

Potential and limitations of quantum extreme learning machines

L. Innocenti ¹✉, S. Lorenzo ¹, I. Palmisano ², A. Ferraro ^{2,3}, M. Paternostro ² & G. M. Palma ^{1,4}

Quantum extreme learning machines (QELMs) aim to efficiently post-process the outcome of fixed — generally uncalibrated — quantum devices to solve tasks such as the estimation of the properties of quantum states. The characterisation of their potential and limitations, which is currently lacking, will enable the full deployment of such approaches to problems of system identification, device performance optimization, and state or process reconstruction. We present a framework to model QELMs, showing that they can be concisely described via single effective measurements, and provide an explicit characterisation of the information exactly retrievable with such protocols. We furthermore find a close analogy between the training process of QELMs and that of reconstructing the effective measurement characterising the given device. Our analysis paves the way to a more thorough understanding of the capabilities and limitations of QELMs, and has the potential to become a powerful measurement paradigm for quantum state estimation that is more resilient to noise and imperfections.

¹Università degli Studi di Palermo, Dipartimento di Fisica e Chimica - Emilio Segrè, via Archirafi 36, I-90123 Palermo, Italy. ²Centre for Theoretical Atomic, Molecular, and Optical Physics, School of Mathematics and Physics, Queen's University Belfast, Belfast BT7 1NN, UK. ³Quantum Technology Lab, Dipartimento di Fisica Aldo Pontremoli, Università degli Studi di Milano, I-20133 Milano, Italy. ⁴NEST, Istituto Nanoscienze-CNR, Piazza S. Silvestro 12, 56127 Pisa, Italy. ✉email: luca.innocenti@unipa.it

Extreme learning machines (ELMs)^{1–3} and reservoir computers (RC)^{4–9} are computational paradigms that leverage fixed, nonlinear dynamics to efficiently extract information from a given dataset. In the classical context, these schemes rely on evolving input data through some nonlinear mapping — typically recurrent neural networks with fixed weights — which augment the dimensionality of the data, easing the extraction of the properties of interest. The main discriminator between RCs and ELMs is whether the reservoir being used can deploy an internal memory. More precisely, RCs hold memory of the inputs seen at previous iterations, making them suitable for temporal data processing⁵. ELMs instead use memoryless reservoirs. Although this makes the training of ELMs easier, it also makes them unsuitable for temporal data processing.

Quantum counterparts to ELMs and RCs — which we will refer to as QELMs and QRCs, respectively — have recently attracted significant interest due to their potential to process quantum information^{10–18,18–23}. Reviews of the state of the art in this context can be found in refs. ^{24–26}, while a study of QRC schemes for the implementation of nonlinear input-output maps with memory on NISQ devices has recently been presented²⁷.

To date, and to the best of our knowledge, a general characterisation of the class of tasks that can be accomplished through QELM-like schemes for the classification, processing or extraction of information encoded in quantum states is lacking. This significantly limits the systematic deployment of such approaches to the issues of quantum-system and quantum-state characterization or validation, which are crucial steps to perform towards the upscaling of quantum technologies and the achievement of the fault-tolerant quantum information processing paradigm.

In this paper, we show that the problem of reconstructing features of a quantum state via an ELM-like setup can be viewed as a linear regression task on the measurement probabilities produced by a suitable positive operator-valued measurement (POVM)^{28,29}. The key observation is that the probability distribution corresponding to an arbitrary measurement of a quantum state is linear in the input density matrix³⁰. This is a fundamental departure from classical ELMs: whereas in the latter case the reservoir is an intrinsically nonlinear operation, the same cannot be said about a quantum reservoir. The latter can always be modelled as a map that linearly processes the input density matrix. In turn, this allows us to identify crucial constraints on the properties that QELM setups can be trained to retrieve. While the learning of classical input information that is nonlinearly encoded in the states^{19,31} is certainly not precluded, our study clarifies how the only possible source of nonlinearity must come from the encoding itself rather than the reservoir dynamics.

We then show that the intrinsic uncertainty arising from the sampling noise on estimated measurement probabilities dramatically affects the performances of any property-reconstruction protocol based on QELMs. This pinpoints a significant fundamental constraint — of strong experimental relevance — to the performance of such schemes. The number of measurement outcomes is also shown to play an important role, affecting the well-conditioning of the associated regression problem, and thus the numerical stability of any estimate. More generally, we show that the efficiency of QELMs is directly tied to the effective POVM summarising both evolution and measurement. This puts the spotlight on the properties of this effective POVM, and on how these are the ones directly affecting performances.

By addressing fundamental features of significant practical repercussions, our study allows us to shape the contours of the class of tasks that can be successfully tackled with QELM architectures, and contributes to the investigation of property-reconstruction protocols, assisted by artificial intelligence, which

is raising growing attention from the quantum-technology community.

Results and discussion

We start by reviewing the basic features of classical ELMs and QELMs, present a general way to model QELMs, and characterise their predictive power in various scenarios.

Notation and theoretical background. An ELM^{1,3} is a supervised machine learning protocol which, given a training dataset $\{(\mathbf{x}_k^{\text{tr}}, \mathbf{y}_k^{\text{tr}})\}_{k=1}^{M_{\text{tr}}} \subset \mathbb{R}^n \times \mathbb{R}^m$, is tasked with finding a target function $f_{\text{target}} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ such that, for each k , $f_{\text{target}}(\mathbf{x}_k^{\text{test}}) \simeq \mathbf{y}_k^{\text{test}}$ with a sufficiently good approximation for previously unseen datapoints $\{(\mathbf{x}_k^{\text{test}}, \mathbf{y}_k^{\text{test}})\}_{k=1}^{M_{\text{test}}}$. As most machine learning algorithms, ELMs are characterised by their model, that is, the way the input-output functional relation is parametrised. For ELMs, the model is a function of the form $\mathbf{x} \mapsto Wf(\mathbf{x})$ with f a fixed — generally nonlinear — function implementing the reservoir dynamics, and W a linear mapping applied to the output of f . The function f is not trained, but rather fixed beforehand, and can for example be implemented as a neural network with fixed random weights. The training algorithm optimises the parameters defining W in order to minimise some distance function — often the standard Euclidean distance — between $Wf(\mathbf{x}_k^{\text{tr}})$ and \mathbf{y}_k^{tr} . As a classical example, one can think of a supervised learning task where \mathbf{x}_k^{tr} are images representing handwritten digits, and \mathbf{y}_k^{tr} the digits the images represent. In this example, n would be the number of pixels in each image, and the goal of the algorithm would be to use the training dataset of labelled images $\{(\mathbf{x}_k^{\text{tr}}, \mathbf{y}_k^{\text{tr}})\}_{k=1}^{M_{\text{tr}}}$ to find the W such that, for all new images \mathbf{x}_k , $Wf(\mathbf{x}_k)$ is the correct digit drawn in \mathbf{x}_k .

The standard way to quantise ELMs is to replace the map f with some quantum dynamics followed by a measurement. To maintain full generality, we consider a completely positive trace-preserving quantum map Λ — which we refer to as a *quantum channel* — followed by a POVM $\{\mu_b : b \in \Sigma\}$, where Σ is the set of possible measurement outcomes³⁰. In the context of QELMs, the training dataset has the form $\{(\rho_k^{\text{tr}}, \mathbf{y}_k^{\text{tr}})\}_{k=1}^{M_{\text{tr}}}$, with ρ_k^{tr} an input state and \mathbf{y}_k^{tr} the output vector that the QELM should associate to ρ_k^{tr} . More precisely, the goal of the training is to find a linear operation W such that

$$\sum_{b \in \Sigma} W_{ab} \text{Tr}(\mu_b \Lambda(\rho_k^{\text{tr}})) \simeq (\mathbf{y}_k^{\text{tr}})_a, \quad (1)$$

with $a = 1 \dots m$ and $k = 1 \dots M_{\text{tr}}$, and with W_{ab} the matrix elements of W . It is also possible to use QELMs as a way to process classical information exploiting complex quantum dynamics. In this case, the training dataset should be considered as a set of the form $\{(\mathbf{s}_k^{\text{tr}}, \mathbf{y}_k^{\text{tr}})\}_k$, in direct analogy with the classical case, where now \mathbf{s}_k are classical vectors suitably encoded in the input quantum states ρ_s . The difference with the classical setup, in this case, is entirely in the specific form of the function mapping inputs to outputs. The capabilities of QELM/QRCs to process classical data depends crucially on the nonlinearity of the encoding $\mathbf{s} \mapsto \rho_s$, as discussed in refs. ^{19,31} [cf. Fig. 1 for a schematic overview of the distinction between ELM and QELM protocols]. We will focus here on the former point of view to derive results that are independent of the specific forms of classical encodings ρ_s , and useful when the goal is to probe property of the input states.

The “classical reservoir function” $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ becomes, in the quantum case, the map

$$\mathbf{P}_{\Lambda, \mu} : \rho \mapsto (\text{Tr}[\mu_b \Lambda(\rho)])_{b=1}^{|\Sigma|} \in \mathbb{R}^{|\Sigma|}, \quad (2)$$

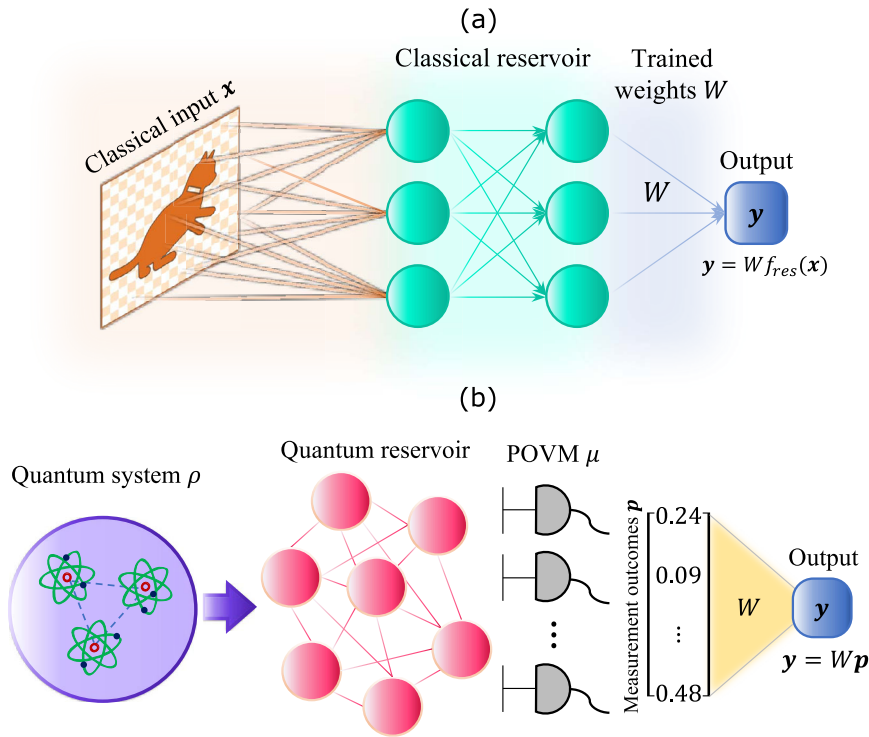


Fig. 1 Schematic overview of ELMs and QELMs. a Classical ELM setup; **b** QELM setup for classical information processing.

Table 1 Summary of the differences between classical ELMs and QELMs.

	ELM	QELM
Training data	$\{(\mathbf{x}_k, \mathbf{y}_k)\}_k$	$\{(\rho_k, \mathbf{y}_k)\}_k$
Model to train	$\mathbf{x} \mapsto Wf(\mathbf{x})$	$\rho \mapsto W\mathbf{p}_{\Lambda, \mu}(\rho)$
Parameters to train	W	W
Cost function	$\ \mathbf{y}_k - Wf(\mathbf{x}_k)\ _2$	$\ \mathbf{y}_k - W\mathbf{p}_{\Lambda, \mu}(\rho_k)\ _2$

These two schemes differ in the type of input states ρ_k fed to the reservoir, and in how the reservoir map itself is implemented: in the classical case, is some nonlinear function often implemented via fixed-weights neural network architectures, whereas in the quantum case it is a quantum channel followed by some measurement.

which sends each input state to the vector of outcome probabilities corresponding to a channel Λ and measurement μ (here $|\Sigma|$ is the dimension of the set of measurement outcomes). Finally, the trained model for QELMs consists of a linear function W applied to the vector of outcome probabilities. This means that, during training, the algorithm optimises the parameters W so as to minimise the distance between $W\mathbf{p}_{\Lambda, \mu}(\rho_k^{\text{tr}})$ and \mathbf{y}_k^{tr} , for all the states and target vectors in the training dataset. In Table 1 we provide a schematic breakdown of the differences between ELMs and QELMs.

In the most general case, the channel Λ is physically implemented by making ρ interact with some reservoir state η and then tracing out some degrees of freedom from the output space. This scenario can be modelled as a channel $\Phi \in C(\mathcal{H}_S \otimes \mathcal{H}_E, \mathcal{H}_E)$ sending states in $\mathcal{H}_S \otimes \mathcal{H}_E$ into states in \mathcal{H}_E , where \mathcal{H}_S and \mathcal{H}_E are the Hilbert spaces of input and reservoir states, respectively, and $C(\mathcal{X}, \mathcal{Y})$ denotes the set of quantum channels sending states in \mathcal{X} to states in \mathcal{Y} . For notational clarity, we will distinguish between the two channels $\Lambda_\eta \in C(\mathcal{H}_S, \mathcal{H}_E)$ and $\mathcal{E}_\rho \in C(\mathcal{H}_E, \mathcal{H}_E)$, defined from Φ as $\Lambda_\eta(\rho) = \mathcal{E}_\rho(\eta) = \Phi(\rho \otimes \eta)$, where η and ρ are states in \mathcal{H}_E and \mathcal{H}_S , respectively. Note that describing the channel as $\Phi \in C(\mathcal{H}_S \otimes \mathcal{H}_E, \mathcal{H}_E)$, means, in

particular, that we assume the output space to have the same dimension as the input reservoir space. One could easily lift this restriction by considering measurements performed on the full space $\mathcal{H}_S \otimes \mathcal{H}_E$, nonetheless we stick to it as it eases our notation.

In the context of open quantum systems, dynamics through a reservoir are often described through channels acting on the reservoir itself, parametrised by the input state. When adopting this point of view, the channel \mathcal{E}_ρ is the one of more direct interest. This is useful for example when studying the memory capabilities of Φ . On the other hand, when one is interested in the retrievability of information encoded in ρ , the linearity of Λ_η is of more direct relevance.

Observable achievability. An observation central to our results is that the mapping from states to probabilities is, regardless of any detail of the dynamics, unavoidably linear

$$\mathbf{p}_{\Lambda, \mu}(\alpha X + \beta Y) = \alpha \mathbf{p}_{\Lambda, \mu}(X) + \beta \mathbf{p}_{\Lambda, \mu}(Y), \quad (3)$$

for any pair of linear maps X, Y and scalars $\alpha, \beta \in \mathbb{C}$. Furthermore, $\mathbf{p}_{\Lambda, \mu}(\rho)$ can be interpreted as a direct measure on the state ρ — that is, the overall process of measuring after an evolution Λ can be reframed as an effective measurement performed directly on ρ . Explicitly, this follows from

$$(\mathbf{p}_{\Lambda, \mu}(\rho))_b = \text{Tr}[\mu_b \Lambda(\rho)] = \text{Tr}[\Lambda^\dagger(\mu_b) \rho] = \text{Tr}[\tilde{\mu}_b \rho], \quad (4)$$

where Λ^\dagger is the adjoint of Λ , and $\tilde{\mu}_b$ denotes said effective measurement which, performed on ρ , reproduces the same measurement outcomes obtained measuring μ_b on $\Lambda(\rho)$. One can equivalently view $\Lambda^\dagger(\mu_b)$ as describing the underlying evolution in the Heisenberg picture. Because the measurement probabilities ultimately depend on the effective POVM $\tilde{\mu}$, we will use the shorthand notation $\mathbf{p}_{\tilde{\mu}} \equiv \mathbf{p}_{\Lambda, \mu}$ when $\tilde{\mu} = \Lambda^\dagger(\mu)$.

A defining feature of QELMs is the restriction to linear post-processing of the measurement probabilities, which has significant implications for their information processing capabilities. To see this, note that applying the linear function W to $\mathbf{p}_{\tilde{\mu}}(\rho)$

produces a vector $\mathbf{y} \equiv W\mathbf{p}_{\tilde{\mu}}(\rho)$, with components

$$y_k = \sum_{b \in \Sigma} (\mathbf{p}_{\tilde{\mu}}(\rho))_b W_{kb} = \text{Tr} \left[\left(\sum_{b \in \Sigma} W_{kb} \tilde{\mu}_b \right) \rho \right]. \quad (5)$$

In other words, any vector \mathbf{y} obtainable via linear post-processing of measurement probabilities has the form $y_k = \text{Tr}(\tilde{\mathcal{O}}_k \rho) \equiv \langle \tilde{\mathcal{O}}_k, \rho \rangle$ for some observable $\tilde{\mathcal{O}}_k$ that is a linear combination of the effective POVM elements. Here and in the following we use the notation $\langle \cdot, \cdot \rangle$ to highlight that expressions of the form $\text{Tr}(A^\dagger B)$ can be interpreted as an inner product between the matrices. It follows that a QELM can learn to retrieve the expectation value of an observable \mathcal{O} iff

$$\mathcal{O} \in \text{span}_{\mathbb{R}}(\{\tilde{\mu}_b : b \in \Sigma\}), \quad (6)$$

that is, if and only if \mathcal{O} can be written as a real linear combination of operators $\tilde{\mu}_b$. It is worth noting that, in this context, we operate under the assumption that $\tilde{\mu}$ — and thus Λ and μ — is known, and therefore the condition is readily verifiable. In particular, a QELM can reproduce the expectation value of arbitrary observables iff $\tilde{\mu}$ is informationally complete (that is, $\text{iff} \{\tilde{\mu}_b : b \in \Sigma\}$ spans the corresponding space of Hermitian operators). Nonetheless, as will be further discussed later, the training procedure does not require knowledge of $\tilde{\mu}$ as it can be seen as a way to estimate the effective measurement $\tilde{\mu}$ itself.

Reconstruction method. Even if we can now readily assess whether a target observable can be retrieved from the information provided in a given QELM setup, the question remains on how exactly this would be done. To fix the ideas, consider a scenario with a single target observable \mathcal{O} , and the effective POVM is some $\tilde{\mu}$ with $|\tilde{\mu}|$ the number of possible outcomes. The problem is thus finding some W — which will be, in this case, a row vector — such that

$$\langle \mathcal{O}, \rho \rangle = W \langle \tilde{\mu}, \rho \rangle \quad (7)$$

for all the elements of the training dataset, which has in this case the form

$$\{(\langle \mathcal{O}, \rho \rangle, \langle \tilde{\mu}, \rho \rangle) : \rho \in \text{TrainingDS}\}, \quad (8)$$

where TrainingDS is the set of states used to generate the training dataset. A convenient way to write this condition is then

$$\langle \mathcal{O}, \rho^{\text{tr}} \rangle = W \langle \tilde{\mu}, \rho^{\text{tr}} \rangle, \quad (9)$$

denoting with ρ^{tr} the vector whose elements are all the training states, with $\langle \mathcal{O}, \rho^{\text{tr}} \rangle$ the vector of expectation values $\langle \mathcal{O}, \rho_k^{\text{tr}} \rangle$, and with $\langle \tilde{\mu}, \rho^{\text{tr}} \rangle$ the matrix with components $\langle \tilde{\mu}_b, \rho_k^{\text{tr}} \rangle$. Equation (9), as a condition for W , is a standard linear regression problem. It is however worth remarking a departure of our task from standard linear regressions: we are not interested in finding any “true value” of W , but rather in finding *some* W which gives the best performances on the test dataset. That means, in particular, that the existence of multiple optimal solutions for W is not an issue.

In the context of QELM, the effective measurement $\tilde{\mu}$ — and thus the matrix $\langle \tilde{\mu}, \sigma \rangle$ — is not known a priori. Instead, during the training phase, only the probabilities $\langle \tilde{\mu}, \rho^{\text{tr}} \rangle$ and expectation values $\langle \mathcal{O}, \rho^{\text{tr}} \rangle$ are given. The task is to solve the corresponding linear system

$$\langle \mathcal{O}, \rho^{\text{tr}} \rangle = W \langle \tilde{\mu}, \rho^{\text{tr}} \rangle \quad (10)$$

for W . Even though without knowing \mathcal{O} and $\tilde{\mu}$ it is not possible to determine a priori the feasibility of the task, if the accuracies during the training phase are sufficiently high one can reasonably expect the condition to be fulfilled. If, on the other hand, the accuracies saturate to a non-optimal amount while increasing the sampling statistics, we can now determine the reason to be \mathcal{O} not being writable as linear combinations of $\tilde{\mu}$.

A standard way to solve Eq. (10) is via the pseudoinverse

$$W = \langle \mathcal{O}, \rho^{\text{tr}} \rangle \langle \tilde{\mu}, \rho^{\text{tr}} \rangle^+ \quad (11)$$

where A^+ denotes the pseudoinverse of A . This solution is exact iff $\text{supp}(\langle \mathcal{O}, \rho^{\text{tr}} \rangle) \subseteq \text{supp}(\langle \tilde{\mu}, \rho^{\text{tr}} \rangle)$, and unique iff $\text{supp}(\langle \mathcal{O}, \rho^{\text{tr}} \rangle) = \text{supp}(\langle \tilde{\mu}, \rho^{\text{tr}} \rangle)$ ³². Given a Hermitian operator X and an informationally complete POVM $\tilde{\mu}$, there is always a dual POVM $\tilde{\mu}^*$, with $|\tilde{\mu}| = |\tilde{\mu}^*|$ that allows the decomposition³³

$$X = \sum_k \langle \tilde{\mu}_k^*, X \rangle \tilde{\mu}_k = \sum_k \langle \tilde{\mu}_k, X \rangle \tilde{\mu}_k^*. \quad (12)$$

The POVM $\tilde{\mu}^*$ is also referred to, in this context, as a dual frame of $\tilde{\mu}$. A particular choice of such a dual basis is constructed as

$$\tilde{\mu}_k^* = S^{-1}(\tilde{\mu}_k), \quad S(X) \equiv \sum_k \tilde{\mu}_k \langle \tilde{\mu}_k, X \rangle, \quad (13)$$

where S is referred to as the frame operator, which is ensured to be invertible, provided $\tilde{\mu}$ is informationally complete, and this basis is the canonical dual frame of $\tilde{\mu}$. With $\tilde{\mu}^*$, we can write

$$\langle \mathcal{O}_i, \rho \rangle = \sum_k \langle \mathcal{O}_i, \tilde{\mu}_k^* \rangle \langle \tilde{\mu}_k, \rho \rangle, \quad (14)$$

which tells us that a general solution to the linear reconstruction problem has the form

$$W = \langle \mathcal{O}, \tilde{\mu}^* \rangle. \quad (15)$$

This provides a very concrete understanding of what the training phase achieves: through training, and solving the associated linear problem, we retrieve a partial description of the measurement process itself, through its dual operators. Note that one can also consider this framework using a complete set of observables \mathcal{O}_i as target, in which case $\langle \mathcal{O}, \tilde{\mu}^* \rangle$ also amounts to a complete characterisation of $\tilde{\mu}^*$, and thus of $\tilde{\mu}$.

The performance of the QELM is quantified by its accuracy on previously unseen “test” states. A standard choice of quantifier is the mean squared error (MSE): given a test state ρ , and assuming that the training produced parameters \mathbf{w} , this reads

$$\text{MSE} = (\langle \mathcal{O}, \rho \rangle - \mathbf{w} \cdot \langle \tilde{\mu}, \rho \rangle)^2. \quad (16)$$

For multiple target observables, the definition is extended straightforwardly: we have

$$\text{MSE} = \|\langle \mathcal{O}, \rho \rangle - W \langle \tilde{\mu}, \rho \rangle\|_2^2, \quad (17)$$

where now $\mathcal{O} = (\mathcal{O}_1, \mathcal{O}_2, \dots)$ is a vector of target observables, and W the matrix obtained from the training phase.

In an ideal scenario, where the probabilities $\langle \tilde{\mu}, \rho \rangle$ are known with perfect accuracy, solving Eq. (9) is not an issue. Assuming that the system is indeed solvable — that is, Eq. (6) is satisfied — then any solution method, e.g. computing the pseudo-inverse of $\langle \tilde{\mu}, \rho^{\text{tr}} \rangle$, will result in some W which maps perfectly well measurement probabilities to expectation values. However, any realistic scenario will result in a radically different outlook. Because the protocol uses measurement probabilities as fundamental building blocks, being mindful of potential numerical instabilities is paramount. In particular, the probability vectors $\langle \tilde{\mu}, \rho \rangle$ will only be known up to a finite accuracy which depends on the finite number N of statistical samples, since the variance of the estimates will scale as N^{-1} .

These statistical fluctuations will both affect the estimation of $\langle \mathcal{O}, \tilde{\mu}^* \rangle$ in the training phase, and the final accuracies in the testing phase. The latter source of noise is present even if $\langle \mathcal{O}, \tilde{\mu}^* \rangle$ is known with perfect accuracy, while the former is due to the use of a finite training dataset.

Reconstruction efficiency. An important factor to consider when using QELMs is the potential numerical instability arising from solving the associated linear system³². While Eq. (11) provides a

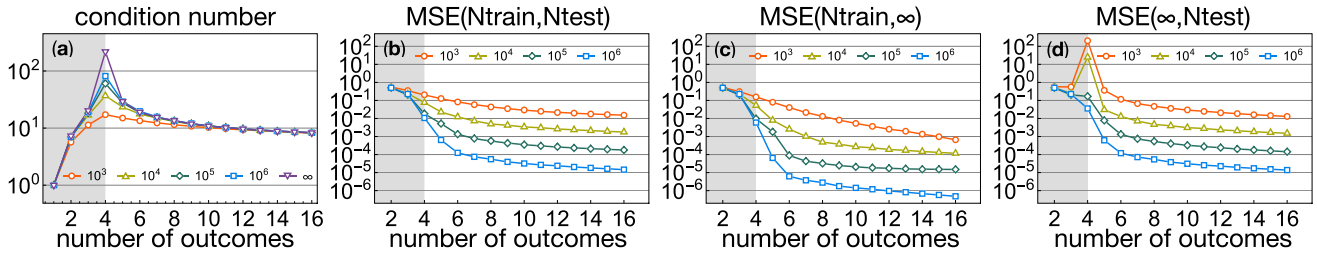


Fig. 2 Mean squared error and condition number vs number of outcomes. Mean squared error (MSE) associated to reconstruction of $\text{Tr}(\mathcal{O}\rho)$ for some single-qubit observable \mathcal{O} in the first scenario configuration, with $M_{\text{tr}} = 100$ and $M_{\text{test}} = 1000$ states used during training and testing phase, respectively. In all plots, different colours refer to different numbers of samples $N_{\text{train}}, N_{\text{test}}$ used to estimate the probabilities. The target observable is chosen at random, and kept fixed in all shown simulations. Choosing different observables does not significantly affect the behaviour of these plots. **a** Condition number of the probability matrix $(\tilde{\mu}, \rho^{\text{tr}})$ as a function of the number of measurement outcomes. **b** MSE as a function of the number of measurement outcomes, when both train and test probabilities are estimated with the same finite precision. **c** As above, but now the test probabilities are estimated with infinite precision. **d** As above, but now the training probabilities are estimated with infinite precision. In this last case, the large error corresponding to four outcomes is due to the amplification of the statistical error in the vector of probabilities \mathbf{p} by the map W . The amount of amplification is described by the condition number in Eq. (19).

general and efficiently computable solution to the learning problem, this solution can be ill-conditioned, i.e. small perturbations of the inputs can result in large perturbations of the outputs. In our context, this happens when $(\tilde{\mu}, \rho^{\text{tr}})$ has small singular values, which might arise due to noise or finite statistics. The issues associated to solving a linear system in a supervised learning context, and some possible ways to tackle them, are discussed in refs. 34–36. Depending on the circumstances, several regularisation techniques can be used to deal with ill-conditioned problems.

A standard way to quantify the potential ill-conditioned nature of a linear system is the condition number³²: Given a linear problem $\mathbf{y} = \mathbf{A}\mathbf{x}$ which one wishes to solve for \mathbf{x} , the condition number of A is

$$\kappa(A) = \frac{s_{\max}}{s_{\min}},$$

where s_{\max} (s_{\min}) is the largest (smallest) singular values of A . The set of solutions to the linear system is the affine space

$$\mathbf{x} \in A^+ \mathbf{y} + \ker(A), \quad (18)$$

where A^+ denotes the pseudo-inverse of A . A simple characterisation of $\kappa(A)$ is that it provides the worst-case scenario estimate of relative error amplification: if $\Delta\mathbf{y}$ is the error associated with \mathbf{y} , the relative error on \mathbf{x} is bounded by

$$\left\| \frac{\Delta\mathbf{x}}{\mathbf{x}} \right\| \leq \kappa(A) \left\| \frac{\Delta\mathbf{y}}{\mathbf{y}} \right\|. \quad (19)$$

Equation (10) is precisely the type of linear system whose numerical stability is estimated via the condition number, in this case $\kappa((\tilde{\mu}, \rho^{\text{tr}}))$, the overarching goal of QELMs is not accurately estimating W , but rather finding any W that results in accurately estimating the target expectation values on the test dataset. In other words, we only care about inaccuracies in the estimation of W in so far as they are reflected in inaccuracies in the MSE $\|(\mathcal{O}, \rho) - W(\tilde{\mu}, \rho)\|_2$. That means the errors we are interested in are those coming from the expression

$$\underbrace{\langle (\mathcal{O}, \rho^{\text{tr}}) | (\tilde{\mu}, \rho^{\text{tr}})^+ | (\tilde{\mu}, \rho) \rangle}_{= W}, \quad (20)$$

where both $(\tilde{\mu}, \rho^{\text{tr}})$ and $(\tilde{\mu}, \rho)$ are estimated up to some finite precision.

An unavoidable source of ill-conditioning is the fundamental statistical nature of the probabilities entering the $|\tilde{\mu}| \times M_{\text{tr}}$ matrix $P \equiv (\tilde{\mu}, \rho^{\text{tr}})$. Let P_N denote the matrix whose elements are the frequencies associated with the corresponding probabilities in P , estimated from N samples. If the input states have dimension

N_{input} (e.g. $N_{\text{input}} = 4$ for 2 qubits), but $|\tilde{\mu}| > N_{\text{input}}$, then P will have some vanishing singular values. Due to the statistical noise, these will become nonzero, albeit remaining relatively small with magnitude of the order of $1/N$, in P_N . This makes the linear inversion problem potentially ill-conditioned, as the eigenspaces corresponding to such singular values do not represent physically relevant information. A simple way to fix this issue is to truncate the singular values, setting to zero those beyond the N_{input} -th one. This strategy does not introduce a significant amount of error, as long as the variances associated to the outcome probabilities are sufficiently smaller than all the other (physically relevant) singular values, which is always the case for sufficiently large N . We will employ this strategy for our simulations.

Another feature is the increase of the condition number $\kappa(P_N)$ on N [cf. Fig. 2a]. This is somewhat counterintuitive, as we would expect estimation to become easier when the probabilities are known more accurately. We refer to Supplementary Note 1 for a detailed discussion of this aspect.

Single-injection examples. Let us consider how our framework applies to the case with single-qubit inputs. Most of the literature focuses on reservoir dynamics defined via some Hamiltonian¹⁰, or on open quantum systems¹¹. Our aim is here to study the performance of QELMs in standard scenarios, and we therefore focus on unitary evolutions, and analyse cases where the reservoir dynamics is a random unitary or isometric evolution rather than a specific Hamiltonian model, in order to gain a better insight into the performances of QELMs in more general contexts. More specifically, we focus on the following three scenarios:

1. The input qubits interact with a high-dimensional state through some random unitary evolution. In this case, the reservoir is a qudit, measured in some fixed computational basis, and the corresponding evolution reads: $\Lambda(\rho) = \text{Tr}_1[V\rho V^\dagger]$ with $V \in \text{U}(\mathbb{C}^2, \mathbb{C}^2 \otimes \mathbb{C}^n)$ a $(2n) \times 2$ isometry, for some $\mathbb{N} \ni n > 2$. In this notation, the initial state of the reservoir is implicitly specified through the choice of isometry V . The corresponding measurement is taken to be $\mu_j = |j\rangle\langle j|$ with $j = 1, \dots, n$, and the effective measurement thus reads

$$\tilde{\mu}_j = \Lambda^\dagger(|j\rangle\langle j|) = V^\dagger(I \otimes |j\rangle\langle j|)V. \quad (21)$$

2. Alternatively, one can consider a scenario involving a single high-dimensional qudit, with no bipartite structure involved. In this case, the “input qubit” is a two-dimensional subspace of the qudit, $\tilde{\rho} = \rho \oplus \eta_0$ with $\rho \in D(\mathbb{C}^2)$ a single-qubit state, and $\eta_0 \in D(\mathbb{C}^{n-2})$ the initial

state of the reservoir degrees of freedom. The dynamics is in this case simply an evolution of the form $\rho \oplus \eta_0 \mapsto U(\rho \oplus \eta_0)U^\dagger$ for some unitary operator $U \in U(\mathbb{C}^n)$. Measurements are again performed in the computational basis, $\mu_k = |k\rangle\langle k|$ with $k = 1, \dots, n$, and thus

$$\tilde{\mu}_k = U^\dagger |k\rangle\langle k| U. \tag{22}$$

In this notation, the degrees of freedom of the input state are also measured after the evolution, but this is not an issue for our purposes.

- As a further example, let us consider a system of qubits interacting through some Hamiltonian H . In this case, the input qubit interacts with N_R reservoir qubits through some Hamiltonian, and the measurement is performed on the reservoir qubits. The dynamics thus has the form

$$\rho \mapsto e^{-iHt}(\rho \otimes \eta_0)e^{iHt} \tag{23}$$

for some evolution time t and initial reservoir state $\eta_0 \in D(\mathbb{C}^{2^{N_R}})$. For our tests, we use a pairwise Hamiltonian for a qubit network of the form

$$H = \sum_{ij=1}^{N_R+1} J_{ij}(\sigma_i^+ \sigma_j^- + \sigma_j^+ \sigma_i^-) + \sum_{i=1}^{N_R} \Delta_i \sigma_i^x, \tag{24}$$

with random coupling constants J_{ij} drawn uniformly at random from the interval $[-1, 1]$, and driving coefficients Δ_i drawn uniformly at random from $[0, 1]$. We consider different network connectivities; in particular we study (1) a linear chain with nearest-neighbour interactions, (2) a fully connected reservoir, with a single node connected to the input, and finally (3) a fully connected reservoir where each node is connected to the input. If measurements are again performed in the computational basis of the reservoir, that is $\mu_k = |k\rangle\langle k|$ with $k = 1, \dots, 2^{N_R}$, the corresponding effective measurements will have the form

$$\tilde{\mu}_k = \text{Tr}_2[(I \otimes \eta_0)e^{iHt}(I \otimes |k\rangle\langle k|e^{-iHt})]. \tag{25}$$

As training objective, we consider the reconstruction of the expectation value of some target observable $\mathcal{O} \in \text{Herm}(\mathbb{C}^2)$. For our simulations, we make the conventional choice $\mathcal{O} = \sigma_x$, with σ_x the Pauli X matrix. Note that choosing different observables or different evolutions, does not significantly affect the results.

To reconstruct arbitrary linear functionals of ρ , the effective measurement must have rank four, that is, it must contain four linearly independent operators. This is required to have tomographically complete knowledge of ρ . This means that, in particular, the reservoir state must be at least four-dimensional.

Figure 2 reports the performances of QELMs trained to retrieve σ_x when the evolution corresponds to an input qubit interacting with a 2^5 -dimensional qudit through a random unitary operator, for different numbers of elements in the effective POVM $\{\tilde{\mu}_k\}$. Let $|\tilde{\mu}|$ denote the number of such elements. In the ideal scenario where training and test probabilities are known with perfect accuracy, the MSE is precisely zero whenever $|\tilde{\mu}| \geq 4$. To get more realistic results, we consider the performance when $\langle \tilde{\mu}, \rho^{\text{tr}} \rangle$ and $\langle \tilde{\mu}, \rho^{\text{test}} \rangle$ are estimated from finite statistics. In these scenarios, the condition number of the matrix $\langle \tilde{\mu}, \rho^{\text{tr}} \rangle$ is also relevant, as it correlates with how much the statistical fluctuations in $\langle \tilde{\mu}, \rho^{\text{test}} \rangle$ can be amplified and lead to estimation inaccuracies. As shown in the figures, the accuracy increases with better statistics, as expected, but also when increasing $|\tilde{\mu}|$. It is worth stressing that this feature does not occur with the ideal probabilities, as in that scenario the MSE is perfectly zero from four measurements onwards. More precisely, we should say that the ideal MSE vanishes almost always when the unitary evolution is drawn

uniformly at random. It is in fact possible to find examples of unitaries which make the reconstruction impossible. Trivial examples would be unitaries that do not correlate input and reservoir degrees of freedom. These cases almost never occur when drawing unitaries at random, however.

Figure 2 shows that, although 4 measurements are in principle sufficient to retrieve the target information, reconstruction in realistic circumstances becomes easier when increasing the dimension of the reservoir, that is, the number of measurement outcomes. In Fig. 2a we see that the numerical problem becomes better conditioned when there are more measurement outcomes. In Fig. 2b, c we appreciate how the accuracy increases when more statistical samples are used, and thus the probabilities approach their ideal values.

Figure 2d shows the MSE when the training parameters are computed from the ideal probability matrix $\langle \tilde{\mu}, \rho^{\text{tr}} \rangle$, while finite statistics is used in the testing phase. In this case, the poor statistical accuracy found for small numbers of outcomes is due to the correspondingly large condition number. This is to be attributed to numerical instability associated with the ideal reconstruction parameters for few measurement outcomes: indeed the large error corresponding to four outcomes is due to the amplification of the statistical error in the vector of probabilities, amplification that is quantified by the condition number Eq. (19). Note that such detrimental effect largely disappears already for $|\tilde{\mu}| \geq 8$. Note that the data shown in Fig. 2d and in the purple triangles in Fig. 2a corresponds to a training performed with perfectly estimated training probabilities. Even if not directly related to performances in practical scenarios, this data is useful to better isolate the different effects caused by inaccuracies during training and testing phases.

Finally, in Fig. 3 we consider how different choices of dynamics influence the reconstruction performances. In particular, we consider input states interacting with the reservoir through a random unitary evolution, a random pairwise Hamiltonian, or a randomly drawn pairwise Hamiltonian with a chain structure, in which each qubit only interacts with its nearest neighbour. Overall, as the degree of connectivity of the network increases, the performance of the reservoir and stability of the linear regression both improve. This is illustrated by the decrease in the MSE and the condition number.

Multiple injections scenario. We focused in the previous section on the achievability of target observables when single copies of an input state are made to interact with a reservoir, which is then measured. As shown, this characterises the amount of exactly retrievable information from functionals that are linear in the input density matrix. Let us consider now the more general scenario where several copies of an input state are allowed as input. We will show that this allows retrieving a broader range of properties of the input states.

Consider a channel \mathcal{E}_ρ applied multiple times to an initial reservoir state η_0 . For n consecutive uses of the channel and initial reservoir state η_0 , the measured state is then $\eta_n = \mathcal{E}_\rho^n(\eta_0)$. This can be rewritten as

$$\eta_n = \Phi(\rho \otimes \Phi(\rho \otimes \eta_{n-2})) = \dots = \tilde{\Phi}(\rho^{\otimes n} \otimes \eta_0), \tag{26}$$

where we have introduced the resulting channel $\tilde{\Phi}$. Following the same argument employed in the single-injection scenario, we see that the possible outputs after linear post-processing of the outcome probabilities are all and only those of the form

$$y = \text{Tr}(\tilde{\mathcal{O}}\rho^{\otimes n}), \tag{27}$$

for some observable $\tilde{\mathcal{O}}$ acting in the space of n copies of ρ .

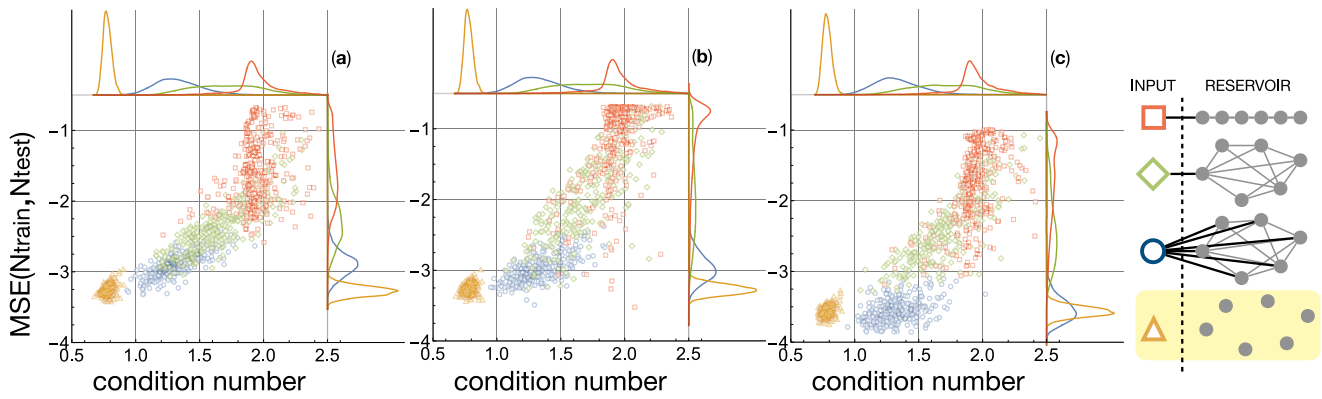


Fig. 3 Relation between mean squared error and condition number. MSE (in logarithmic scale) obtained by training random reservoirs corresponding to different types of dynamics to retrieve a fixed target observable, shown against the condition number. The target one-qubit observable \mathcal{O} is **a** the σ_x Pauli matrix, **b** the σ_z Pauli matrix, and **c** a one-qubit observable sampled at random. In each case, we plot data corresponding to a reservoir dynamics that is (red squares) a random one-dimensional spin chain with nearest-neighbour interactions, (green diamond) a random fully connected spin Hamiltonian, where the input is only connected to a single node of the reservoir, (blue circles) a random fully connected spin Hamiltonian, and (orange triangles) a random unitary evolution. Each point shows simulation results obtained using $M_{tr} = 100$ training states, $M_{test} = 1000$ test states, and a reservoir comprised of 6 qubits. The statistics is fixed to $N_{train} = N_{test} = 10^4$ samples used to estimate each measurement probability. Except for the random unitary case (orange triangles), representative of the first scenario, the other configuration are examples of the third scenario and the Hamiltonian parameters J_{ij} and Δ_i from Eq. (24) are sampled uniformly at random in the interval $[0, 1]$.

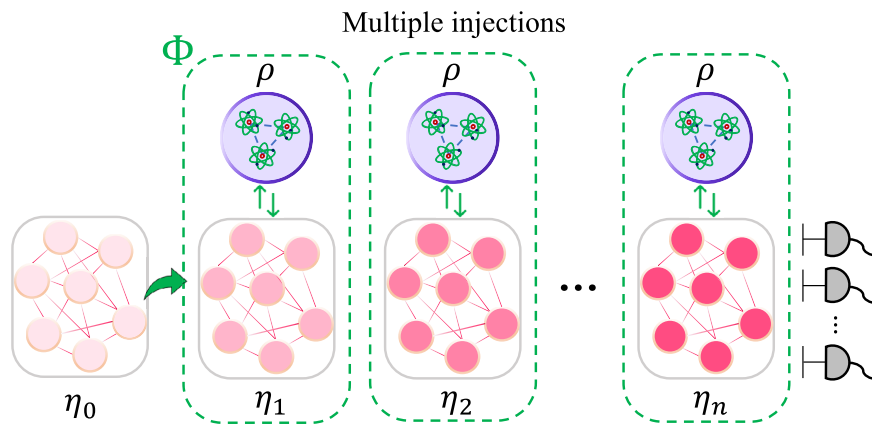


Fig. 4 Summary multiple injection QELMs. The reservoir interacts with multiple copies of the same quantum state ρ and progressively acquires information about it. The measurement is performed on the final reservoir state η_n .

Training these models thus proceeds similarly to the linear case: the probabilities are estimated from measurements performed after each series of n injections, and these probabilities are then used to solve Eq. (7) and thus find the optimal train parameters W . Figure 4 shows a scheme of the multiple-injection model here described.

After n injections, the space on which the effective POVM acts has dimension

$$d_{n,m} \equiv \binom{m+n-1}{n} \tag{28}$$

with $\dim(\mathcal{H}) \equiv m$ the dimension of each input state. This is the number of degrees of freedom characterising a symmetric tensor of the form $\rho^{\otimes n}$ with $\rho \in \mathcal{D}(\mathcal{H})$. This is also the space where the target observables \mathcal{O} live. It follows that, in order to be able to reconstruct arbitrary functionals of ρ up to the maximum order of n , the measurement must contain at least $d_{n,m}$ linearly independent components.

Consider for example the task of estimating the purity of a given state. Observe that the map $\rho \mapsto \text{Tr}(\rho^2)$ can be written as $\text{Tr}(\rho^2) = \text{Tr}[\text{SWAP}(\rho \otimes \rho)]$. As per our previous observations,

this means that the purity can be retrieved from a QELM provided that at least two injections are used, and that the effective measurement $\tilde{\mu}$ is such that SWAP can be expressed as a real linear combination of the measurement operators $\tilde{\mu}_b$.

Multiple injections examples. To showcase the reconstruction of nonlinear functionals of the input state, in Fig. 5 we consider targets functionals of the form $\text{Tr}(\mathcal{O}\rho^k)$, $\text{Tr}(e^\rho)$, and $\sqrt{1 - \text{Tr}(\rho^2)}$. We focus on the number of injections required for the reconstruction in each case, and thus assume ideal training and target probabilities.

In Fig. 5a we give the MSE associated with the reconstruction of $\text{Tr}(\mathcal{O}\rho^k)$, $k = 1, \dots, 7$, for a random one-qubit observable \mathcal{O} , for different numbers of injections. In these simulations, the reservoir is an 8-qubit system, with an additional qubit used for the input states, reset to ρ for each injection. As expected from our previous discussion, we observe that the reconstruction is only successful when the number of injections n is larger than the degree k of the polynomial of the target observable.

Furthermore, note how the reconstruction fails again when the number of injections increases too much. This upper bound for the

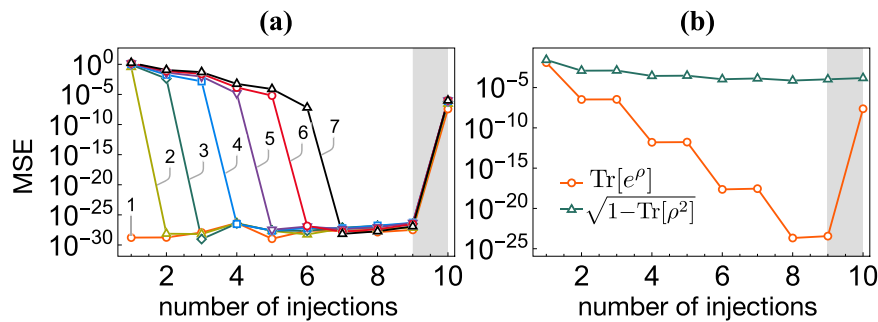


Fig. 5 Reconstruction error vs number of injections. **a** Reconstruction MSE for polynomial targets $\text{Tr}(\mathcal{O}_p^k)$ with $k=1, \dots, 7$. As previously discussed, reconstruction is not possible unless the number of injections is greater than or equal to the degree of the target function (for polynomial target functions). At the same time, the number of independent measurements bounds the number of injections that can be used without degrading the information. Here, the reservoir consists of 8 qubits, which is why the reconstruction fails when more than 9 injections are used, consistently with Eq. (28): with 8 qubits there are $2^8 = 256$ available measurements, while the 9-injection space has dimension $\binom{2^8+9-1}{9} = 220 < 256$. Using 10 injections implies a dimension $\binom{2^8+10-1}{10} = 286 > 256$. **b** Reconstruction MSE for the nonlinear targets $\text{Tr}(e^\rho)$ and $\sqrt{1 - \text{Tr}(\rho^2)}$. Due to the nonlinearity of these functions, ideal reconstruction is never feasible. The significantly better performances obtained for $\text{Tr}(e^\rho)$ are due to the coefficients of its Taylor series vanishing faster than those of $\sqrt{1 - \text{Tr}(\rho^2)}$.

reconstruction is due to the finite dimension of the reservoir — or, equivalently, the finite number of measurement operators in $\tilde{\mu}$. In fact, reconstructing $\text{Tr}(\mathcal{O}_p^k)$ from a measurement performed after n injections amounts to reconstructing a specific observable $\tilde{\mathcal{O}}$ acting on the space of states of the form $\rho^{\otimes n}$. If the measurements are not suitably chosen, as is the case in QELM-like scenarios, this means that the number of (linearly independent) measurements must be sufficient to reconstruct all possible observables on such a space, whose dimensionality is $d_{n,m}$.

In Fig. 5b we treat the case of nonlinear functionals of ρ . The performance achieved in approximating $\sqrt{1 - \text{Tr}(\rho^2)}$ is poor due to the slow convergence of the Taylor expansion of the functional. The step-like behaviour that is evident in the MSE associated with the reconstruction of $\text{Tr}(e^\rho)$, which is also present in the case of $\sqrt{1 - \text{Tr}(\rho^2)}$ although less evidently, can be explained by noticing that the trace of odd powers of ρ is a polynomial of the same degree of the previous even ones.

Conclusion

We provided a complete characterisation of the information exactly retrievable from linear post-processing of measurement probabilities in QELM schemes. This sheds light on the tight relation between the capabilities of a device to retrieve nonlinear functionals of input states, and the memory of the associated quantum channel.

We found that the estimation efficiency of QELM protocols is entirely reflected in the properties of an effective POVM describing the entire apparatus, comprised of a dynamical evolution and a measurement stage. In particular, we showed that the effective POVM contains all of the information required to determine which observables can be estimated, and to what accuracy, as well as which kinds of effective POVMs, induced by different types of dynamics, result in different degrees of estimation accuracies. In turn, this clarifies the class of dynamics that result in POVMs that are effective for efficient and accurate property estimations. We further found that the inevitable sampling noise, intrinsic to any measurement data coming from a quantum device, crucially affects estimation performances, and cannot be neglected when discussing the protocols.

Our work paves the way for a number of future endeavours on this line of research, including an extension of our analysis to time-trace signals for dynamical QRCs, and the in-depth analysis of POVM optimality for quantum state estimation. Moreover, the

translation of our findings into performance-limiting factors of recently designed experimental scenarios for QELMs and QRCs, and the identifications of ways to counter them, will be paramount for the grounding of the role that such architectures could play in the development of schemes for quantum property validation. At the same time, our study of QELMs for quantum state estimation purposes fits tightly with, and has the potential to improve on, several experimental detection strategies which rely on some form of linear regression to estimate target states^{37–41}.

Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

Received: 15 September 2022; Accepted: 8 May 2023;

Published online: 25 May 2023

References

- Huang, G.-B., Zhu, Q.-Y. & Siew, C.-K. 2004 *IEEE international joint conference on neural networks (IEEE Cat. No. 04CH37541)*, Vol. 2, 985–990 (IEEE, 2004).
- Huang, G.-B., Wang, D. H. & Lan, Y. Extreme learning machines: a survey. *Int. J. Machine Learn. Cybern.* **2**, 107–122 (2011).
- Wang, J., Lu, S., Wang, S.-H. & Zhang, Y.-D. *Multimedia Tools and Applications* 1–50, <https://link.springer.com/article/10.1007/s11042-021-11007-7> (2021).
- Lukoševičius, M. & Jaeger, H. Reservoir computing approaches to recurrent neural network training. *Comput. Sci. Rev.* **3**, 127–149 (2009).
- Lukoševičius, M. *Neural Networks: Tricks of the Trade*. 659–686, https://link.springer.com/chapter/10.1007/978-3-642-35289-8_36 (2012).
- Angelatos, G., Khan, S. A. & Türeci, H. E. Reservoir computing approach to quantum state measurement. *Phys. Rev. X* **11**, 041062 (2021).
- Jaeger, H. The “echo state” approach to analysing and training recurrent neural networks—with an erratum note. *German National Research Center for Information Technology GMD Technical Report*. Vol. 148, p. 13 (German National Research Center for Information Technology, 2001).
- Jaeger, H. & Haas, H. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science* **304**, 78–80 (2004).
- Van der Sande, G., Brunner, D. & Soriano, M. C. Advances in photonic reservoir computing. *Nanophotonics* **6**, 561–576 (2017).
- Fujii, K. & Nakajima, K. Harnessing disordered-ensemble quantum dynamics for machine learning. *Phys. Rev. Appl.* **8**, 024030 (2017).
- Ghosh, S., Opala, A., Matuszewski, M., Paterek, T. & Liew, T. C. H. Quantum reservoir processing. *npj Quant. Inf.* **5**, 1–6 (2019).
- Ghosh, S., Krisnanda, T., Paterek, T. & Liew, T. C. Realising and compressing quantum circuits with quantum reservoir computing. *Commun. Phys.* **4**, 1–7 (2021).

13. Martínez-Peña, R., Nokkala, J., Giorgi, G. L., Zambrini, R. & Soriano, M. C. Information processing capacity of spin-based quantum reservoir computing systems. *Cognit. Comput.* 1–12, <https://doi.org/10.1007/s12559-020-09772-y> (2020).
14. Kutvonen, A., Fujii, K. & Sagawa, T. Optimizing a quantum reservoir computer for time series prediction. *Sci. Rep.* **10**, 1–7 (2020).
15. Tran, Q. H. & Nakajima, K. Higher-order quantum reservoir computing. *arXiv* <https://doi.org/10.48550/arXiv.2006.08999> (2020).
16. Martínez-Peña, R., Giorgi, G. L., Nokkala, J., Soriano, M. C. & Zambrini, R. Dynamical phase transitions in quantum reservoir computing. *Phys. Rev. Lett.* **127**, 100502 (2021).
17. Krisnanda, T., Ghosh, S., Paterek, T. & Liew, T. C. Creating and concentrating quantum resource states in noisy environments using a quantum neural network. *Neural Netw.* **136**, 141–151 (2021).
18. Rafayelyan, M., Dong, J., Tan, Y., Krzakala, F. & Gigan, S. Large-scale optical reservoir computing for spatiotemporal chaotic systems prediction. *Phys. Rev. X* **10**, 041037 (2020).
19. Nokkala, J. et al. Gaussian states of continuous-variable quantum systems provide universal and versatile reservoir computing. *Commun. Phys.* **4**, 1–11 (2021).
20. Nakajima, K., Fujii, K., Negoro, M., Mitarai, K. & Kitagawa, M. Boosting computational power through spatial multiplexing in quantum reservoir computing. *Phys. Rev. Appl.* **11**, 034021 (2019).
21. Mujal, P. Quantum reservoir computing for speckle disorder potentials. *Condens. Matter* **7**, 17 (2022).
22. Mujal, P., Martínez-Peña, R., Giorgi, G. L., Soriano, M. C. & Zambrini, R. Time-series quantum reservoir computing with weak and projective measurements. *npj Quant. Inf.* **9**, 16 (2023).
23. Martínez-Peña, R. & Ortega, J.-P. Quantum reservoir computing in finite dimensions. *Phys. Rev. E* **107**, 035306 (2023).
24. Tanaka, G. et al. Recent advances in physical reservoir computing: a review. *Neural Netw.* **115**, 100–123 (2019).
25. Fujii, K. & Nakajima, K. *Reservoir Computing*, 423–450 (Springer, 2021).
26. Mujal, P. et al. Opportunities in quantum reservoir computing and extreme learning machines. *Adv. Quant. Technol.* **4**, 2100027 (2021).
27. Chen, J., Nurdin, H. I. & Yamamoto, N. Temporal information processing on noisy quantum computers. *Phys. Rev. Appl.* **14**, 024065 (2020).
28. Ghosh, S., Opala, A., Matuszewski, M., Paterek, T. & Liew, T. C. Reconstructing quantum states with quantum reservoir networks. *IEEE Trans. Neural Netw. Learn. Syst.* **32**, 3148–3155 (2020).
29. Tran, Q. H. & Nakajima, K. Learning temporal quantum tomography. *Phys. Rev. Lett.* **127**, 260401 (2021).
30. Watrous, J. *The Theory of Quantum Information* (Cambridge University Press, 2018).
31. Govia, L. C. G., Ribeill, G. J., Rowlands, G. E. & Ohki, T. A. Nonlinear input transformations are ubiquitous in quantum reservoir computing. *Neuromorph. Comput. Eng.* **2**, 014008 (2022).
32. Higham, N. J. *Accuracy and Stability of Numerical Algorithms* (Society for Industrial and Applied Mathematics, 2002).
33. Casazza, P. G. & Lynch, R. G. A brief introduction to hilbert space frame theory and its applications. *arXiv* <https://doi.org/10.48550/arXiv.1509.07347> (2015).
34. Rosasco, L., Caponnetto, A., Vito, E., Odone, F. & Giovannini, U. Learning, regularization and ill-posed inverse problems. *Adv. Neural Inf. Process. Syst.* **17**, 1145–1152 (2004).
35. De Vito, E. et al. Learning from examples as an inverse problem. *J. Machine Learn. Res.* <https://www.jmlr.org/papers/v6/devito05a.html> (2005).
36. Zhao, G., Shen, Z., Miao, C. & Man, Z. On improving the conditioning of extreme learning machine: a linear case. In *2009 7th International Conference on Information, Communications and Signal Processing (ICICSP)*, 1–5 (IEEE, 2009).
37. Zia, D. et al. Regression of high-dimensional angular momentum states of light. *Phys. Rev. Res.* **5**, 013142 (2023).
38. Suprano, A. et al. Enhanced detection techniques of orbital angular momentum states in the classical and quantum regimes. *New J. Phys.* **23**, 073014 (2021).

39. Suprano, A. et al. Dynamical learning of a photonics quantum-state engineering process. *Adv. Photon.* **3**, 066002–066002 (2021).
40. Stricker, R. et al. Experimental single-setting quantum state tomography. *PRX Quant.* **3**, 040310 (2022).
41. García-Pérez, G. et al. Learning to measure: adaptive informationally complete generalized measurements for quantum algorithms. *PRX Quant.* **2**, 040342 (2021).

Acknowledgements

LI acknowledges support from MUR and AWS under project PON Ricerca e Innovazione 2014–2020, “calcolo quantistico in dispositivi quantistici rumorosi nel regime di scala intermedia” (NISQ - Noisy, Intermediate-Scale Quantum). IP is grateful to the MSCA Cofund project CITI-GENS (Grant nr. 945231). MP acknowledges the support by the European Union’s Horizon 2020 FET-Open project TEQ (766900), the Horizon Europe EIC Pathfinder project QuCoM (Grant Agreement No. 101046973), the Leverhulme Trust Research Project Grant UltraQuTe (grant RGP-2018-266), the Royal Society Wolfson Fellowship (RSWF/R3/183013), the UK EPSRC (EP/T028424/1), and the Department for the Economy Northern Ireland under the US-Ireland R&D Partnership Programme.

Author contributions

L.I., S.L., I.P., A.F., M.P. and G.M.P. developed the core ideas of the research. L.I., S.L. and I.P. performed the numerical and analytical calculations. L.I., S.L., I.P., A.F., M.P. and G.M.P. all contributed to interpreting and overwriting the obtain numerical and analytical results. All authors contributed to the writing of the manuscript. M.P. and G.M.P. supervised the research.

Competing interests

The authors declare no competing interest.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-023-01233-w>.

Correspondence and requests for materials should be addressed to L. Innocenti.

Peer review information *Communications Physics* thanks the anonymous reviewers for their contribution to the peer review of this work. A peer review file is available.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023