Single-shot quantum machine learning

Erik Recio-Armengol,[1,2] Jens Eisert,[3,4,5] and Johannes Jakob Meyer[3]

[1] *ICFO-Institut de Ciencies Fotoniques, The Barcelona Institute of Science and Technology, 08860 Castelldefels (Barcelona), Spain*
[2] *Eurecat, Centre Tecnològic de Catalunya, Multimedia Technologies, 08005 Barcelona, Spain*
[3] *Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, 14195 Berlin, Germany*
[4] *Fraunhofer Heinrich Hertz Institute, 10587 Berlin, Germany*
[5] *Helmholtz-Zentrum Berlin für Materialien und Energie, 14109 Berlin, Germany*

(Dated: June 3, 2024)

• **Introduction** Quantum machine learning aims to improve learning methods through the use of quantum computers. If it is to ever realize its potential, many obstacles need to be overcome. A particularly pressing one is the measurement problem that arises because the outputs of quantum learning models are inherently random. As such, many executions of quantum learning models have to be aggregated to obtain an actual prediction. In this work, we analyze when quantum learning models can evade this issue and produce predictions in a near-deterministic way – paving the way to *single-shot quantum machine learning* (see Figure 1). We give a rigorous definition of single-shotness in quantum classifiers and show that the degree to which a quantum learning model is near-deterministic is constrained by the distinguishability of the embedded quantum states used in the model. Opening the black box of the embedding, we show that if the embedding is realized by quantum circuits, a certain depth is necessary for single-shotness to be even possible. We conclude by showing that quantum learning models cannot be single-shot in a generic way and trainable at the same time.

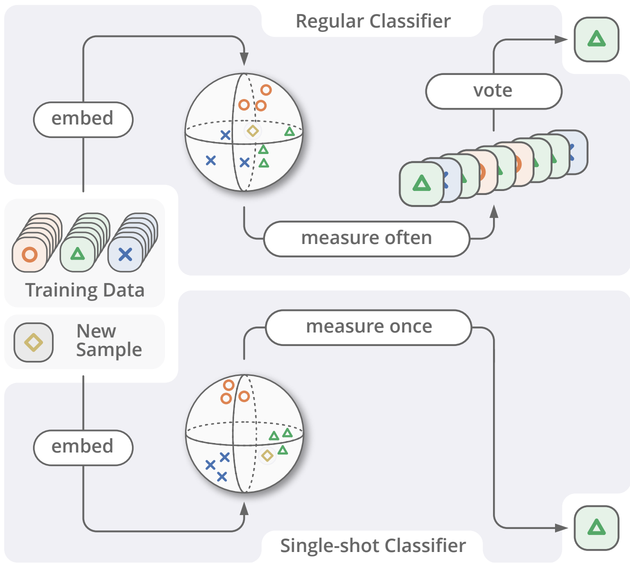

Figure 1. A depiction of the intuitive difference between a regular quantum classifier (top panel) and a single-shot one (bottom panel). In both cases, data is embedded into a quantum system through a quantum feature map. In a regular classifier, the procedure that extracts the label has to be repeated often and then aggregated into a prediction by a majority vote. In a single-shot classifier on the other hand, a single pass through the quantum model is sufficient to extract the label near-deterministically.

• **Quantum Classifiers** A classification task can be formalized by considering data from a data domain $\boldsymbol{x} \in \mathcal{X}$ with labels from a label domain $y \in \mathcal{Y}$ which we assume to be discrete. A quantum classifier is defined as a function that maps classical data to quantum states via a quantum feature map and assigns labels by processing measurement outcomes. Formally, the probability of obtaining the label y upon preparing $\rho(\boldsymbol{x})$ is given by $\mathrm{Tr}[\rho(\boldsymbol{x})\Pi_y]$, where Π_y is a positive operator-valued measure (POVM) element describing the combination of measurement and post-processing on a mathematical level. The predicted label can be formalized as $f(\boldsymbol{x}) = \arg\max_{y'} \mathrm{Tr}[\rho(\boldsymbol{x})\Pi_{y'}]$. The probabilistic nature of quantum classifiers arises because they output random variables rather than deterministic labels, necessitating multiple runs and majority voting to determine the most likely label.

• **Quantum Multi-Hypothesis Testing** Quantum classifiers perform tasks analogous to quantum multi-hypothesis testing, where the objective is to distinguish between multiple quantum states. It comes to little surprise that the performance of quantum classifiers can be analyzed through the lens of hypothesis testing. The Bayesian multi-hypothesis testing error is given by:

$$P_e^*(\{p_j\rho_j\}_{j=1}^r) = \min_{\{\Pi_j\}} \left\{ 1 - \sum_{j=1}^r p_j \mathrm{Tr}[\Pi_j \rho_j] \right\}, \tag{1}$$

where $\{p_j\rho_j\}$ are the quantum states we are tasked to distinguish with associated probabilities. The usefulness of hypothesis testing stems from the fact that we have good lower bounds on the multi-hypothesis testing error. For the binary case, we have the Helstrom-Holevo Theorem [1] that relates the minimal error of binary hypothesis testing $P_e^*(p\rho, (1-p)\sigma)$ to the trace distance of the quantum states.

• **Single-shot Quantum Machine Learning** The concept of single-shot quantum machine learning aims to overcome the measurement problem by ensuring that a single execution of a quantum classifier is sufficient to obtain a prediction with high probability. A classifier is defined as δ-single-shot if the expected probability that the label predicted from infinite shots matches the output label for one shot is at least $1-\delta$:

Definition 1 (Single-shot probabilistic classifier (Bayesian)). *For datapoints $\boldsymbol{x}$ distributed according to a distribution $p(\boldsymbol{x})$ supported on a data domain $\mathcal{X}$, we say a probabilistic classifier f is δ-single shot if*

$$\mathop{\mathbb{E}}_{\boldsymbol{x}\sim p(\boldsymbol{x})} \left\{ \mathbb{P}[\hat{f}(\boldsymbol{x}) = f(\boldsymbol{x})] \right\} \geq 1-\delta, \tag{2}$$

where the probability is taken over the random distribution of datapoints and the inherent randomness of the classifier.

Here $f(x)$ is the majority label assigned by the classifier, and $\hat{f}(x)$ is the random label output by the classifier run only once. Notice that we are not worried with outputting the true labels of the data, but only with the consistency of the classifier, regardless of its correctness. Nevertheless, this property depends on the specific learning problem and the data distribution. In the full paper, we also have a definition that works without considering a distribution on the data (which we might not know) and we study the interplay between these two definitions.

• **Reduction to Multi-Hypothesis Testing** To quantify the limits of the Bayesian single-shot property, we split the parameter space according to the majority labels assigned by the classifier. We formally define the associated set as $\tilde{\mathcal{X}}_{y'} \coloneqq \{\boldsymbol{x} \in \mathcal{X} \mid f(\boldsymbol{x}) = y'\}$.

Theorem 2 (Bayesian error probability lower bound). *Let f be a probabilistic classifier on a data domain $\mathcal{X}$ taking discrete values in $\mathcal{Y}$. For datapoints $\boldsymbol{x}$ distributed according to a distribution $p(\boldsymbol{x})$, we define the average states for the classes assigned by f as*

$$\tilde{p}(y) \coloneqq \int_{\tilde{\mathcal{X}}_y} \mathrm{d}p(\boldsymbol{x}), \quad \tilde{\rho}_y \coloneqq \frac{1}{\tilde{p}(y)} \int_{\tilde{\mathcal{X}}_y} \mathrm{d}p(\boldsymbol{x})\, \rho(\boldsymbol{x}). \tag{3}$$

If the classifier f has the Bayesian single-shot property, the error probability δ has to fulfill

$$\delta \geq P_e^*(\{\tilde{p}_y \tilde{\rho}_y\}_{y\in\mathcal{Y}}). \tag{4}$$

Therefore, the error probability δ is bounded by the optimal error of the corresponding multi-hypothesis testing problem (and, therefore, by the trace distance between the average embedded states).

• **Ultimate Limits of Single-shot Quantum Machine Learning** The feasibility of achieving single-shot properties in quantum classifiers is fundamentally constrained by the distinguishability of the embedded quantum states.

▷ **Noiseless circuits** For noiseless data-reuploading circuits we can find a bound that relates the single-shot error probability to the length of the circuit.

Theorem 3 (Lower bound on the single-shot error probability (Bayesian)). *Let $\hat{f}(\boldsymbol{x})$ be a quantum classifier with the task of classifying r groups of classical data for which we know its probability distribution. The single-shotness error probability is bounded by the worst average distance between two data classes:*

$$\delta \geq \min_{1\leq i\leq r} \max_{1\leq j\neq i\leq r} \frac{\min(p_i,p_j)}{p_i+p_j} \left(1 - \frac{L\Delta d_{avg}^{ij}}{2}\right), \tag{5}$$

where

$$d_{avg}^{ij} \coloneqq (p_i+p_j) \int_{\mathcal{X}_i} \frac{dp(\boldsymbol{x}_i)}{p_i} \int_{\mathcal{X}_j} \frac{dp(\boldsymbol{x}_j)}{p_j} \|\boldsymbol{x}_i - \boldsymbol{x}_j\|_1 \tag{6}$$

is the expected distance of two datapoints sampled according to the data distribution, conditioned on them being from the classes i and j.

Specializing the above result to two classes and setting $p_i, p_j = 1/2$ and rearranging we find that the depth L of the circuit required to achieve a low single-shot error probability is inversely proportional to the average distance between classes in the input data space: $L \geq 2\frac{1-2\delta}{\Delta d_{\mathrm{avg}}}$ where Δ is the maximum spectral spread of the generators used in the circuit, and d_{avg} is the average distance between different classes.

▷ **Noisy Circuits** In practical implementations, quantum circuits are affected by noise, which impacts their ability to produce easily distinguishable quantum states. Considering an error model with local depolarizing noise, we find an upper bound on the trace distance between noisy embedded states that decays exponentially with the number of computational steps. We also provide a third asymptotic bound that improves for shallow circuits. Combining all three bounds we have derived for the noisy case (Figure 2) gives us an upper bound on the achievable distance between the embedded average states per majority class predicted by the quantum classifier and as such allows us to analyze the possibility of single-shot classification by noisy quantum classifiers.

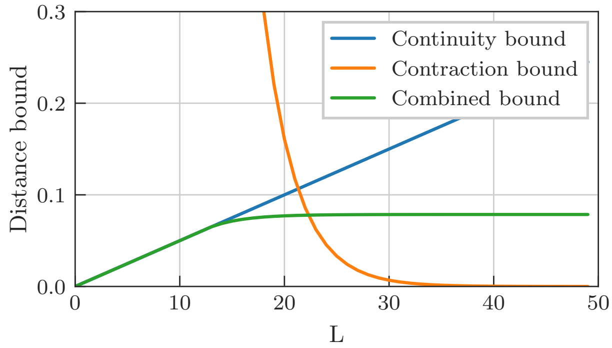

Figure 2. Comparison of the three different bounds where we can identify the three regimes: linear, geometric and exponential decay.

• **Generically accurate single-shot models are hard to learn** Achieving single-shotness requires that quantum states with the same predicted label are sufficiently separated from those with different predicted labels. This separation ensures that the states cluster appropriately, facilitating accurate single-shot predictions. Theoretically, embedding all sufficiently separated inputs into mutually orthogonal directions in Hilbert space can enable single-shot classification for any dataset labeling. This could be achieved with numerous encoding gates interspersed with deep layers of random unitary gates. However, we build on arguments presented in Ref. [2] to show that such an embedding often results in poor generalization performance, essentially as unseen datapoints would be mapped into orthogonal complements, leading to classifiers not outperforming random guessing on new data. This highlights a critical challenge in quantum learning models: balancing expressivity and generalization. Highly expressive models that map inputs to orthogonal directions struggle to generalize to new data.

• **Challenges and Future Directions** Achieving generic single-shot properties in quantum classifiers is challenging due to the inherent trade-off between expressivity and generalization performance. Overly expressive models tend to overfit the training data, resulting in poor generalization to unseen data. Future research should focus on developing training procedures that enforce single-shot properties while maintaining good generalization. Additionally, exploring single-shot quantum machine learning for regression tasks and other problem domains could provide further insights and applications.

• **Conclusion** This study introduces and formalizes the concept of single-shot quantum machine learning, highlighting its potential to mitigate the measurement problem in quantum machine learning. By leveraging connections to quantum multi-hypothesis testing, we derive fundamental limits on the depth of quantum circuits necessary for single-shot classification. Our findings suggest that, while single-shot properties can be achieved under certain conditions, practical implementations must carefully balance expressivity and generalization. Further research is needed to develop robust single-shot quantum learning models and extend these concepts to broader applications.

[1] Carl W. Helstrom. Quantum detection and estimation theory. 1(2):231–252. ISSN 1572-9613. doi:10.1007/BF01007479. URL `https://doi.org/10.1007/BF01007479`.

[2] Hao-Chung Cheng, Min-Hsiu Hsieh, and Ping-Cheng Yeh. The learnability of unknown quantum measurements. URL `http://arxiv.org/abs/1501.00559`.