

Machine learning with minimal use of quantum computers: Provable advantages in Learning Under Quantum Privileged Information (LUQPI)

<https://arxiv.org/abs/2601.22006>

Vasily Bokov^{1,2,3}, Lisa Kohl⁴, Sebastian Schmitt³, and Vedran Dunjko^{1,2}

¹ $\langle \text{aQa}^L \rangle$, Leiden University, The Netherlands

² LIACS, Leiden University, Leiden, The Netherlands

³ Honda Research Institute Europe GmbH, Offenbach, Germany

⁴ Cryptology Group, CWI Amsterdam, The Netherlands

Abstract. Quantum machine learning (QML) is often listed as a promising candidate for useful applications of quantum computers, in part due to numerous proofs of possible quantum advantages. A central question is how small a role quantum computers can play while still enabling provable learning advantages over classical methods. We study an especially restricted setting where a quantum computer is used *only* as a feature extractor: it acts independently on individual data points, without access to labels or global dataset information, and furthermore assumed to be available only as means to augment the training set, and is not available in deployment. The training and deployment are carried out by fully classical learners on a dataset augmented with quantum-generated features.

We formalize this model by adapting the classical framework of Learning Under Privileged Information (LUPI) to the quantum case, which we call Learning Under Quantum Privileged Information (LUQPI). Within this framework, we show that even such minimally involved quantum feature extraction available only for the training data can nonetheless yield exponential quantum-classical separations for suitable concept classes and distributions, under reasonable computational assumptions. We further situate LUQPI in a taxonomy of related quantum and classical settings and also show how standard classical machinery—most notably the SVM+ algorithm—can exploit quantum-augmented data. In the latter direction, we present numerical experiments in a physically motivated many-body setting, where privileged quantum features are ground state expectation values, and observe consistent performance gains for LUQPI-style models over strong classical baselines.

This work presents formal grounding of a long-standing idea, that data driven methods can distil useful information from hard-to-characterize sources (e.g. noisy or otherwise limited quantum computers) where analytic treatment or explicit programming is out of reach.

1 Introduction

Quantum machine learning (QML) is often regarded as a promising area where quantum computers might outperform classical information processing, with a number of results now suggesting that quantum models can, at least in principle, beat classical learners on carefully designed tasks [8,6,11]. Beyond establishing that some advantage is possible, an intriguing question is: *what is the minimal way in which a quantum computer needs to be involved in the learning pipeline to obtain such an advantage?* Recent work has shown provable speed-ups when a quantum computer performs the entire training while deployment is classical [11], or when only the inference stage is quantum [12,14].

In this work, we push this “minimal involvement” perspective much further. We ask whether one can obtain provable advantages when the quantum computer is used only as a *feature extractor*: it acts on each individual input datapoint and has no access to labels nor any global property of the dataset. In particular, the quantum device never sees training labels and never performs end-to-end optimization. It is only allowed to compute additional features that are then handed to a classical learner in an ‘augmented’ dataset, and critically, the feature extraction is not available in deployment, i.e., the trained model is fully classical in every sense. Conceptually, our approach mirrors (quantum) topological data analysis [13,5]: a hard-to-compute quantum procedure extracts features from individual inputs that reveal structure, making downstream classical learning easier, while keeping the quantum role as restricted as possible. A similar picture arises in chemistry, where full ground-state energies may be out of reach, but the data can be augmented with harder-to-characterize features such as trapped-state properties [3].

Formally, the model we study aligns with the classical framework of Learning Under Privileged Information (LUPI) [17], in which, during training, a learner has access to additional information that will not be available at deployment time for new, to-be-labeled, points. Our work can be understood as an instantiation of this framework with an efficient quantum algorithm that analyzes data points in an i.i.d. fashion, and hence we refer to our setting as Learning Under Quantum Privileged Information (LUQPI).

We highlight that LUQPI constitutes a significantly more restrictive usage scenario of quantum computers compared to previous quantum-enhanced approaches, many of which require the quantum algorithm to access labels, to process several data points jointly, or to control the full training loop.

The contributions of this paper are as follows.

- (i) We formally define advantageous learning scenarios with quantum feature extraction and introduce two natural versions: *quantum online*, where feature extraction is available also at inference, and—our key and much more restrictive model—the *offline* LUQPI setting, in which the quantum device is used only during training. Perhaps surprisingly, provable advantages persist even under this stronger restriction.
- (ii) We prove that even in this substantially constrained LUQPI setting, one can construct concept classes for which quantum feature extraction enables exponential advantages over any efficient classical learner under reasonable complexity-theoretic assumptions. The advantage even holds against non-uniform learners (explained below).

- (iii) We provide extensive numerical experiments for a physically motivated problem where quantum-privileged information can be computed from ground states of many-body systems. These experiments show that providing privileged features during training can, in certain cases, improve the performance of classical learners, even when those features are unavailable at deployment (i.e. in the test phase).

2 The LUQPI Framework

We build our framework on top of the classical PAC learning approach, extended to accommodate quantum-generated features available only at training time. A learning task is represented by a concept class—a family of target functions the learner aims to identify from labeled examples—and is *efficiently learnable* if some algorithm, given polynomially many labeled examples, outputs a hypothesis with small error in polynomial time.

For each input size n , let $\mathcal{X}_n$ be a domain of instances, $\mathcal{Y}$ a label space, and $\mathcal{C} = \bigcup_n \mathcal{C}_n$ a concept class with $\mathcal{C}_n \subseteq \mathcal{Y}^{\mathcal{X}_n}$. A learner receives i.i.d. labelled examples $(x, c(x))$ with $x \sim D_n$ from a distribution class $\mathcal{D} = \bigcup_n D_n$, and must output a hypothesis h with $\Pr_{x \sim D_n}[h(x) \neq c(x)] \leq \epsilon$ with probability $\geq 1 - \delta$, using $\text{poly}(n, 1/\epsilon, \log(1/\delta))$ samples and time. A *quantum learning advantage* holds when $(\mathcal{C}, \mathcal{D})$ is quantumly but not classically efficiently learnable. We distinguish *uniform* (BPP) from *non-uniform* (P/poly) classical learners, the latter receiving polynomial-sized advice. Advantages against non-uniform learners are a strictly stronger statement, ruling out even classical learners with arbitrary precomputed advice (capturing, e.g., algorithms pre-trained on related tasks).

A *feature extractor* family $\mathcal{E} = \{E_n : \mathcal{X}_n \rightarrow \mathcal{F}_n\}_n$ converts raw inputs into additional features, and the *extended example oracle* $\text{EX}_{\text{ext}}(c, D_n, E_n)$ returns samples $((x, E_n(x)), c(x))$ with $x \sim D_n$. We require E_n to be quantumly efficiently evaluable with polynomial-sized output, while learning and use of features are classically efficient. We formalise two regimes: in the **online** setting the extractor is available at both training and deployment (the features $E_n(x)$ are recomputed for each fresh input), while in the **offline** setting, which we call *LUQPI*, the quantum device is used only once, in a data-preprocessing stage, to compute features for the training set. In both cases the quantum subroutine acts on single points and never sees labels, leading to the following definition.

Definition 1 (Quantum-offline advantageous feature extraction, informal). $\mathcal{E}$ is offline-advantageous for $\mathcal{C}$ under $\mathcal{D}$ if (Hardness) $\mathcal{C}$ is not efficiently PAC learnable by any polynomial-time classical learner under $\mathcal{D}$, and (Learnability without deploy-time features) there is a classical learner $\mathcal{A}_{\text{off}}$ that, given $\text{poly}(n, 1/\epsilon, \log(1/\delta))$ samples from $\text{EX}_{\text{ext}}(c, D_n, \mathcal{E})$, outputs an efficiently classically computable $\tilde{h}$ with $\Pr_{x \sim D_n}[h(x) \neq c(x)] \leq \epsilon$ with probability $\geq 1 - \delta$, without computing E_n on new points.

3 Provable Advantages in LUQPI

We now sketch how to establish provable learning advantages within the LUQPI framework, showing that quantum feature extraction can yield exponential separations over classical learners even with such minimal quantum involvement. The construction we present is deliberately contrived, aimed at the cleanest separation proof, while we expect the practical utility of LUQPI to extend well beyond it, as our numerical experiments illustrate. Prior approaches to quantum-enhanced learning either require quantum access at deployment (as in quantum kernels [8] or extreme learning machines [20]), rely on feature maps that process multiple points or observe labels. Earlier discrete-cube-root constructions [7, 11] can be cast as offline feature extraction, but only under contrived distributions or by giving up advantages against non-uniform learners. Our construction is, to our knowledge, the first to achieve a quantum-offline advantage relative to a fixed natural distribution and against non-uniform classical learners—the strongest combination in this landscape.

We work in a cyclic group $G = (G, g, q)$ of order q with generator g , instantiated deterministically via a safe-prime group generator $\text{GroupGen}(1^n)$ (see full paper [1]), so the group is fixed per size. Our construction leverages two classical facts about such groups: the *discrete logarithm* problem (DLOG), which asks to recover z from $g^z \in G$, is believed to be classically hard but is efficiently solvable on a quantum computer via Shor’s algorithm [15]; and the *Decisional Diffie–Hellman* (DDH) assumption, which posits that group elements of the form g^{ab} are computationally indistinguishable from uniformly random ones given g^a and g^b , underpinning the security of ElGamal encryption. We introduce the *ElGamal Encrypted Key (EEK)* concept class, where each concept encrypts its input under a hidden secret key. Building on ElGamal encryption, we encode the secret key bitwise in the exponent: for $y = y_1 \cdots y_n \in \{0, 1\}^n$ and $\iota(y) = \sum_i y_i 2^{n-i} \bmod q$,

$$c_y(h_1, \dots, h_n) = \left(g^{\iota(y)}, \{ h_i^{\iota(y)} \cdot g^{y_i} \}_{i \in [n]} \right), \quad h_i \in G. \quad (1)$$

Each concept is classically easy to evaluate given y , yet recovering y from samples is hard. On the quantum side, the secret key can be extracted from a single example by solving discrete logarithms via Shor’s algorithm, leading to the following theorem.

Theorem 1 (LUQPI learnability, informal). *The feature map $E_n(h_1, \dots, h_n) = (\text{DLOG}_g(h_1), \dots, \text{DLOG}_g(h_n))$, computable quantumly via Shor’s algorithm, recovers the secret key bit-by-bit from a single labelled example, so the EEK class is LUQPI learnable.*

Classically, on the other hand, learning the same class is provably hard under a circular variant of the DDH assumption, as stated in the next theorem.

Theorem 2 (Classical hardness, informal). Assuming the circular DDH assumption holds relative to *GroupGen* (proven secure in the Generic Group Model [10,4]), the EEK class is hard to learn classically under the uniform distribution over G_n .

Combined with Theorem 1, this gives an advantageous quantum-offline LUQPI separation for a fixed natural distribution; under the non-uniform variant of the assumption, the separation holds against non-uniform classical learners. More details and the explicit construction of the prove is shown in [1].

4 LUQPI in Practice: Rydberg Phase Classification

To test LUQPI in a more natural setting, we study ground-state phase detection in a 1D Rydberg atom chain [2,9], which exhibits Disordered, Z_2 , and Z_3 phases. The specific learning task is: given Hamiltonian parameters (R_0/a , Δ/Ω), predict the phase *without* computing the ground state at deployment. The privileged information, available only during training, consists of the expectation values O_{Z_2} and O_{Z_3} of the order parameters for the Z_2 and Z_3 phases, computed from the ground state.

We implement our LUQPI algorithm with SVM+ [17,16], which models the SVM slack variables as a function of the privileged features, and compare it against (i) a classical SVM baseline (no privileged information) and (ii) a transformer-based conditional generative model [19,18] that learns to predict quantum observables explicitly. Ground states are obtained with DMRG simulations and experiments use small training sets (15–100 samples) over 30 randomised splits with 95% confidence intervals. We consider three sampling strategies: uniform (points drawn across the full parameter space) and light- and hard-boundary sampling, which increasingly concentrate samples near phase boundaries.




Fig. 1: Performance of LUQPI (SVM+) against classical baselines on Rydberg phase classification under hard-boundary sampling. **a)** Test accuracy versus training-set size for SVM, SVM+, and the transformer (top), and the corresponding SVM+ accuracy gain over SVM with 95% confidence intervals (bottom); SVM+ achieves the highest accuracy throughout the low-data regime. **b)** Predicted phase diagrams in the (R_b, Δ) plane for a representative run with 30 training points (black circles), comparing ground truth, standard SVM, SVM+, and the transformer; red crosses mark misclassified points. SVM+ reduces misclassifications, with the largest improvements concentrated near phase boundaries.

Under uniform sampling, the transformer performs best while SVM and SVM+ are comparable, as most samples lie far from boundaries where order parameters add little information (not fully shown here but in [1]). The benefit of privileged information emerges under boundary-biased sampling. For instance, in the hard-boundary low-data regime (15–40 samples) SVM+ attains the best overall accuracy, with a 2–4% mean gain over SVM and outperforming the transformer (Fig. 1, a). As the misclassification maps show (Fig. 1, b), these gains concentrate near phase boundaries and in underrepresented phases—precisely where privileged quantum features are most informative.

5 Summary

We introduced Learning Under Quantum Privileged Information (LUQPI), a minimal framework in which a quantum computer only augments the training set, label-independently and point-by-point, while both training and deployment remain fully classical. Even under such a restrictive usage scenario for quantum computers, we prove an exponential quantum-classical separation against non-uniform learners under a fixed natural distribution, and demonstrate how it can be implemented in practical settings. Our work presents a valuable step toward the efficient use of quantum hardware as a source of features that make classically hard learning tasks tractable for purely classical learners.

References

1. Bokov, V., Kohl, L., Schmitt, S., Dunjko, V.: Machine learning with minimal use of quantum computers: Provable advantages in learning under quantum privileged information (luqi) (2026). <https://doi.org/10.48550/ARXIV.2601.22006>, <https://arxiv.org/abs/2601.22006>
2. Browaeys, A., Lahaye, T.: Many-body physics with individually controlled rydberg atoms. *Nature Physics* **16**, 132–142 (2020)
3. Cao, C., Gambetta, F.M., Montanaro, A., Santos, R.A.: Unveiling quantum phase transitions from traps in variational quantum algorithms. *npj Quantum Information* **11**(1) (2025). <https://doi.org/10.1038/s41534-025-01038-5>, <http://dx.doi.org/10.1038/s41534-025-01038-5>
4. Corrigan-Gibbs, H., Kogan, D.: The discrete-logarithm problem with preprocessing. In: *Annual International Conference on the Theory and Applications of Cryptographic Techniques (EUROCRYPT)*. pp. 415–447. Springer (2018)
5. Gyurik, C., Cade, C., Dunjko, V.: Towards quantum advantage via topological data analysis. *Quantum* **6**, 855 (2022)
6. Gyurik, C., Dunjko, V.: On establishing learning separations between classical and quantum machine learning with classical data (2022)
7. Gyurik, C., Dunjko, V.: Exponential separations between classical and quantum learners (2024)
8. Havlíček, V., Córcoles, A.D., Temme, K., Harrow, A.W., Kandala, A., Chow, J.M., Gambetta, J.M.: Supervised learning with quantum-enhanced feature spaces. *Nature* **567**, 209–212 (2019)
9. Huang, H.Y., Kueng, R., Torlai, G., Albert, V.V., Preskill, J.: Provably efficient machine learning for quantum many-body problems. *Science* **377** (2022). <https://doi.org/10.1126/science.abk3333>
10. Ishai, Y., Li, H., Lin, H.: Succinct homomorphic MACs from groups and applications. <https://eprint.iacr.org/2024/2073> (2024), *cryptology ePrint Archive*, Paper 2024/2073
11. Jerbi, S., Gyurik, C., Marshall, S.C., Molteni, R., Dunjko, V.: Shadows of quantum machine learning. *Nature Communications* **15** (2024). <https://doi.org/10.1038/s41467-024-49877-8>
12. Kurkin, A., Shen, K., Pielawa, S., Wang, H., Dunjko, V.: Universality and kernel-adaptive training for classically trained, quantum-deployed generative models (2025). <https://doi.org/10.48550/ARXIV.2510.08476>, <https://arxiv.org/abs/2510.08476>
13. Lloyd, S., Garnerone, S., Zanardi, P.: Quantum algorithms for topological and geometric analysis of data. *Nature Communications* **7** (2016). <https://doi.org/10.1038/ncomms10138>
14. Recio-Armengol, E., Ahmed, S., Bowles, J.: Train on classical, deploy on quantum: scaling generative quantum machine learning to a thousand qubits (2025). <https://doi.org/10.48550/ARXIV.2503.02934>, <https://arxiv.org/abs/2503.02934>
15. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing* **26**, 1484–1509 (1997)
16. Vapnik, V., Izmailov, R.: Learning using privileged information: Similarity control and knowledge transfer. *Journal of Machine Learning Research* **16**, 2023–2049 (2015)
17. Vapnik, V., Vashist, A.: A new learning paradigm: Learning using privileged information. *Neural Networks* **22**, 544–557 (2009)
18. Wang, H., Weber, M., Izaac, J., Lin, C.Y.Y.: Generative quantum states. <https://github.com/PennyLaneAI/generative-quantum-states> (2022)
19. Wang, H., Weber, M., Izaac, J., Lin, C.Y.Y.: Predicting properties of quantum systems with conditional generative models (2022)
20. Xiong, W., Facelli, G., Sahebi, M., Agnel, O., Chotibut, T., Thanasilp, S., Holmes, Z.: On the role of quantum reservoir computing in machine learning. *Quantum Machine Intelligence* **7** (2025). <https://doi.org/10.1007/s42484-025-00239-7>