

# Autonomous quantum machine learning via Riemannian gradient-based optimization

Kasidit Srimahajariyapong, Thiparat Chotibut, Suwannachad Suwannajitt, Tithivich Paka, Ricard Puig, Yudai Suzuki, Marek Gluza, Zoë Holmes, Supanut Thanasilp

*This abstract is based on Ref. [1]*

**Introduction.** Training quantum machine learning models typically requires extracting classical information through quantum measurements. In variational approaches [2, 3] such as parametrized quantum circuits [4], this entails repeated estimation of loss functions and gradients within hybrid quantum–classical feedback loops, whose statistical errors are fundamentally constrained by quantum measurements. On the other hand, although less apparent, non-variational models face similar limitations. Quantum kernel methods require estimating data-independent quantities such as Gram matrices [5–7], while quantum reservoir computing relies on expectation values of observables [8–10].

The core of the problem in all these cases are arguably the use of a classical optimizer to solve a quantum problem. This means that information in the quantum model needs to be repeatedly converted to a classical cost function for classical optimization. The classical cost function (or kernel values) is obtained statistically by repeatedly measuring the output of a quantum circuit. If cost function gradients are very small then the statistical noise resulting from this quantum-to-classical conversion process will wash out any meaningful information about how to train the quantum model. However, if we could train on the exact cost function values, then variational quantum algorithms are guaranteed to converge in the over-parameterized regime [11]. Thus it is the combination of the exponential concentration and the inevitability of shot noise that inhibits the scaling of these methods. This motivates the bold question: *could we banish the classical optimizer from the optimization loop and optimize entirely within the quantum domain?*

The concept of autonomous, self-training operation has begun to emerge. Approaches based on encoding optimization problem into dynamical or dissipative dynamics have been proposed in Ref. [12–14]. In addition, although not originally framed as self-training protocols, a series of works on circuit synthesis for Riemannian gradient-based optimization (known as double-bracket quantum algorithms) can be effectively viewed as operating under the same principles in the related field of variational quantum algorithm [15–18]. Together, these developments provide the first solid step into fully quantum, self-contained learning dynamics. However, for all intents and purposes, *these methods have not yet been explored in the context of quantum machine learning.*

**Summary of our results.** In this work, we propose and analyze the performance of autonomous self-training quantum models based on Riemannian gradient optimization, summarised in Fig. 1. First, we define autonomous quantum machine learning as a regime in which, after encoding the training data, no further interaction between quantum and classical resources is required; the optimal model is obtained entirely through coherent quantum evolution without intermediate information extraction. Hence, the training process proceeds fully within the quantum domain.

We then introduce a family of quantum models based on Riemannian gradient optimization, for which the gradient expressions for generic nonlinear loss functions are rigorously derived from first principles. This framework is flexible and can accommodate a wide range of learning tasks, including classification, regression, and generative modeling, through appropriate choices of loss functions. However, the resulting training dynamics do not, in general, correspond to the fully autonomous paradigm, as estimating certain intermediate quantities may still require measurement and classical processing.

Moving forward, we identify sufficient conditions under which autonomous quantum machine learning naturally emerges within this framework. We establish a guarantee of exponential decrease in the training loss for sufficiently small step sizes. Crucially, we further show that if the training data and the underlying learning problem possess hidden symmetry, the trained model automatically becomes aware of this symmetry without any additional information being provided. This constitutes another key advantage of our approach, enabling it to exploit symmetries present in the training data even when they are not known *a priori*.

While the training loss decreases exponentially in the number of iterations, similarly to other ‘double-bracket’ methods, our circuits have a recursive structure which means that the circuit depth grows exponentially with the number of training iterations, with the base of the exponent increasing with the size of the training dataset. This could pose a potential limitation if good prediction accuracy cannot be achieved with a small number of training iterations and limited training data. Encouragingly, when we consider canonical examples of unitary learning with hidden symmetry [19, 20] and quantum phase classification [21] (tasks that naturally fit within the autonomous learning paradigm), our results suggest that effective learning can still be achieved at low, constant circuit depths. Overall,




**Figure 1: Overview.** Pictorial comparison between (a) a hybrid classical-quantum loop in typical quantum machine learning and (b) a fully quantum loop in autonomous quantum machine learning. (c) Riemannian gradient optimization as an example of autonomous quantum machine learning.




Figure 2: **Performance of phase classification for the generalized cluster model** (a) The training (solid) and testing (dashed) loss functions are plotted against the number of training iterations (b) the test accuracy is plotted against the number of training iterations. (c) The predicted phases in  $J_1 - J_2$  plane after 1 iteration are shown for all testing points. Each color represents a different phase. with black crosses ( $\times$ ) indicating the sparse 5 training samples.

our approach provides an alternative framework for quantum machine learning based on autonomous, self-training quantum dynamics.

**Quantum machine learning with Riemannian gradients.** We consider a setting where a set of  $N$  training input quantum states and output labels  $\{\rho_i, y_i\}_{i=1}^N$  is given. The goal is to find a unitary operator  $U$  such that it minimizes a **loss function of the form**  $\mathcal{L}(U) = \sum_i f(\ell_i(U), y_i)$  where  $\ell_i(U) = \text{Tr}(U\rho_i U^\dagger O_i)$  denotes the expectation value of a (potential index dependent) observable  $O_i$ , and  $f(\cdot, \cdot)$  is an arbitrary non-linear differentiable function which compares how close two arguments are. Note that the outputs  $\{y_i\}_i$  can be scalar, vector quantities or even quantum states themselves. This setting covers a wide range of learning tasks, including classification [21], regression [19, 20], and (even though less obvious) generative modeling [22].

In the standard parametrized circuit approach, the unitary operator  $U(\theta)$  is parametrized and updated by tuning the parameters  $\theta$  via a classical optimizer. In contrast, here the unitary operator is directly updated via the Riemannian gradient over the unitary manifold  $\nabla(\mathcal{L})_U$  [23, 24]. This gradient is defined as the unique tangent vector at  $U$  on the unitary manifold such that it points into the steepest direction for the function  $\mathcal{L}$  at  $U$ . In our work, we derive the Riemannian gradient for the generic loss (with the step-size  $s$ ) through the first principles for a *single* step leading to  $\nabla(\mathcal{L})_U := g = -\sum_i [U\rho_i U^\dagger, O_i] \left( \frac{\partial f}{\partial \ell_i} \right)_{s_k=0}$ .

By recursively applying the Riemannian gradient for each training iteration, this results in an iterative prescription for the trained quantum circuit on iteration  $k$  (with  $U_0$  as some fixed initial unitary) as

$$U_{k+1} = e^{-s_k g_k} U_k ; \quad \text{given} \quad g_k = -\sum_i [\rho_i^{(k)}, O_i] \left( \frac{\partial f}{\partial \ell_i} \right)_{s_k=0} \quad \text{and} \quad \rho_i^{(k)} = U_k \rho_i U_k^\dagger . \quad (1)$$

Crucially, this recursive relation between consecutive iterations provides an explicit quantum circuit corresponding to the optimized unitary for a training iteration of our interest. Here, since  $\{(\partial f / \partial \ell_i)_{s_k=0}\}_i$  for  $k$  iteration can in general depend on the previous updated state at  $k-1$  iteration, the circuit implementation may still require the classical data extraction to estimate  $\{(\partial f / \partial \ell_i)_{s_k=0}\}_i$  for each intermediate iteration. Nevertheless, we show a sufficient condition that if the function  $f$  is linear in the expectation value  $\{\ell_i\}_i$ , the description of the trained circuit at any iteration requires no quantum measurements. More precisely, if  $\frac{\partial f}{\partial \ell_i} = c_i$  where  $c_i$  is some constant (independent of the intermediate states), then one can recursively use Eq. (1) to determine the whole circuit architecture that corresponds to the trained model (for  $k$  iteration), allowing for a fully autonomous execution of the training process on the quantum hardware i.e., *just run the given quantum circuit*. Prominent examples of such learning tasks include phase classification [21] and learning unitary [19, 20].

Further analysis on this autonomous framework shows two key properties. First, for sufficiently small step sizes, we show that the training loss is decreased exponentially with increasing training iterations. That is, if  $s_k \in \mathcal{O}(\|g_k\|_2^2 / M_k)$  where  $M_k$  is the maximum curvature within the range  $[0, s_k]$  for all  $k$ , then the loss converges according to:  $\mathcal{L}(U_k) - \mathcal{L}(U_\infty) \sim h^{-k}$  with the convergence factor  $h > 1$ . Next, the framework has an inherent mechanism for becoming aware of the problem's symmetry directly from the training data. More formally, let  $S$  be any operator, and suppose the data  $\{\rho_i, O_i\}$  together with initial unitary  $U_0$  satisfy  $[S, \rho_i] = [S, O_i] = [S, U_0] = 0$  for all  $i$ . Then, the updated unitary also satisfies  $[S, U_k] = 0$  for all  $k$ . As a result, our approach naturally exploits symmetries present in the training data even when they are not informed.

However, one must be mindful of the circuit depth resulting from the trained model. The recursive update relation leads to multiplicative growth of the circuit depth with increasing training iterations once the full circuit is unfolded. The resulting circuit depth scales as  $\sim \mathcal{O}(N^k)$ . In practice, however, other techniques, such as stochastic gradient methods and circuit compression, may be employed to significantly reduce the circuit depth.

**Canonical examples.** Although various techniques can be employed to reduce the circuit depth, it is encouraging to observe that effective training is possible even without relying on such heuristic approaches. To demonstrate this, as well as to illustrate how our framework operates in practice, we consider two learning tasks: (i) quantum phase classification and (ii) learning an unknown quantum unitary.

(i) Quantum phase classification. Consider a Hamiltonian  $H(\alpha)$  that exhibits different phases depending on the

parameter  $\alpha$ . Given  $N$  pairs of ground states of this Hamiltonian (each corresponding to a different  $\alpha$ ) and their associated phase labels  $\{\rho_i, y_i\}_i$ , the task is to train a model that predicts the phase of an unseen ground state. Here the loss function can be expressed as  $\mathcal{L}(U) = 1 - \frac{1}{N} \sum_{i=1}^N \text{Tr}(P_{y_i} U \rho_i U^\dagger)$  where  $P_{y_i}$  is a class-dependent projector mapping the output state onto the orthogonal subspace associated with the label  $y_i$ , with the additional condition that  $\sum_{y_i} P_{y_i} = \mathbb{1}_s$ , where  $\mathbb{1}_s$  denotes the identity operator on this subspace. Because this loss function is linear, the Riemannian gradient optimization proceeds fully autonomously. After obtaining the trained unitary  $U_*$ , the phase of an unseen state is predicted by selecting the label corresponding to the largest probability  $\text{Tr}(P_{y_i} U_* \rho_i U_*^\dagger)$ .

We consider the 1D generalized cluster model  $H(J_1, J_2) = \sum_{i=1}^n (Z_i - J_1 X_i X_{i+1} - J_2 X_{i-1} Z_i X_{i+1})$  where  $X_i$  and  $Z_i$  denote the Pauli- $X$  and Pauli- $Z$  operators acting on qubit  $i$ , respectively, and  $J_1$  and  $J_2$  are coupling parameters. This model exhibits four distinct topological and symmetry-breaking phases depending on the values of  $J_1$  and  $J_2$ . The input quantum states correspond to many-body ground states at different parameter coordinates. We simulate an  $n = 8$  qubit system using a step size of  $s = 0.1$ . The algorithm is provided with an extremely sparse training dataset consisting of only five data points, indicated by the black crosses ( $\times$ ) in Fig. 2(c). Figure 2(a) shows monotonic convergence of the training loss together with a decrease in the testing loss, reaching final values of 0.04 and 0.22, respectively, after  $k = 20$  training iterations. Despite this severe data sparsity, the algorithm achieves a test accuracy of 0.95 after only  $k \sim 2$  training iterations across a dense grid of 10,201 unseen test samples, as shown in Fig. 2(b), accurately reconstructing the phase boundaries (Fig. 2(c)). The high prediction accuracy despite the nonzero testing loss arises because correct classification requires only that the probability assigned to the correct phase be the largest among all classes. Note that the underlying mechanism for distinguishing the phases can be shown to rely solely on global comparisons of overlaps between different quantum states. From this perspective, our framework operates under the same underlying principle as global fidelity kernels, but in a fully autonomous manner.

(ii) *Learning an unknown quantum unitary.* We now proceed to consider a slightly more complex learning task for which our approach is particularly well suited. In particular, we aim to leverage the model's inherent ability to detect hidden symmetries. Given only a training dataset, determining whether the underlying learning problem possesses symmetry is, in general, a difficult task. In contrast, our approach automatically identifies and exploits such symmetries directly from the training data. To demonstrate this, we consider unitary learning with some symmetry.

Our goal is to learn an unknown dynamics  $V$  through the training data pairs of input and output states  $\{|\psi_i\rangle, V|\psi_i\rangle\}_i$  (where  $|\psi_i\rangle$  is drawn from some distribution) such that after training  $U_*|\psi\rangle \approx V|\psi\rangle$  for an unseen  $|\psi\rangle$  drawn from the same distribution. The loss function in this case can be expressed as  $\mathcal{L}(U) = 1 - \frac{1}{N} \sum_{i=1}^N F(U|\psi_i\rangle, V|\psi_i\rangle)$  where  $F(|\Psi\rangle, |\Phi\rangle)$  is a fidelity between two states  $|\Psi\rangle$  and  $|\Phi\rangle$ . This form of the loss again satisfies the autonomous sufficient condition.

Here we consider the quantum dynamics,  $V(t) = e^{-itH}$  generated by the Hamiltonian  $H = \sum_{i=1}^{n-1} (X_i X_{i+1} + Y_i Y_{i+1})$ . In particular, this Hamiltonian evolution conserves the number operator  $(n - \sum_i Z_i)/2$  counting the number of qubits in the  $-1$  eigenspace of  $Z_i$  (number of excitations), since  $[\sum_i Z_i, H] = 0$ . Consequently, the Hilbert space can be decomposed into invariant subspaces labeled by the number of excitations. We further consider that the input states of the form  $|\psi\rangle = \frac{1}{\mathcal{Z}} \sum_x (a_x + ib_x)|x\rangle$  where the coefficients  $a_x$  and  $b_x$  are drawn from a normal distribution i.e.,  $a_x, b_x \in \mathcal{N}(0, 1)$  and  $|x\rangle$  represents a computational basis with Hamming weight 1, and  $\mathcal{Z}$  is a normalization. Essentially, these input states are drawn from a Haar random ensemble of the  $n$ -dimensional single-excitation subspace.

To make a comparison and show the power of symmetry aware, we compare this training dataset to a problem-agnostic dataset, where the training dataset are Haar states drawn from the full  $2^n$ -dimensional Hilbert space. As shown in Fig. 3, a symmetric dataset successfully recovers the unitary action on the invariant subspace using  $\mathcal{O}(\text{poly}(n))$  samples. Specifically with  $N = 20$  at  $k = 20$  iterations, the model achieves a reasonable good testing loss of  $10^{-3}$  and demonstrates robust generalization, i.e., the testing loss decrease alongside the training loss. Conversely,  $N = 20$  samples from the problem-agnostic training dataset are insufficient to achieve generalization as the testing loss does not change its order of magnitude despite a decrease in the training loss.

**Conclusion.** We introduced a framework for autonomous quantum machine learning based on Riemannian gradient optimization, in which training can proceed entirely within the quantum domain under suitable conditions. Beyond establishing guarantees on the training dynamics and uncovering an inherent mechanism for exploiting hidden symmetries, we identified the circuit-depth scaling as a key practical challenge. Nevertheless, our canonical examples demonstrate that effective learning can still be achieved at low circuit depths, highlighting the potential of autonomous quantum learning as an alternative paradigm for quantum machine learning.




Figure 3: **Performance of unitary learning** (a) The testing loss functions are plotted against the number of training samples. (b) The training (solid) and testing (dashed) loss functions using  $N = 20$  are plotted against the number of training iterations (c) The commutators norm  $\|[\sum_i Z_i, U_k]\|_2$  between the symmetric generator and the trained unitary are plotted against the number of training iteration for verifying that the algorithm is aware of the problem's symmetry from the training data.

## References

- [1] Kasidit Srimahajariyapong, Thiparat Chotibut, Suwannachad Suwannajitt, Tithivich Paka, Ricard Puig, Yudai Suzuki, Marek Gluza, Zoë Holmes, and Supanut Thanasilp. Autonomous quantum machine learning via riemannian gradient-based optimization. *manuscript in preparation*.
- [2] 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. *Nature Reviews Physics*, 3(1):625–644, 2021.
- [3] Kishor Bharti, Tobias Haug, Vlatko Vedral, and Leong-Chuan Kwek. Noisy intermediate-scale quantum algorithm for semidefinite programming. *Physical Review A*, 105(5):052445, 2022.
- [4] Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001, 2019.
- [5] Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Physical Review Letters*, 122(4):040504, 2019.
- [6] Vojtěch Havlíček, 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(7747):209–212, 2019.
- [7] Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R McClean. Power of data in quantum machine learning. *Nature Communications*, 12(1):1–9, 2021.
- [8] Keisuke Fujii and Kohei Nakajima. Harnessing disordered-ensemble quantum dynamics for machine learning. *Physical Review Applied*, 8(2):024030, 2017.
- [9] Kohei Nakajima, Keisuke Fujii, Makoto Negoro, Kosuke Mitarai, and Masahiro Kitagawa. Boosting computational power through spatial multiplexing in quantum reservoir computing. *Physical Review Applied*, 11(3):034021, 2019.
- [10] Pere Mujal, Rodrigo Martí nez-Peña, Johannes Nokkala, Jorge García-Bení, Gian Luca Giorgi, Miguel C. Soriano, and Roberta Zambrini. Opportunities in quantum reservoir computing and extreme learning machines. *Advanced Quantum Technologies*, 4(8):2100027, 2021.
- [11] Xuchen You, Shouvanik Chakrabarti, and Xiaodi Wu. A convergence theory for over-parameterized variational quantum eigensolvers. *arXiv preprint arXiv:2205.12481*, 2022.
- [12] Danial Motlagh, Modjtaba Shokrian Zini, Juan Miguel Arrazola, and Nathan Wiebe. Ground state preparation via dynamical cooling. *arXiv preprint arXiv:2404.05810*, 2024.
- [13] Ahmet Burak Catli, Sophia Simon, and Nathan Wiebe. Exponentially better bounds for quantum optimization via dynamical simulation. *arXiv preprint arXiv:2502.04285*, 2025.
- [14] Clara C Wanjura and Florian Marquardt. Quantum equilibrium propagation for efficient training of quantum systems based on onsager reciprocity. *Nature Communications*, 16(1):6595, 2025.
- [15] Marek Gluza. Double-bracket quantum algorithms for diagonalization. *Quantum*, 8:1316, April 2024.
- [16] Yudai Suzuki, Bi Hong Tiang, Jeongrak Son, Nelly HY Ng, Zoë Holmes, and Marek Gluza. Double-bracket algorithm for quantum signal processing without post-selection. *Quantum*, 9:1954, 2025.
- [17] Marek Gluza, Jeongrak Son, Bi Hong Tiang, René Zander, Raphael Seidel, Yudai Suzuki, Zoë Holmes, and Nelly HY Ng. Double-bracket quantum algorithms for quantum imaginary-time evolution. *Physical review letters*, 136(2):020601, 2026.
- [18] Andrew Wright, Reyhaneh Aghaei Saem, Supanut Thanasilp, Yudai Suzuki, and Zoë Holmes. Double-bracket quantum algorithms for thermal state preparation. *arXiv preprint arXiv:2606.05947*, 2026.
- [19] Hsin-Yuan Huang, Sitan Chen, and John Preskill. Learning to predict arbitrary quantum processes. *PRX Quantum*, 4(4):040337, 2023.
- [20] Matthias C Caro, Hsin-Yuan Huang, Nicholas Ezzell, Joe Gibbs, Andrew T Sornborger, Lukasz Cincio, Patrick J Coles, and Zoë Holmes. Out-of-distribution generalization for learning quantum dynamics. *Nature Communications*, 14(1):3751, 2023.
- [21] Matthias C Caro, Hsin-Yuan Huang, M Cerezo, Kunal Sharma, Andrew Sornborger, Lukasz Cincio, and Patrick J Coles. Generalization in quantum machine learning from few training data. *Nature Communications*, 13, 2022.
- [22] Manuel S Rudolph, Sacha Lerch, Supanut Thanasilp, Oriel Kiss, Oxana Shaya, Sofia Vallecorsa, Michele Grossi, and Zoë Holmes. Trainability barriers and opportunities in quantum generative modeling. *npj Quantum Information*, 10(1):116, 2024.
- [23] Roeland Wiersema and Nathan Killoran. Optimizing quantum circuits with riemannian gradient flow. *Physical Review A*, 107(6):062421, 2023.
- [24] René Zander, Raphael Seidel, Li Xiaoyue, and Marek Gluza. Role of riemannian geometry in double-bracket quantum imaginary-time evolution. In *International Conference on Geometric Science of Information*, pages 105–114. Springer, 2025.