Inductive Supervised Quantum Learning

Alex Monràs,[1] Gael Sentís,[2,3] and Peter Wittek[4,5]

[1]Física Teòrica: Informació i Fenòmens Quàntics, Universitat Autònoma de Barcelona, ES-08193 Bellaterra (Barcelona), Spain
[2]Departamento de Física Teórica e Historia de la Ciencia, Universidad del País Vasco UPV/EHU, E-48080 Bilbao, Spain
[3]Naturwissenschaftlich-Technische Fakultät, Universität Siegen, 57068 Siegen, Germany
[4]ICFO-The Institute of Photonic Sciences, Castelldefels E-08860 Spain
[5]University of Borås, Borås S-50190 Sweden

(Received 27 September 2016; published 12 May 2017)

In supervised learning, an inductive learning algorithm extracts general rules from observed training instances, then the rules are applied to test instances. We show that this splitting of training and application arises naturally, in the classical setting, from a simple independence requirement with a physical interpretation of being nonsignaling. Thus, two seemingly different definitions of inductive learning happen to coincide. This follows from the properties of classical information that break down in the quantum setup. We prove a quantum de Finetti theorem for quantum channels, which shows that in the quantum case, the equivalence holds in the asymptotic setting, that is, for large numbers of test instances. This reveals a natural analogy between classical learning protocols and their quantum counterparts, justifying a similar treatment, and allowing us to inquire about standard elements in computational learning theory, such as structural risk minimization and sample complexity.

DOI: 10.1103/PhysRevLett.118.190503

Real-world problems often demand optimizing over massive amounts of data. Machine learning algorithms are particularly well suited to deal with such daunting tasks: by mimicking a learning process, data are handled in a tractable way and approximately optimal solutions are inferred. Quantum machine learning, an emergent line of research that introduces quantum resources in learning algorithms [1–5], brings this pragmatic approach to quantum information processing, with a strong emphasis on speedup [6–9]. Quantum mechanics, however, also alters the limitations on what is physically possible in a classical learning setup, thus potentially changing the structure of learning algorithms at a fundamental level and opening a door for increasing performance. In particular, handling quantum data collectively typically allows us to outperform local approaches in many information processing tasks [10–15]. Investigating the potential advantage of using quantum resources in learning algorithms crucially demands to establish the ultimate limits achievable within the framework of quantum machine learning. This Letter tackles the question for general inductive supervised learning scenarios.

In machine learning, we are given a sample of a distribution called training instances [16,17]. The training instances may have further fine structure, and we often think of them as pairs consisting of an input object and a matching output value or label: this is the scenario of supervised machine learning where a classifying function is induced from the training instances and then used to assign labels to a number of unlabeled instances that we call test instances. Not all forms of supervised learning are inductive: transductive learning refers to a problem in which labeled training instances are available, as well as unlabeled instances [18]. The task is to propagate the labels to the unlabeled ones; that is, we do not require inducing a function that we can use infinitely many times. In this case, the geometry of both the labeled and unlabeled instances will influence the outcome. At variance with supervised learning, in which the training occurs in a single step, reinforcement learning algorithms are trained on an instance basis with the possibility of changing the distribution by the subsequent querying, and the quantum generalization of the scenario has already been studied [19].

In this Letter we develop a framework for inductive supervised quantum learning, that is, when training and test instances are given in a quantum form, and we contrast its structure with its classical analog. We first show for the classical case that a natural independence requirement among test instances, i.e., that the learning algorithm be nonsignaling, induces the standard splitting of inductive learning algorithms into a training phase and a test phase. We then prove that the same splitting holds asymptotically in the quantum case, despite having access to coherent collective quantum operations. In other words, we show that, in a fully quantum setting, the following three statements are equivalent in terms of performance in the asymptotic limit: (i) a supervised learning algorithm learns a function which is applied to every test instance; (ii) a supervised learning algorithm satisfies a nonsignaling criterion; (iii) a supervised learning scenario splits into separate training and test phases.

More formally, we derive a de Finetti theorem for nonsignaling quantum channels and use it to prove that the performance of any quantum learning algorithm, under the restriction of being nonsignaling, approaches that of a protocol that first measures the training instances and then infers labels on the test instances, in the limit of many test instances. Our result reveals a natural analogy between classical and quantum learning protocols, justifying a similar treatment that we have been taken for granted. Ultimately, the result provides a solid basis to generalize key concepts in statistical learning theory, such as structural risk minimization [20], to quantum scenarios.

Classical inductive learning.—Consider a supervised learning problem characterized by an unknown joint probability distribution P_{XY}, where X and Y are random variables that model the test data and the label associated with it, respectively. We denote its respective marginals by P_X and P_Y. We are given a finite set of independent, identically distributed unlabelled test instances $\{x_i\}_{i=1}^n$ and a set of correctly labeled examples called the training set. The training set is generated by sampling the distribution P_{XY}, and we model it by the random variable $A = \{(X_1, Y_1), \ldots, (X_m, Y_m)\}$. We are then set to solve the task of assigning a label y_i to each test instance x_i, based on the information contained in the training set A. We define a *learning protocol* that implements this task by a stochastic map $\mathbb{P}(y_1, \ldots, y_n|A, x_1, \ldots, x_n)$.

The natural figure of merit for assessing the performance of a learning protocol is the *expected risk*, defined in terms of the conditional expected risk or average score per test instance $\mathbf{E}[\mathbb{P}|A] = \sum_i \mathbb{E}[s_{y_i,y_i'}\mathbb{P}(y_{1:n}|A, x_{1:n})]$, where y_i' are the true labels, accessible to a referee for evaluation purposes, $s_{j,k} = 1 - \delta_{j,k}$, and we have introduced the short-hand notation $x_{1:n} = \{x_1, \ldots, x_n\}$ (likewise for y and y'). The expectation is in terms of variables $x_{1:n}, y'_{1:n}$ over the distribution P_{XY}, i.e., for a generic function g, $\mathbb{E}[g(x, y, y')] = \sum_{x,y,y'} g(x, y, y')P_{XY}(x, y')$. The expected risk is then defined as the average conditional expected risk over realizations a of the training set, i.e., $\mathbf{E}[\mathbb{P}] = \sum_a p_A(a)\mathbf{E}[\mathbb{P}|a]$.

It is convenient to define the marginal maps of $\mathbb{P}$:

$$\mathbb{P}_i(y_i|A, x_{1:n}) \equiv \sum_{y_{1:i-1}, y_{i+1:n}} \mathbb{P}(y_{1:n}|A, x_{1:n}). \qquad (1)$$

We call a learning protocol *inductive* if it satisfies the condition

$$\begin{aligned}&\mathbb{P}_i(y_i|A, x_{1:i-1}, x_i, x_{i+1:n})\\ &= \mathbb{P}_i(y_i|A, x'_{1:i-1}, x_i, x'_{i+1:n}),\\ &\quad \forall\ x'_{1:i-1}, x'_{i+1:n},\end{aligned} \qquad (2)$$

for all i, namely, that $\mathbb{P}_i(y_i|A, x_{1:n})$ is actually independent of all the X random variables but X_i, for all i. Equation (2) can be interpreted as a *nonsignaling* condition among the test instances as far as the learning protocol is concerned. Note, however, that each marginal map $\mathbb{P}_i$ is still affected by the training set A. This definition encompasses the standard assumption of inductive learning, where a classifier f is extracted from the training set, and only f determines the label to be assigned to each test instance.

In contraposition, consider a transductive learning scenario, where the topology of all of the unlabelled instances can have an impact on the assignment of any of the labels. The independence condition in Eq. (2) is thus violated; hence, a transductive protocol is potentially signaling.

The following lemma pinpoints the feature of classical inductive learning that is relevant for our goal. In the next section we explore its extension to quantum settings.

Lemma 1: For every inductive learning protocol $\mathbb{P}$ that assigns labels y_i to test instances x_i, there exists a set F of classifying functions $f : X \to Y$ and stochastic maps $T(f|A)$, $Q(y|x, f) = \delta_{y,f(x)}$ such that the inductive protocol $\tilde{\mathbb{P}}$,

$$\tilde{\mathbb{P}}(y_{1:n}|A, x_{1:n}) = \sum_f \left(\prod_{i=1}^n Q(y_i|x_i, f)\right) T(f|A),$$

has expected risk $\mathbf{E}[\mathbb{P}|A] = \mathbf{E}[\tilde{\mathbb{P}}|A]$ for all A.

The proof can be found in Ref. [21]. In words, this lemma shows that every conceivable inductive learning protocol can be regarded, with no effect on its performance, as a two-phase operation: a training phase (represented by the stochastic map T), where a classifier f is extracted from the training set A, followed by a test phase (represented by Q), where f is applied to each test instance x_i and output labels y_i are assigned. The key ingredient behind its proof is that the risk is a symmetric function of the joint inputs and outputs, thus randomly permuting the inputs—and its corresponding outputs—does not affect the expected risk. Under this randomization, the resulting protocol $\bar{\mathbb{P}}$ remains nonsignaling, and thus applying the marginal protocol $\bar{\mathbb{P}}_{1|a}(y|x) \coloneqq \bar{\mathbb{P}}_1(y|a, x)$ on each of the test instances will yield the same expected risk as the original protocol $\mathbb{P}$. It is enlightening for our purpose to realize that all test instances are independently acted upon by maps $\bar{\mathbb{P}}_{1|a}$ that use the same sample of the training set a. As we show in the next section, this is the element of the proof that fails to hold in the quantum case.

On a fundamental note, consider the converse of Lemma 1: any learning protocol that splits into a training phase and a test phase satisfies the nonsignaling condition (2), so one can arguably think of the nonsignaling condition as a definitory trait of inductive learning. The advantage of this approach is manyfold. On one hand, it allows one to focus on a much simpler set of features which fully characterize the performance of the protocol. In addition, the training phase can be extended to provide further information relevant to assess, in advance of the test phase, the expected performance of the protocol. This is the case in, e.g., structural risk minimization, where not only a function is chosen but also an estimator of the expected risk itself is provided [20], and confidence intervals are obtained.

Quantum learning.—Quantum information cannot be cloned; hence, the argument supporting Lemma 1 breaks down. However, it is possible to approximately clone a quantum state, and the quality of the clones, will depend on how many copies must be produced, reaching an asymptotic limit in which each copy contains no more information than that which can be obtained by a single quantum measurement on the original system. This idea is reflected in the seminal paper [25], which asserts that asymptotic cloning is equivalent to state estimation succeeded by state preparation. Intuitively, this principle hints at a plausible inductive strategy in a learning scenario where both the training set and all test instances are given as quantum states: perform a quantum measurement $\mathcal{M}$ on the training set A, distribute the measurement outcome across all test instances $B_{1:n}$, and then use it to handle each test instance independently. This approach has the property of being nonsignaling by construction. We will show that any symmetric nonsignaling protocol can be well approximated by this strategy when the number of test instances is large.

In analogy to a classical learning problem, where an unknown probability distribution p_{XY} must be mimicked by attaching appropriate labels to given random variables X, one may consider the most general quantum learning problem as the task of mimicking bipartite quantum states ρ_{XY} by the action of a quantum channel on the marginal ρ_X. The learning protocol can be thought of as a collective quantum channel $\mathbb{Q}$ which takes a training register A and the set of test instance registers $X^{\otimes n}$ as inputs, and yields a corresponding set of output registers $Y^{\otimes n}$ (see Fig. 1 for an illustrative description of the setup).

Definition 1.—A quantum learning protocol for a training set A and n quantum states $\rho_{XY} \in \mathcal{X} \otimes \mathcal{Y}$ is a multipartite quantum channel $\mathbb{Q}\colon \mathcal{A} \otimes \mathcal{X}_{1:n} \to \mathcal{Y}_{1:n}$. A nonsignaling quantum learning protocol is a quantum channel $\mathbb{Q}\colon \mathcal{A} \otimes \mathcal{X}_{1:n} \to \mathcal{Y}_{1:n}$, such that $\mathrm{tr}_{\mathcal{Y}_{1:i-1}\mathcal{Y}_{i+1:n}}[\mathbb{Q}(\rho_{AX_{1:n}})]$ is only a function of $\rho_{AX_{1:\hat{i}:n}} \coloneqq \mathrm{tr}_{\mathcal{X}_{1:i-1}\mathcal{X}_{i+1:n}}[\rho_{AX_{1:n}}]$, $\forall i$.

This approach serves as a good starting point for generalizing several quantum learning problems, both discriminative and generative. In particular, a quantum state classification problem may be expressed as $\rho_{XY} = \sum_y p_y \rho_X^{(y)} \otimes |y\rangle\langle y|$, where register X contains the quantum state, and Y holds a classical label corresponding to the state in X. This naturally encompasses the programmable quantum discriminator [26,27], but admits a much wider class of setups. Another relevant approach is that of quantum state tomography, i.e., where a classical label x is taken as a predictor for certain quantum states $\rho_Y^{(x)}$, thus $\rho_{XY} = \sum_x p_x |x\rangle\langle x| \otimes \rho_Y^{(x)}$. The task of the protocol is to learn from the training set each one of the quantum states and then produce a similar copy for each X instance. More generally, one could consider the task of generating genuine bipartite quantum states ρ_{XY} starting from their reduced states ρ_X [28].

Definition 2.—Given a *risk* observable $S \in \mathcal{Y} \otimes \mathcal{Y}'$, the expected risk of protocol $\mathbb{Q}$ is the expectation value of the

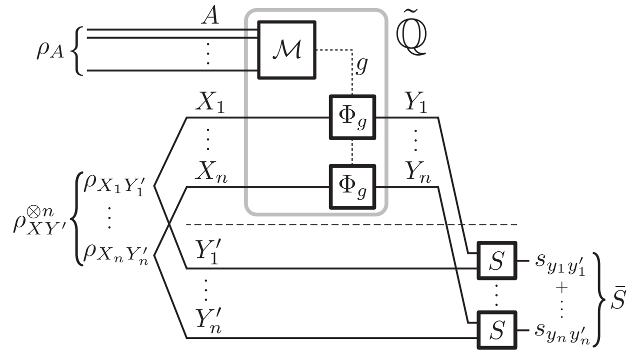

FIG. 1. Diagrammatic representation of a generic quantum learning protocol $\mathbb{Q}$ (gray box), as per Definition 1, which is approximated by $\tilde{\mathbb{Q}}$, given in Eq. (3). Both setups take training and test instances ρ_A and $\rho_{XY'}^{\otimes n}$ as inputs. We distinguish two agents in the diagrams: the performer of the learning protocol or "learner," placed above the dashed horizontal line, and the "referee," placed below. The learner sends the output registers $Y_{1:n}$ of the learning channel to the referee, who contrasts them with the registers $Y'_{1:n}$ and evaluates the average risk of the channel $\bar{S}$ (see Definition 2). This referee plays a role similar to the classical tester in the final phase of quantum-enhanced reinforcement learning [19]. While the most general approach (a collective quantum channel $\mathbb{Q}$) in principle acts globally on all its inputs, its approximation $\tilde{\mathbb{Q}}$ comprises two separate phases: first (training phase) a measurement $\mathcal{M}$ is performed on the training set ρ_A, and second (test phase) the classical information g obtained from the measurement is distributed among all test instances, and corresponding quantum channels Φ_g are applied locally to each one of them.

symmetrized risk observable $\bar{S} \in (\mathcal{Y} \otimes \mathcal{Y}')_{1:n}$ on the output of the channel $\mathbb{Q}$, $\mathbf{E}[\mathbb{Q}] = \mathrm{tr}[\mathbb{Q} \otimes \mathrm{id}_{\mathcal{Y}'_{1:n}}(\rho_A \otimes (\rho_{XY'})^{\otimes n})\bar{S}]$.

For every quantum protocol $\mathbb{Q}$ we can define the symmetrized protocol $\bar{\mathbb{Q}} = (1/n!) \sum_{\sigma \in S_n} \Pi_\sigma^{(\mathcal{Y})\dagger} \circ \mathbb{Q} \circ (\mathrm{id}_{\mathcal{A}} \otimes \Pi_\sigma^{(\mathcal{X})})$. $\bar{\mathbb{Q}}$ is nonsignaling if $\mathbb{Q}$ is. In analogy with the classical case, a nonsignaling quantum channel $\mathbb{Q}\colon \mathcal{A} \otimes \mathcal{X}^{\otimes n} \to \mathcal{Y}^{\otimes n}$ naturally admits a notion of *marginalization*, thereby inducing channels for a reduced number of registers, $k \leq n$, $\mathbb{Q}_k\colon \mathcal{A} \otimes \mathcal{X}^{\otimes k} \to \mathcal{Y}^{\otimes k}$ (Lemma 1 in Ref. [21]). Then, the expected risk can be expressed in terms of $\bar{\mathbb{Q}}_1$, $\mathbf{E}[\mathbb{Q}] = \mathrm{tr}[\bar{\mathbb{Q}}_1 \otimes \mathrm{id}_{\mathcal{Y}'} (\rho_A \otimes \rho_{XY'}) S_{YY'}]$, or the conditional channel $\bar{\mathbb{Q}}_{1|\rho_A}\colon \rho_X \mapsto \bar{\mathbb{Q}}_1(\rho_A \otimes \rho_X)$.

It is clear that the line of reasoning so far is a simple reformulation of the ideas involved in the classical arguments. If one could implement the protocol $\bar{\mathbb{Q}}_1$ on each of the test instances then one could perform with an average performance $\mathbf{E}[\mathbb{Q}]$. At this point, however, we encounter the fundamental roadblock that motivates this work. The map $\rho_A \otimes \rho_{X_1} \otimes \cdots \otimes \rho_{X_n} \mapsto \bar{\mathbb{Q}}_{1|\rho_A}(\rho_{X_1}) \otimes \cdots \otimes \bar{\mathbb{Q}}_{1|\rho_A}(\rho_{X_n})$ is nonlinear in ρ_A, so it does not reflect a physically realizable transformation. This reflects the nonclonable nature of quantum information [29,30]: it is the impossibility of cloning the training set that prevents the simultaneous application of the map $\bar{\mathbb{Q}}_{1|\rho_A}$ on the n test instances. Therefore, a generic quantum channel $\mathbb{Q}: \mathcal{A} \otimes \mathcal{X}_{1:n} \to \mathcal{Y}_{1:n}$ that distributes symmetrically the system A across n identical parties X, can, at best, perform some sort of approximate cloning, which then is acted upon independently and symmetrically. Since, asymptotically, this cloning operation becomes a measure-and-prepare process, we can reduce it to a measurement on the training set, and consider the preparation of the *clones* as part of the task to be performed on each test instance. This argument can be made formal as a de Finetti-type theorem for quantum channels [31]. To conclude this section, we make our statement rigorous.

Theorem 1: Let $\mathbb{Q}: \mathcal{A} \otimes \mathcal{X}_{1:n} \to \mathcal{Y}_{1:n}$ be a nonsignaling quantum channel, and let $S \in \mathcal{Y} \otimes \mathcal{Y}'$ be a local operator. Then, there exists a POVM $\mathcal{M}(dg)$ on $\mathcal{A}$ and a set of quantum channels $\Phi_g: \mathcal{X} \to \mathcal{Y}$ such that the quantum channel $\tilde{\mathbb{Q}}$,

$$\tilde{\mathbb{Q}} = \int \hat{\mathcal{M}}(dg) \otimes \Phi_g^{\otimes n}, \tag{3}$$

satisfies $|\mathbf{E}[\mathbb{Q}] - \mathbf{E}[\tilde{\mathbb{Q}}]| \leq \kappa/n^{1/6} + [O(1/n^{1/3})]$, where the coefficient κ depends on the dimensions of the spaces A, X, and Y.

The main ingredient behind this theorem is the quantum de Finetti theorem for quantum states, which can be found in Ref. [23]. We refer the reader to Ref. [21] for a detailed derivation.

Theorem 1 shows that, for any local operator S, its symmetrized expectation under the action of a nonsignaling quantum channel $\mathbb{Q}$ can be approximated by a one-way quantum channel $\tilde{\mathbb{Q}}$ of local operations and classical communication (LOCC). This channel amounts to performing a measurement $\mathcal{M}(dg)$ yielding outcome g over the training set, and applying simultaneously Φ_g on each of the test instances (see Fig. 1). The resulting performance of both protocols, as measured by their expected risk, converge to each other as n tends to infinity.

Discussion.—The main result reported in this Letter, Theorem 1, is a natural consequence of the symmetry implicit in the problem. Given the fact that the performance on a multiple-instance inductive learning task is symmetric under simultaneous exchange of the test-answer pairs, a randomized permutation of the test instances will yield the same average performance. Therefore, each protocol performs equally well as its randomized permutation protocol. We have used this symmetry and the fact that the quantum information contained in the training set cannot be perfectly distributed over an arbitrarily large number of parties, to show that any such protocol must, effectively, be well approximated by first performing a measurement over the training set, and then distributing the outcome.

Previous works have already dwelt on this issue, that is, the contrast between coherent quantum operations and separate training and test phases for learning tasks, in various specific scenarios. Examples are a quantum pattern matching algorithm [33], quantum learning of unitary operations [34], and quantum learning for state classification [14,27]. It is worth stressing that, whereas the results so far have been case specific, we approach the problem from a very general standpoint, allowing us to discuss the broad class of inductive quantum learning protocols within a common framework.

The simplification of general quantum protocols to schemes that use LOCC has several relevant implications. As quantum information technologies advance, coherent collective manipulation of quantum information will become accessible on a practical scale. Nevertheless, the demonstration of a scalable, general-purpose, quantum computer is still beyond the foreseeable future. For this reason, reducing collective approaches to simpler, local ones is of utmost importance. With the result reported in this Letter, the degree of coherence required for implementing several inductive quantum learning protocols is greatly reduced, from requiring joint coherent manipulation of both the training set and all test instances, to only the training set.

Outlook.—Designing quantum algorithms to learn from quantum information poses a serious challenge. Analytical results are scarce, and numerical computations quickly become intractable. A prominent example is quantum state discrimination, which has no known closed-form solution in general scenarios [35], and only highly symmetric cases are exactly solvable [36]. Reducing a generic quantum learning protocol to a single-instance one-way LOCC protocol greatly simplifies the task. We expect that our result will allow us to derive performance bounds for a variety of relevant quantum learning tasks. Also, our quantitative approximation bounds allow for single-copy algorithms to be used as benchmarks for coherent multi-instance ones.

Another benefit of this reduction is the ability to access, without disturbance, the state of the learner in between the training and test phases. This information is essential for several machine learning tasks. For structural risk minimization [20], one uses an estimate of the expected risk, produced by evaluating the performance of a given classifier on the training set. In the quantum setup, this approach is not directly applicable. As the training set can only be accessed once, one can either extract information to determine the best classifier, or to assess the performance of one given classifier. However, both tasks will generally be incompatible. Therefore, a "quantum black box"—e.g., a fully quantum processor that takes all the inputs (training and tests)—will, despite being the most general approach, provide only the required answers. It is unclear how one can adapt a generic quantum black box to provide an assessment of its own performance. Our result, nevertheless, opens the door to assessing the performance of any classifier by suitably processing the intermediate measurement outcome g. We expect the result reported here will shed light on the potential and limitations of learning from quantum sources, and ultimately serve as a starting ground for developing a fully quantum theory of risk bounds in statistical learning.

A few comments on the degree of generality of our result are in order. The convergence rate of our approximation is potentially not tight, and we expect better bounds to be achievable. For simplicity, the approach presented here uses the operator form of Chebyshev inequality (Lemma 5 in Ref. [21]), which ultimately hinders us from obtaining a better bound. We expect a more detailed study will yield better approximations. More importantly, our result can be extended in various ways. A potentially very relevant practical problem is to learn quantum operations rather than *states*. This, however, can be easily addressed within the Choi matrix formalism. A related result for learning quantum unitary operations already shows the same splitting reported here [34]. Indeed, the formalism of quantum combs [37] provides the theoretical framework for this extension, but, essentially, the most general such process will also be described by a suitable multipartite quantum system $\omega_{AB_{1:n}}$, where A will now consist of input and output ports, and the maps Φ_g will be potential implementations of the learned operations.

The authors acknowledge useful discussions with Andreas Winter and Giulio Chiribella. We are grateful to the anonymous referees who provided extensive feedback that greatly improved the Letter. A. M. is supported by the ERC (Advanced Grant IRQUAT, Project No. ERC-267386). G. S. is supported by the Spanish MINECO (project FIS2015-67161-P) and the ERC (Starting Grant No. 258647/GEDENTQOPT). P. W. acknowledges financial support from the ERC (Consolidator Grant QITBOX), MINECO (Severo Ochoa Grant SEV-2015-0522 and FOQUS), Generalitat de Catalunya (SGR 875), and Fundació Privada Cellex.

[1] P. Wittek, *Quantum Machine Learning: What Quantum Computing Means to Data Mining* (Academic Press, New York, NY, 2014).
[2] M. Schuld, I. Sinayskiy, and F. Petruccione, Contemp. Phys. **56**, 172 (2015).
[3] J. Adcock, E. Allen, M. Day, S. Frick, J. Hinchliff, M. Johnson, S. Morley-Short, S. Pallister, A. Price, and S. Stanisic, arXiv:1512.02900.
[4] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, arXiv:1611.09347.
[5] S. Arunachalam and R. de Wolf, arXiv:1701.06806.
[6] E. Aïmeur, G. Brassard, and S. Gambs, Mach. Learn. **90**, 261 (2013).
[7] P. Rebentrost, M. Mohseni, and S. Lloyd, Phys. Rev. Lett. **113**, 130503 (2014).
[8] X.-D. Cai, D. Wu, Z.-E. Su, M.-C. Chen, X.-L. Wang, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan, Phys. Rev. Lett. **114**, 110504 (2015).
[9] S. Lloyd, S. Garnerone, and P. Zanardi, Nat. Commun. **7**, 10138 (2016).
[10] C. H. Bennett, D. P. DiVincenzo, C. A. Fuchs, T. Mor, E. Rains, P. W. Shor, J. A. Smolin, and W. K. Wootters, Phys. Rev. A **59**, 1070 (1999).
[11] V. Giovannetti, S. Lloyd, and L. Maccone, Nature (London) **412**, 417 (2001).
[12] V. Giovannetti, S. Lloyd, and L. Maccone, Science **306**, 1330 (2004).
[13] J. Niset, A. Acín, U. L. Andersen, N. J. Cerf, R. García-Patrón, M. Navascués, and M. Sabuncu, Phys. Rev. Lett. **98**, 260404 (2007).
[14] G. Sentís, M. Guţă, and G. Adesso, EPJ Quantum Techno. **2**, 17 (2015).
[15] G. Sentís, E. Bagan, J. Calsamiglia, G. Chiribella, and R. Muñoz-Tapia, Phys. Rev. Lett. **117**, 150502 (2016).
[16] L. Devroye, L. Györfi, and G. Lugosi, *A Probabilistic Theory of Pattern Recognition* (Springer, New York, 1996).
[17] T. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed. (Springer, New York, 2008).
[18] A. Gammerman, V. Vovk, and V. Vapnik, *Proceedings of the UAI-98, 14th Conference on Uncertainty in Artificial Intelligence* (Morgan Kaufmann Publishers Inc., San Francisco, CA, 1998), p. 148.
[19] V. Dunjko, J. M. Taylor, and H. J. Briegel, Phys. Rev. Lett. **117**, 130501 (2016).
[20] V. Vapnik, *The Nature of Statistical Learning Theory* (Springer, New York, 1995).
[21] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.118.190503 for detailed proofs of Lemma 1 and Theorem 1, which includes Refs. [22–24].
[22] I. Bengtsson and K. Życzkowski, *Geometry of Quantum States: An Introduction to Quantum Entanglement* (Cambridge University Press, Cambridge, England, 2006).
[23] M. Christandl, R. König, G. Mitchison, and R. Renner, Commun. Math. Phys. **273**, 473 (2007).
[24] M. Christandl and B. Toner, J. Math. Phys. (N.Y.) **50**, 042104 (2009).
[25] J. Bae and A. Acín, Phys. Rev. Lett. **97**, 030402 (2006).
[26] G. Sentís, E. Bagan, J. Calsamiglia, and R. Muñoz-Tapia, Phys. Rev. A **82**, 042312 (2010).
[27] G. Sentís, J. Calsamiglia, R. Muñoz-Tapia, and E. Bagan, Sci. Rep. **2**, 708 (2012).
[28] All these tasks are meant within the framework described in Fig. 1, and therefore their performance is assessed only through registers $\mathcal{Y}$ and $\mathcal{Y}'$. In the fully quantum case of generating bipartite states ρ_{XY} from ρ_X, this means that a risk observable would measure the dissimilarity between the states $\mathbb{Q}(\mathrm{tr}_{Y'}\rho_{XY'})$ and $\mathrm{tr}_X\rho_{XY'}$, instead of comparing (perhaps more naturally in other contexts) the produced bipartite state with a copy of the target bipartite state saved as reference.
[29] D. Dieks, Phys. Lett. A **92**, 271 (1982).
[30] W. K. Wootters and W. H. Zurek, Nature (London) **299**, 802 (1982).
[31] We note that a de Finetti theorem for fully symmetric quantum channels can be found in the literature [32].
[32] O. Fawzi and R. Renner, Commun. Math. Phys. **340**, 575 (2015).
[33] M. Sasaki and A. Carlini, Phys. Rev. A **66**, 022303 (2002).

[34] A. Bisio, G. Chiribella, G. M. D'Ariano, S. Facchini, and P. Perinotti, Phys. Rev. A **81**, 032324 (2010).
[35] K. Nakahira, K. Kato, and T. S. Usuda, Phys. Rev. A **91**, 052304 (2015).
[36] G. Chiribella, G. M. D'Ariano, P. Perinotti, and M. F. Sacchi, Phys. Rev. A **70**, 062105 (2004).
[37] G. Chiribella, G. M. D'Ariano, and P. Perinotti, Phys. Rev. Lett. **101**, 060401 (2008).