kwek, and Alán Aspuru-Guzik, Noisy intermediate-scale quantum algorithms, *Rev. Mod. Phys.* **94**, 015004 (2022).
- [17] P. Oscar Boykin, Tal Mor, Vwani Roychowdhury, Farrokh Vatan, and Rutger Vrijen, Algorithmic cooling and scalable NMR quantum computers, *Proc. Natl. Acad. Sci. U.S.A.* **99**, 3388 (2002).
- [18] Leonard J. Schulman and Umesh V. Vazirani, *Molecular scale heat engines and scalable quantum computation, in Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, STOC '99* (Association for Computing Machinery, New York, NY, USA, 1999), pp. 322–329.
- [19] Daniel K. Park, Nayeli A. Rodríguez-Briones, Guanru Feng, Robabeh Rahimi, Jonathan Baugh, and Raymond Laflamme, *Heat Bath Algorithmic Cooling with Spins: Review and Prospects* (Springer, New York, 2016), p. 227.
- [20] Yuval Elias, Tal Mor, and Yossi Weinstein, Semioptimal practicable algorithmic cooling, *Phys. Rev. A* **83**, 042340 (2011).
- [21] Daniel K. Park, Guanru Feng, Robabeh Rahimi, StÉphane Labruyère, Taiki Shibata, Shigeaki Nakazawa, Kazunobu Sato, Takeji Takui, Raymond Laflamme, and Jonathan Baugh, Hyperfine spin qubits in irradiated malonic acid: Heat-bath algorithmic cooling, *Quantum Inf. Process.* **14**, 2435 (2015).
- [22] Nayeli Azucena Rodríguez-Briones, Novel heat-bath algorithmic cooling methods, Ph.D. thesis, University of Waterloo, 2020.
- [23] Junan Lin, Nayeli A. Rodríguez-Briones, Eduardo Martín-Martínez, and Raymond Laflamme, Thermodynamic analysis of algorithmic cooling protocols: Efficiency metrics and improved designs, *Phys. Rev. A* **110**, 022215 (2024).
- [24] Ralph Silva, Pharnam Bakhshinezhad, and Fabien Clivaz, Optimal unitary trajectories under commuting target and cost observables; applications to cooling, [arXiv:2412.07291](https://arxiv.org/abs/2412.07291).
- [25] Nayeli Azucena Rodríguez-Briones and Raymond Laflamme, Achievable polarization for heat-bath algorithmic cooling, *Phys. Rev. Lett.* **116**, 170501 (2016).
- [26] Lindsay Bassman Otelie, Antonella De Pasquale, and Michele Campisi, Dynamic cooling on contemporary quantum computers, *PRX Quantum* **5**, 030309 (2024).
- [27] Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Information Science and Statistics (Springer, New York, 2006).
- [28] Shai Shalev-Shwartz and Shai Ben-David, *Understanding Machine Learning: From Theory to Algorithms* (Cambridge University Press, Cambridge, 2014).
- [29] Roberto Giuntini, Federico Holik, Daniel K. Park, Hector Freytes, Carsten Blank, and Giuseppe Sergioli, Quantum-inspired algorithm for direct multi-class classification, *Appl. Soft Comput.* **134**, 109956 (2023).
- [30] Zoë Holmes, Kunal Sharma, M. Cerezo, and Patrick J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, *PRX Quantum* **3**, 010313 (2022).
- [31] Edward Grant, Marcello Benedetti, Shuxiang Cao, Andrew Hallam, Joshua Lockhart, Vid Stojevic, Andrew G. Green,

- and Simone Severini, Hierarchical quantum classifiers, *npj Quantum Inf.* **4**, 65 (2018).
- [32] Arthur Pesah, M. Cerezo, Samson Wang, Tyler Volkoff, Andrew T. Sornborger, and Patrick J. Coles, Absence of barren plateaus in quantum convolutional neural networks, *Phys. Rev. X* **11**, 041011 (2021).
- [33] Tak Hur, Leeseok Kim, and Daniel K. Park, Quantum convolutional neural network for classical data classification, *Quantum Mach. Intell.* **4**, 3 (2022).
- [34] Matt Lourens, Ilya Sinayskiy, Daniel K. Park, Carsten Blank, and Francesco Petruccione, Hierarchical quantum circuit representations for neural architecture search, *npj Quantum Inf.* **9**, 79 (2023).
- [35] Juhyeon Kim, Joonsuk Huh, and Daniel K. Park, Classical-to-quantum convolutional neural network transfer learning, *Neurocomputing* **555**, 126643 (2023).
- [36] M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, *Nat. Commun.* **12**, 1791 (2021).
- [37] Nayeli A. Rodríguez-Briones, Jun Li, Xinhua Peng, Tal Mor, Yossi Weinstein, and Raymond Laflamme, Heat-bath algorithmic cooling with correlated qubit-environment interactions, *New J. Phys.* **19**, 113047 (2017).
- [38] Nayeli A. Rodríguez-Briones, Eduardo Martín-Martínez, Achim Kempf, and Raymond Laflamme, Correlation-enhanced algorithmic cooling, *Phys. Rev. Lett.* **119**, 050502 (2017).
- [39] Álvaro M. Alhambra, Matteo Lostaglio, and Christopher Perry, Heat-bath algorithmic cooling with optimal thermalization strategies, *Quantum* **3**, 188 (2019).
- [40] Philip Taranto, Patryk Lipka-Bartosik, Nayeli A. Rodríguez-Briones, Martí Perarnau-Llobet, Nicolai Friis, Marcus Huber, and Pharnam Bakhshinezhad, Efficiently cooling quantum systems with finite resources: Insights from thermodynamic geometry, *Phys. Rev. Lett.* **134**, 070401 (2025).
- [41] Tak Hur, Israel F. Araujo, and Daniel K. Park, Neural quantum embedding: Pushing the limits of quantum supervised learning, *Phys. Rev. A* **110**, 022411 (2024).
- [42] Carlos Ortiz Marrero, Mária Kieferová, and Nathan Wiebe, Entanglement-induced barren plateaus, *PRX Quantum* **2**, 040316 (2021).
- [43] Taylor L. Patti, Khadijeh Najafi, Xun Gao, and Susanne F. Yelin, Entanglement devised barren plateau mitigation, *Phys. Rev. Res.* **3**, 033090 (2021).
- [44] Jose M. Fernandez, Seth Lloyd, Tal Mor, and Vwani Roychowdhury, Algorithmic cooling of spins: A practicable method for increasing polarization, *Int. J. Quantum Inf.* **2**, 461 (2004).
- [45] Jonathan Baugh, Osama Moussa, Colm A. Ryan, Ashwin Nayak, and Raymond Laflamme, Experimental implementation of heat-bath algorithmic cooling using solid-state nuclear magnetic resonance, *Nature* **438**, 470 (2005).
- [46] C. A. Ryan, O. Moussa, J. Baugh, and R. Laflamme, Spin based heat engine: Demonstration of multiple rounds of algorithmic cooling, *Phys. Rev. Lett.* **100**, 140501 (2008).
- [47] Rodolfo R. Soldati, Durga B. R. Dasari, Jörg Wrachtrup, and Eric Lutz, Thermodynamics of a minimal algorithmic cooling refrigerator, *Phys. Rev. Lett.* **129**, 030601 (2022).
- [48] Yosi Atia, Yuval Elias, Tal Mor, and Yossi Weinstein, Algorithmic cooling in liquid-state nuclear magnetic resonance, *Phys. Rev. A* **93**, 012325 (2016).
- [49] Krishna Shende, Arvind, and Kavita Dorai, State-independent robust heat-bath algorithmic cooling of nuclear spins, *Phys. Rev. Appl.* **21**, 024017 (2024).
- [50] Xueyuan Hu and Valerio Scarani, Thermalization of finite complexity and its application to heat-bath algorithmic cooling, *Phys. Rev. A* **112**, 042209 (2025).
- [51] Jake Xuereb, Benjamin Stratton, Alberto Rolandi, Jiming He, Marcus Huber, and Pharnam Bakhshinezhad, Cooling a qubit using  $n$  others, *PRX Quantum* **6**, 040368 (2025).
- [52] B. L. WELCH, The generalization of “STUDENT’S” problem when several different population variances are involved, *Biometrika* **34**, 28 (1947).
- [53] Tak Hur and Daniel K. Park, Understanding Generalization in Quantum Machine Learning with Margins, in *Proceedings of the 42nd International Conference on Machine Learning*, Proceedings of Machine Learning Research, edited by Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (PMLR, Vancouver, Canada, 2025), Vol. 267, pp. 26338–26360, <https://proceedings.mlr.press/v267/hur25a.html>.
- [54] Sofiene Jerbi, Lukas J. Fiderer, Hendrik Poulsen Nautrup, Jonas M. Kübler, Hans J. Briegel, and Vedran Dunjko, Quantum machine learning beyond kernel methods, *Nat. Commun.* **14**, 517 (2023).
- [55] Martin Larocca, Supanut Thanasilp, Samson Wang, Kunal Sharma, Jacob Biamonte, Patrick J. Coles, Lukasz Cincio, Jarrod R. McClean, Zoë Holmes, and M. Cerezo, A review of barren plateaus in variational quantum computing, *Nat. Rev. Phys.* **7**, 174 (2025).
- [56] Carsten Blank, Daniel K. Park, June-Koo Kevin Rhee, and Francesco Petruccione, Quantum classifier with tailored quantum kernel, *npj Quantum Inf.* **6**, 1 (2020).
- [57] Daniel K. Park, Carsten Blank, and Francesco Petruccione, *Robust quantum classifier with minimal overhead*, in *2021 International Joint Conference on Neural Networks (IJCNN)* (2021), pp. 1–7.
- [58] M. Schuld, M. Fingerhuth, and F. Petruccione, Implementing a distance-based classifier with a quantum interference circuit, *Europhys. Lett.* **119**, 60002 (2017).
- [59] Daniel K. Park, Carsten Blank, and Francesco Petruccione, The theory of the quantum kernel-based binary classifier, *Phys. Lett. A* **384**, 126422 (2020).
- [60] Carsten Blank, Adenilton J. da Silva, Lucas P. de Albuquerque, Francesco Petruccione, and Daniel K. Park, Compact quantum kernel-based binary classifier, *Quantum Sci. Technol.* **7**, 045007 (2022).
- [61] Nicolas M. de Oliveira, Daniel K. Park, Israel F. Araujo, and Adenilton J. da Silva, Quantum variational distance-based centroid classifier, *Neurocomputing* **576**, 127356 (2024).
- [62] Jungyun Lee and Daniel K. Park, Quadratic speed-ups in quantum kernelized binary classification, *Adv. Quantum Technol.* **7**, 2400126 (2024).
- [63] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd, Quantum support vector machine for big data classification, *Phys. Rev. Lett.* **113**, 130503 (2014).

- [64] Morten Kjaergaard, Mollie E. Schwartz, Jochen Braumüller, Philip Krantz, Joel I.-J. Wang, Simon Gustavsson, and William D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [65] Petar Jurcevic, Ali Javadi-Abhari, Lev S. Bishop, Isaac Lauer, Daniela F. Bogorin, Markus Brink, Lauren Capelluto, Oktay Günlük, Toshinari Itoko, Naoki Kanazawa, *et al.*, Demonstration of quantum volume 64 on a superconducting quantum computing system, *Quantum Sci. Technol.* **6**, 025020 (2021).
- [66] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A. Buell, *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [67] Jay M. Gambetta, Jerry M. Chow, and Matthias Steffen, Building logical qubits in a superconducting quantum computing system, *npj Quantum Inf.* **3**, 2 (2017).
- [68] Kenneth Wright, Kristin M. Beck, Sea Debnath, J. M. Amini, Y. Nam, N. Grzesiak, J.-S. Chen, N. C. Pimenti, M. Chmielewski, C. Collins, *et al.*, Benchmarking an 11-qubit quantum computer, *Nat. Commun.* **10**, 5464 (2019).
- [69] Juan M. Pino, Jennifer M. Dreiling, Caroline Figgatt, John P. Gaebler, Steven A. Moses, M. S. Allman, C. H. Baldwin, Michael Foss-Feig, David Hayes, Karl Mayer, *et al.*, Demonstration of the trapped-ion quantum CCD computer architecture, *Nature* **592**, 209 (2021).
- [70] Pascal Scholl, Michael Schuler, Hannah J. Williams, Alexander A. Eberharter, Daniel Barredo, Kai-Niklas Schymik, Vincent Lienhard, Louis-Paul Henry, Thomas C. Lang, Thierry Lahaye, *et al.*, Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms, *Nature* **595**, 233 (2021).
- [71] Loïc Henriët, Lucas Beguin, Adrien Signoles, Thierry Lahaye, Antoine Browaeys, Georges-Olivier Raymond, and Christophe Jurczak, Quantum computing with neutral atoms, *Quantum* **4**, 327 (2020).
- [72] Simon J. Evered, Dolev Bluvstein, Marcin Kalinowski, Sepehr Ebadi, Tom Manovitz, Hengyun Zhou, Sophie H. Li, Alexandra A. Geim, Tout T. Wang, Nishad Maskara, *et al.*, High-fidelity parallel entangling gates on a neutral-atom quantum computer, *Nature* **622**, 268 (2023).
- [73] Michael A. Nielsen and Isaac L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, New York, 2011), 10th ed.
- [74] Heinz-Peter Breuer and Francesco Petruccione, *The Theory of Open Quantum Systems* (OUP, Oxford, 2002).
- [75] Ali Javadi-Abhari, Matthew Treinish, Kevin Krsulich, Christopher J. Wood, Jake Lishman, Julien Gacon, Simon Martiel, Paul D. Nation, Lev S. Bishop, Andrew W. Cross, Blake R. Johnson, and Jay M. Gambetta, Quantum computing with Qiskit, arXiv preprint [arXiv:2405.08810](https://arxiv.org/abs/2405.08810)
- [76] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, Scikit-learn: Machine learning in Python, *J. Mach. Learn. Res.* **12**, 2825 (2011).
- [77] R. A. Fisher, Iris. UCI Machine Learning Repository, 1936. <https://doi.org/10.24432/C56C76>.
- [78] Stefan Aeberhard and M. Forina, Wine. UCI Machine Learning Repository, 1992. <https://doi.org/10.24432/C5PC7J>.
- [79] E. Alpaydin and C. Kaynak, Optical Recognition of Handwritten Digits. UCI Machine Learning Repository, 1998. <https://doi.org/10.24432/C50P49>.
- [80] Terry Sejnowski and R. Gorman, Connectionist Bench (Sonar, Mines vs. Rocks). UCI Machine Learning Repository, 1988. <https://doi.org/10.24432/C5T01Q>.
- [81] Jack W. Smith, Using the ADAP learning algorithm to forecast the onset of diabetes mellitus, in *Proceedings of the Annual Symposium on Computer Application in Medical Care* (IEEE Computer Society Press, 1988), pp. 261–265.
- [82] N. A. Rodríguez-Briones and D. K. Park, qDNA-yonsei. GitHub Repository, 2025. <https://github.com/qDNA-yonsei/BQR4QML>.