

Tackling Sampling Noise in Physical Systems for Machine Learning Applications: Fundamental Limits and Eigentasks

Fangjun Hu^{1,*}, Gerasimos Angelatos^{1,2,†}, Saeed A. Khan^{1,†}, Marti Vives^{1,3}, Esin Türeci⁴, Leon Bello¹,
Graham E. Rowlands², Guilhem J. Ribeill², and Hakan E. Türeci¹

¹Department of Electrical and Computer Engineering, Princeton University,
Princeton, New Jersey 08544, USA

²Raytheon BBN, Cambridge, Massachusetts 02138, USA

³Q-CTRL, Santa Monica, California 90401, USA

⁴Department of Computer Science, Princeton University, Princeton, New Jersey 08544, USA

(Received 4 January 2023; revised 30 August 2023; accepted 5 September 2023; published 30 October 2023)

The expressive capacity of physical systems employed for learning is limited by the unavoidable presence of noise in their extracted outputs. Though present in physical systems across both the classical and quantum regimes, the precise impact of noise on learning remains poorly understood. Focusing on supervised learning, we present a mathematical framework for evaluating the resolvable expressive capacity (REC) of general physical systems under finite sampling noise and provide a methodology for extracting its extrema, the eigentasks. Eigentasks are a native set of functions that a given physical system can approximate with minimal error. We show that the REC of a quantum system is limited by the fundamental theory of quantum measurement and obtain a tight upper bound for the REC of any finitely sampled physical system. We then provide empirical evidence that extracting low-noise eigentasks can lead to improved performance for machine learning tasks such as classification, displaying robustness to overfitting. We present analyses suggesting that correlations in the measured quantum system enhance learning capacity by reducing noise in eigentasks. The applicability of these results in practice is demonstrated with experiments on superconducting quantum processors. Our findings have broad implications for quantum machine learning and sensing applications.

DOI: 10.1103/PhysRevX.13.041020

Subject Areas: Quantum Physics, Quantum Information

I. INTRODUCTION

A physical system receiving an input stimulus typically evolves in response to it, such that its degrees of freedom become dependent on said input after a certain period of interaction with it. This everyday observation has a profound implication: any dynamical system can be viewed as performing a transformation of its input, realizing an input-output map [1]. This functional map can, in principle, be optimized, inspiring an emerging approach to learning with analog physical systems, which we collectively refer to as physical neural networks (PNNs) [2–4]. PNNs employ a wide variety of analog physical systems to compute a trainable transformation on an input signal [5–17]. More precisely, the role of an idealized (i.e., completely

deterministic, noise-free) physical system in these approaches is that of a high-dimensional feature generator. Given inputs \mathbf{u} , the measured degrees of freedom $x_k(\mathbf{u})$ for $k \in [K]$, generated by the system, act as an input-dependent vector of features. These features are used to approximate a function $f(\mathbf{u})$ via a learned linear projection with sufficient accuracy, as dictated by a chosen loss function (see Fig. 1). Different characteristics of the physical system, described by a set of hyperparameters θ , may determine its ability to approximate a particular function. Consequently, the relationship between a specific physical system and the classes of functions it can express with high accuracy is a fundamental question in this paradigm of machine learning [17–22].

No physical system, however, exists in isolation and is, therefore, necessarily subject to noise. Noise can enter at the input, whereby it evolves under the same dynamical law governing the evolution of the physical system. There may also be variability in this very dynamics of the physical system itself. Finally, there is typically noise associated with the measurement of output features from the physical system. As a consequence of these noise sources, the resulting feature map is stochastic: even under identical preparations and inputs \mathbf{u} , the outcome of a measurement $X_k^{(s)}(\mathbf{u})$ of a feature k can vary between repetitions, each of

*Corresponding author: fhu@princeton.edu

†These authors contributed equally to this work.

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

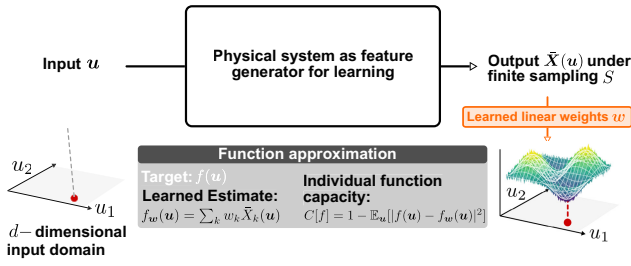


FIG. 1. Framework of learning with physical systems we consider in this work: inputs \mathbf{u} are transformed to a set of outputs $\tilde{X}(\mathbf{u})$ via a parametrized feature generator, implemented using an arbitrary physical system. Outputs are combined with appropriate weights \mathbf{w} to approximate a desired function f . Capacity $C[f]$ quantifies the error in this approximation f_w . We consider normalized functions $\mathbb{E}_u[f^2] = 1$, where \mathbb{E}_u denotes the expectation over the input domain with respect to a chosen measure.

which is referred to as a “shot” s . By empirically averaging the outcomes of S shots, one can generally reduce this stochasticity. We refer to the resulting noise as “sampling noise.” Theoretical analysis and experimental implementations of PNNs have already demonstrated that sampling noise can have a substantial role in the ultimate performance of a physical learning machine [11, 12, 15]. However, it is also known that this role may be more subtle than a limitation on performance across the board, as evidenced, for example, in the effective use of noise for regularization to aid generalizability in learning [23–25].

Often, heuristic descriptions are used to theoretically model such sampling noise and explain its effects on learning [8, 11, 18, 26]. However, when considering physical *quantum* systems for learning, a fundamental microscopic model for sampling noise is provided and, in fact, imposed by the quantum theory of measurement. For a quantum system prepared in an initial state density matrix $\hat{\rho}_0$ and evolving under an input-parametrized quantum channel $\mathcal{U}(\boldsymbol{\theta}, \mathbf{u})$, the final state is $\hat{\rho}(\mathbf{u}) = \mathcal{U}(\boldsymbol{\theta}, \mathbf{u})\hat{\rho}_0$ [27]. Sampling noise in measured features from this quantum system is constrained by the choice of measurement projectors \hat{M}_k associated with $\hat{\rho}(\mathbf{u})$. Unless the physical transformation defined by the quantum channel is optimized to yield only specific highly localized $\hat{\rho}(\mathbf{u})$ in the eigenspace of \hat{M}_k (as in quantum algorithms such as Grover’s or Shor’s [28])—a significant design restriction—or an excessively large number of shots S is used—a significant hardware restriction—such quantum sampling noise is an intrinsic component of learning with quantum systems.

Therefore, a framework is required that can account for sampling noise across generic physical systems and provide tools for learning when sampling noise is unavoidable. In this paper, we address the following question directly: what is the resolvable function space of an arbitrary physical system when regarded as an input-output machine *in the presence of sampling noise*? This simple objective

leads us to a general mathematical framework with important consequences for statistical learning theory, which we now overview. Our analysis is centered around a specific metric, the resolvable expressive capacity (REC), which is a generalization of the information processing capacity introduced in Ref. [18] (see also the earlier work in Ref. [29]) to account for the presence of sampling noise. Specifically, the REC is a quantitative measure of the accuracy with which K system-specific orthogonal functions can be constructed from K stochastic features $\tilde{X}_k(\mathbf{u})$. Remarkably, this accuracy has a tight, calculable S -dependent upper bound. The special functions, referred to as the eigentasks $y^{(k)}(\mathbf{u})$ of the physical system, define the maximally resolvable function space under S shots, which sets the stage for the introduction of a learning methodology in the presence of sampling noise.

Crucially, our framework can be applied to an arbitrary physical system via the solution of a simple matrix eigenproblem. The matrices in question are standard Gram matrix \mathbf{G} and covariance matrix \mathbf{V} , which can be estimated using stochastic samples from the system as a function of inputs \mathbf{u} over the domain of interest; the analysis can therefore be directly implemented in experimental settings without an internal model of the system. The solution of this linear eigenproblem yields both the eigentasks $\{y^{(k)}(\mathbf{u})\}$ and associated “noise-to-signal” eigenvalues $\{\beta_k^2\}$, which codify the normalized noise power in a construction of $y^{(k)}(\mathbf{u})$ from finite- S $\tilde{X}_k(\mathbf{u})$. The REC of the system is then only a function of $\{\beta_k^2\}$ and S .

In the second part of this paper, we develop eigentask learning, a means of learning in physical systems where sampling noise dominates, by using the noise-ordered eigentasks to construct a maximally resolvable basis of measured features. Our approach affects a systematic removal of high-noise features during training, which we demonstrate in experiments. These experimental demonstrations provide empirical evidence of robustness to overfitting in supervised learning, enhancing generalizability in the presence of sampling noise. Such a learning scheme built on avoiding features identified as having large noise may, in fact, be at play in natural physical systems such as biological neural circuits [30]. A well-studied example is that of neural vision: here, input visual stimuli drive stochastic dynamics of sensory neurons in the visual cortex, which must together elicit a target response, such as the brain correctly distinguishing two images. Studies have shown that the dynamics of individual neurons under nominally identical stimuli can exhibit great variability on a shot-by-shot basis [31]; however, in spite of the significant noise, the overall driven behavior remains capable of distinguishing visual stimuli with high fidelity. Studies analyzing the robust neural code despite noisy neural activity have found emergent global coding directions in the population activity that evade “modes” with maximal noise [26, 30]. The eigentask construction introduced here can be viewed as a generalization of

this idea of noise-ordered modes to function spaces over an arbitrary input domain and for an arbitrary physical system.

The eigentask learning framework is sensitive to not just the properties of the noise itself (encoded in \mathbf{V}), but also any dependence between it and the noise-free features via the Gram matrix \mathbf{G} . Such a situation typically prevails when the dominant source of sampling noise is either part of or evolves under the same input-output map defined by the physical system, as opposed to a completely uncorrelated noise source downstream. A simple example illustrating this, and one we analyze in detail, is that of a classical optical system, with features measured via photo-detection. Here, the sampled features—the integrated photocurrents—are subject to shot noise whose variance is related to the mean of the photocurrents themselves.

In the quantum regime of operation of a physical system—our ultimate focus—this relation between sampling noise and the state of the physical system emerges in the most general description of quantum measurement as a positive operator-valued measure (POVM). We formulate the computation of REC and eigentasks for arbitrary quantum systems in the presence of this fundamental sampling noise structure; we focus on qubit-based quantum systems (including gate-based circuits and quantum annealers) operated as untrained PNNs under static inputs [so-called extreme learning machines (ELMs) [8,32]], but our analysis is applicable to far more general quantum sensing and learning platforms. To validate the theoretical findings and emphasize their ready applicability to experimental scenarios, we implement an ELM through a parametrized quantum circuit encoding on an IBMQ superconducting processor, demonstrating the calculation of REC, the construction of eigentasks, and the application of eigentask learning to a classification task. In all cases, excellent agreement is seen with numerical simulation, and direct correlation is observed between REC and success at the considered classification task. This invites the exploration of principles to maximize the finite-sampling REC of a quantum system; for the qubit-based systems analyzed here, we show that an increase in measured quantum correlations can aid this goal.

The remainder of this paper is organized as follows. Section II provides the general theory of REC and eigentasks with respect to sampling noise in generic supervised physical learning systems and presents a calculation for a basic classical optical PNN. Section III applies REC theory and eigentask construction to machine learning with quantum systems, which is then validated and demonstrated with experiments performed on a seven-qubit IBMQ superconducting processor in Sec. IV. Finally, conclusions are presented in Sec. V.

II. THEORETICAL ANALYSIS

A. Sampling noise in learning with physical systems

The most general approach to supervised learning from classical data using a generic physical system is depicted

schematically in Fig. 1. A table with symbols and abbreviations used in the text can be found in Table I of Appendix A. We consider a scheme that begins with “embedding” the classical input data \mathbf{u} , sampled from a distribution $p(\mathbf{u})$, into the physical system to be used for learning. The form of this embedding is unrestricted beyond the requirement of being physical, and its precise nature influences the REC and eigentasks; some concrete examples are provided shortly.

In order to access information from the physical system after its interaction with the input, measurements must be performed on its K accessible degrees of freedom. For a fixed input \mathbf{u} , a single measurement or “shot” s yields *single-shot* random-valued features $\{X_k^{(s)}(\mathbf{u})\}$ for each $k \in [K]$. We define the *measured features* \bar{X}_k as S -shot sample means of $\{X_k^{(s)}\}$:

$$\bar{X}_k(\mathbf{u}) = \frac{1}{S} \sum_{s=1}^S X_k^{(s)}(\mathbf{u}), \quad (1)$$

whose expectation (equivalently, via the central limit theorem, the S -infinite limit) is given by

$$x_k(\mathbf{u}) \equiv \mathbb{E}_{\mathcal{X}}[\bar{X}_k(\mathbf{u})] = \lim_{S \rightarrow \infty} \bar{X}_k(\mathbf{u}), \quad (2)$$

where \mathbf{u} is regarded as a free variable. To be more precise, the expectation is evaluated over the product distribution of S independent and identically distributed (IID) vectors $\mathcal{X}(\mathbf{u}) \equiv \{(X_0^{(s)}(\mathbf{u}), X_1^{(s)}(\mathbf{u}), \dots, X_{K-1}^{(s)}(\mathbf{u}))\}_{s \in [S]}$, conditioned on a fixed \mathbf{u} .

With the definition of their expectation in Eq. (2), the measured features $\bar{\mathbf{X}}(\mathbf{u}) \in \mathbb{R}^K$, a column vector consisting of $\bar{X}_k(\mathbf{u})$, can be conveniently decomposed by extracting its deterministic mean value, together with a zero-mean, input-dependent noise term $\boldsymbol{\zeta}(\mathbf{u})$:

$$\bar{\mathbf{X}}(\mathbf{u}) = \mathbf{x}(\mathbf{u}) + \frac{1}{\sqrt{S}} \boldsymbol{\zeta}(\mathbf{u}). \quad (3)$$

Here, $\boldsymbol{\zeta}$ encodes the statistics of sampling; it generally has nontrivial cumulants of all orders, of which the covariances take the particular S -independent form $\boldsymbol{\Sigma}(\mathbf{u}) \in \mathbb{R}^{K \times K}$:

$$\boldsymbol{\Sigma}_{jk}(\mathbf{u}) \equiv \text{Cov}_{\mathcal{X}}[\zeta_j(\mathbf{u}), \zeta_k(\mathbf{u})] \quad (4)$$

and depend on only input \mathbf{u} . We note that Eq. (3) is exact. The factor of $1/\sqrt{S}$ is merely extracted for convenience of the analysis to follow and is *not* meant to suggest an expansion for large S at this stage; cumulants of $\boldsymbol{\zeta}$ beyond second order inherit a complicated S dependence.

The general input-output relationship $\mathbf{u} \rightarrow \bar{\mathbf{X}}_k(\mathbf{u})$ above can be made concrete by considering three example physical systems, depicted in Fig. 2. For an *optical system*,

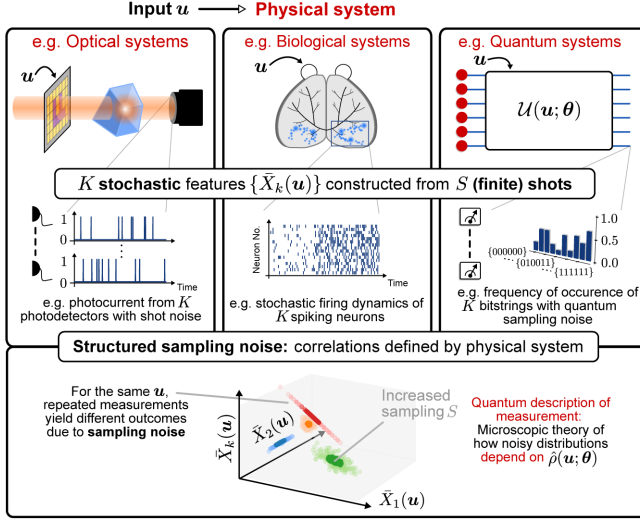


FIG. 2. Three distinct examples of physical systems for learning. Extracted information takes the form of K stochastic features \bar{X} obtained under S shots. For quantum systems, the geometric structure of distributions of these measured features (bottom) is fundamentally determined by quantum sampling noise, which depends on the quantum state $\hat{\rho}(\mathbf{u}; \theta)$ and, hence, on the nature of the mapping from input \mathbf{u} to this state. We show four calculated distributions differing only in the values of inputs \mathbf{u} to highlight this dependence.

the input \mathbf{u} could, for instance, be embedded as a collection of pixel values on a spatial light modulator (SLM) in the path of a propagating beam of light. The individual single-shot features $\{X_k^{(s)}(\mathbf{u})\}$ could be generated by integrating the photocurrent from each pixel of a number-resolving CCD camera for a certain hold time. For a *biological neural circuit*, the input \mathbf{u} might be a static-in-time visual stimulus, representing the electromagnetic (EM) field intensity incident on photoreceptors in the eye, and $\{X_k^{(s)}(\mathbf{u})\}$ can be the action potential of the k th neuron integrated over a certain time period, e.g., measured through Ca^{2+} imaging [33]. Finally, for a *superconducting quantum processor*, inputs may be embedded via a suitable quantum channel, implemented, for example, via parametrized quantum gates. The single-shot features are simply the indicator functions of the possible outcome labels after quantum measurement. In all cases, the measured features $\bar{X}_k(\mathbf{u})$ may be obtained by repeating each experiment S times with the same \mathbf{u} and constructing the S -shot histogram.

The randomness of the measured features derives from the quantum mechanical or the thermodynamical nature of the processes that the physical system is subject to during its evolution but, more importantly, in the measurement and detection phase. In the case of neural circuits, for instance, even when great care is exercised by presenting identical stimuli, the timing of action potentials of individual neurons can vary significantly over repeated trials on a scale that can be physiologically relevant. This noise can be

traced to various sources [31] including dynamical changes of internal states of neurons between trials and random processes neurons are subject to. The source of sampling noise for the optical system discussed in Sec. IID is the shot noise related to the discrete nature of energy exchange between the EM field and the photodetector, an electronic system. For an ideal quantum computing system, the noise process we consider is due to shot noise in projective measurement, which we refer to as quantum sampling noise. We note that in qubit systems there are many other potential noise sources, but in modern quantum processors these ought to be subleading at least for shallow circuits. Indeed, in experiments reported in Sec. IV B, we observe that sampling noise dominates even at the maximum available S . Quantum sampling noise is still the limiting source of noise after the advancement of error-corrected quantum computers.

A last important source of noise is the noise in the input signal to be processed. Visual neural circuits, for instance, involve the absorption of photons that arrive at the photoreceptors from EM sources that are subject to quantum mechanical or thermodynamical fluctuations. Here, we are not concerned with a precise description of the physical nature of the input stimuli and account for it by assuming an underlying probability distribution $p(\mathbf{u})$ from which the inputs are sampled. The most complete treatment of such a process requires a quantum mechanical description of both the signal generating system and its coherent coupling to the physical system that processes it, as introduced and analyzed in Ref. [34].

B. Resolvable expressive capacity and eigentasks

Returning to the situation depicted in Fig. 1, supervised learning in physical systems can generically be cast as encoding data in the system and then using measurement outputs to approximate a desired function $f(\mathbf{u})$ (here assumed to be square-integrable $\mathbb{E}_{\mathbf{u}}[f^2] < \infty$), where the expectation over input data $\mathbb{E}_{\mathbf{u}}$ is defined with respect to the distribution $p(\mathbf{u})$: $\mathbb{E}_{\mathbf{u}}[f] \equiv \int d\mathbf{u} p(\mathbf{u}) f(\mathbf{u})$. The introduction of the symbol $\mathbb{E}_{\mathbf{u}}$ for expectation over \mathbf{u} is necessitated by the use of two types of averages in the analysis of the loss function: over the output samples ($\mathbb{E}_{\mathcal{X}}$) and over the input domain ($\mathbb{E}_{\mathbf{u}}$).

Within the PNN approach considered here, $f(\mathbf{u})$ is approximated for finite S as $f_{\mathbf{W}}(\mathbf{u}) = \mathbf{W}^T \bar{\mathbf{X}}(\mathbf{u}) = \sum_k W_k \bar{X}_k(\mathbf{u})$. To quantify the fidelity of this approximation, we introduce a statistical variant of the function capacity [18,22,35], which is the normalized mean-squared accuracy of the estimate $f_{\mathbf{W}}$:

$$C[f] = 1 - \min_{\mathbf{W} \in \mathbb{R}^K} \frac{\mathbb{E}_{\mathbf{u}}[\mathbb{E}_{\mathcal{X}}[(f(\mathbf{u}) - f_{\mathbf{W}}(\mathbf{u}))^2]]}{\mathbb{E}_{\mathbf{u}}[f(\mathbf{u})^2]}. \quad (5)$$

This quantity differs from that introduced in Refs. [18,22,35] in that the squared error term $(f(\mathbf{u}) - f_{\mathbf{W}}(\mathbf{u}))^2$ is stochastic; thus, both the expectation over the output samples \mathcal{X} and the

expectation over the inputs \mathbf{u} are needed to ensure that Eq. (5) is a deterministic value. Minimizing the error in the approximation of $f(\mathbf{u})$ by $f_{\mathbf{w}}(\mathbf{u})$ over the input domain to determine capacity therefore requires finding

$$\mathbf{w} = \operatorname{argmin}_{\mathbf{w} \in \mathbb{R}^K} [\mathbb{E}_{\mathcal{X}}[(f - \mathbf{W}^T \bar{\mathbf{X}})^2]]. \quad (6)$$

This minimization can always be expressed analytically via a pseudoinverse operation (see Appendix C 1). This function capacity is constructed such that $0 \leq C[f] \leq 1$, with the upper limit indicating a perfect approximation.

The choice of a linear estimator and a mean-squared error loss function may appear restrictive at first glance, but the generality of our formalism averts such limitations. The use of a linear estimator applied directly to readout features appears to preclude nonlinear postprocessing of measurements; this is intentional and simply meant to ensure the calculated function capacity is a measure of the ability of the physical system itself and not of a nonlinear processing layer. Furthermore, the mean-squared loss effectively describes the first term in a Taylor expansion of a wide range of arbitrary nonlinear postprocessing and nonquadratic loss functions. The most well-known example is that of logistic regression for supervised classification problems, where the sigmoid function $\sigma(\mathbf{W}^T \bar{\mathbf{X}})$ [i.e., $\sigma(z) = 1/(1 + \exp(-z))$] is used for postprocessing, while the cross-entropy loss function is used for optimization (for further details and analysis of nonlinear postprocessing, see Appendix C 5).

To extend the notion of capacity to a task-independent metric representing how much classical information about an input can be extracted from a system in the presence of sampling noise, we sum the function capacity over a basis of functions $\{f_{\ell}\}_{\ell \in \mathbb{N}}$ which are complete and orthonormal with respect to the input distribution, i.e., equipped with the inner product $\langle f_{\ell}, f_{\ell'} \rangle_p = \int f_{\ell}(\mathbf{u}) f_{\ell'}(\mathbf{u}) p(\mathbf{u}) d\mathbf{u} = \delta_{\ell\ell'}$. The total REC is then $C_T \equiv \sum_{\ell=0}^{\infty} C[f_{\ell}]$, which effectively quantifies how many linearly independent functions can be expressed from a linear combination of $\{\bar{X}_k(\mathbf{u})\}$. Our main result—proven in detail in Appendix C 4—is that, given any $S \in \mathbb{N}^+$, the REC for a physical system whose measured features are stochastic variables of the form of Eq. (3) is given by

$$C_T(\boldsymbol{\theta}) = \operatorname{Tr} \left(\left(\mathbf{G} + \frac{1}{S} \mathbf{V} \right)^{-1} \mathbf{G} \right) = \sum_{k=0}^{K-1} \frac{1}{1 + \beta_k^2(\boldsymbol{\theta})/S}. \quad (7)$$

Here, we make explicit the dependence on $\boldsymbol{\theta}$, the hyperparameters of the input embedding to indicate the important dependence of the S -shot REC on the input encoding.

The first equality, arrived at through straightforward algebraic manipulation, is written in terms of the expected feature Gram and covariance matrices $\mathbf{G} \equiv \mathbb{E}_{\mathbf{u}}[\mathbf{x}\mathbf{x}^T]$ and $\mathbf{V} \equiv \mathbb{E}_{\mathbf{u}}[\boldsymbol{\Sigma}]$, respectively. First, we are able to conclude that $\lim_{S \rightarrow \infty} C_T = \operatorname{Rank}\{\mathbf{G}\} \leq K$, recovering the bound in

Ref. [18]. Importantly, the rank of the Gram matrix is always equal to the maximal number of linearly independent functions in the set $\{x_k(\mathbf{u})\}$ (see Appendix C 2). In this article, we consider only the case where \mathbf{G} is full rank, which is the most interesting case: maximizing the rank of \mathbf{G} maximizes the highest achievable (i.e., infinite- S) REC for a physical system. Furthermore, this condition is typically met unless the physical system is constrained by special symmetries; in such cases where some features $\{x_k(\mathbf{u})\}$ are linearly dependent, the matrix inverse in Eq. (7) should be modified to a pseudoinverse. We also later demonstrate that both \mathbf{G} and \mathbf{V} can be estimated efficiently and accurately in experiment and, consequently, under finite S (see Appendix D). The second equality in Eq. (7) remarkably provides a closed-form expression for C_T at any S , which is independent of the specific choice of the generally infinite set $\{f_{\ell}\}_{\ell \in \mathbb{N}}$ (and, thus, not subject to numerical challenges associated with its evaluation over such a set [18]). Instead, the REC is entirely captured by the function capacity of K distinct functions and for a given physical system is fully characterized by the spectrum of eigenvalues $\{\beta_k^2\}_{k \in [K]}$ satisfying the generalized eigenvalue problem

$$\mathbf{V}\mathbf{r}^{(k)} = \beta_k^2 \mathbf{G}\mathbf{r}^{(k)}. \quad (8)$$

In the above, all quantities depend on $\boldsymbol{\theta}$ and, therefore, on the specific physical system and input embedding via the Gram (\mathbf{G}) and covariance (\mathbf{V}) matrices. Associated with each β_k^2 is an eigenvector $\mathbf{r}^{(k)}$ living in the space of measured features, defining a set of K orthogonal functions via the linear transformation

$$y^{(k)}(\mathbf{u}) = \sum_j r_j^{(k)} x_j(\mathbf{u}). \quad (9)$$

We refer to $\{y^{(k)}\}$ as *eigentasks*, as they form the minimal set of orthonormal functions ($\mathbb{E}_{\mathbf{u}}[y^{(j)} y^{(k)}] = \delta_{jk}$) which saturates the available REC of a physical system and, thus, the accessible information content present in its measured features. Specifically, the capacity to approximate a given $y^{(k)}$ with S shots is $C[y^{(k)}] = 1/(1 + \beta_k^2/S)$: the REC in Eq. (7) is simply a sum of eigentask capacities. This further highlights that a given parametrized system can only approximate a target function to the degree that this function can be written as a linear combination of $\{y^{(k)}\}$. The eigentasks therefore serve as a powerful basis for learning, as explored in Sec. IV C.

C. Resolvable expressive capacity and eigentasks in practice: Measured eigentasks

Our use of the expectation over distributions of the input, $\mathbb{E}_{\mathbf{u}}[\cdot]$, and finitely sampled measured features, $\mathbb{E}_{\mathcal{X}}[\cdot]$, in principle implies the availability of an infinite number of

input and measured samples, respectively. Of course, for the practical implementation of any PNN, both these values are finite. However, as we demonstrate via calculations of the REC and eigentasks using both theoretical and experimental systems, our framework can be applied when these values are constrained to be finite.

More precisely, we note that, in practice, only a finite number of values N can be IID sampled from the input distribution; namely, $\mathbf{u}^{(n)} \sim p(\mathbf{u})$ for any $n \in [N]$. For each discrete input, one set of measured output features constructed from finite S is obtained, a single sample from the distribution $\mathcal{X}(\mathbf{u}^{(n)})$. The collection of both input and output samples constitutes the *complete dataset*, which we denote as $\mathcal{D} \equiv \{(\mathbf{u}^{(n)}, \mathcal{X}(\mathbf{u}^{(n)}))\}_{n \in [N]}$. Our calculation of REC and eigentasks has some dependence on \mathcal{D} via N and S .

In particular, the practically computed optimal weights in the capacity calculation are not the deterministic weights \mathbf{w} but \mathbf{w}^* computed on a given set of input samples and measured features and, hence, depend on the dataset \mathcal{D} :

$$\mathbf{w}^*(\mathcal{D}) \equiv \underset{\mathbf{w} \in \mathbb{R}^k}{\operatorname{argmin}} \frac{1}{N} \sum_{n=1}^N (f(\mathbf{u}^{(n)}) - \mathbf{W}^T \bar{\mathbf{X}}(\mathbf{u}^{(n)}))^2. \quad (10)$$

Weights $\mathbf{w}^*(\mathcal{D})$ vary due to changes in \mathcal{D} . Generally, when N and S are simultaneously finite and S is fixed, a study of the N -scaling behavior of the difference between \mathbf{w} in Eq. (6) and the average optimized weight $\mathbb{E}_{\mathcal{D}}[\mathbf{w}^*(\mathcal{D})]$, as well as the variation of $\mathbf{w}^*(\mathcal{D})$ for different \mathcal{D} , falls in the realm of training and generalization errors over the input domain, an important area of research in theoretical machine learning [36,37]. We leave this problem for future work; for all calculations and experiments in this paper, we consider the case—always realized in practice (and of particular relevance where *sampling*, and, thus, the time and resource cost of processing with physical systems, is concerned)—where the dataset consists of a finite number N of input samples.

We do address the important problem of REC and eigentask calculation when this fixed value of N is finite and using only a given set of measured features constructed under finite sampling S . As alluded to earlier, in Appendix D, we demonstrate how the eigenproblem Eq. (8) can be constructed for finite N and S and present corrections to the eigenvalues and eigenvectors due to the finiteness of S . Numerical examples presented in Appendix D demonstrate a favorable match between this correction method and numerical simulations of eigenvalues and eigenvectors (see Figs. 8 and 9).

Importantly, we define a set of *measured* eigentasks $\bar{\mathbf{y}}^{(k)}(\mathbf{u}) = \sum_j r_j^{(k)} \bar{\mathbf{X}}_j(\mathbf{u})$ constructed from a given set of measured features. For these measured eigentasks, we find (see Appendix C 3) that $\{\mathbf{r}^{(k)}\}$ specify a unique linear transformation that simultaneously orthogonalizes not only

the signal, but also the associated noise: $\mathbb{E}_{\mathbf{u}}[\mathbb{E}_{\mathcal{X}}[\bar{\mathbf{y}}^{(j)} \bar{\mathbf{y}}^{(k)}]] = \delta_{jk}(1 + \beta_k^2/S)$. The term β_k^2/S is therefore the mean-squared error, or noise power, associated with the approximation of eigentask $\mathbf{y}^{(k)}$; equivalently, $\bar{\mathbf{y}}^{(k)}$ has a signal-to-noise ratio of S/β_k^2 . This leads to a natural interpretation of $\{\beta_k^2\}$ as noise-to-signal (NSR) eigenvalues. The eigentasks, ordered in increasing noise strength $0 \leq \beta_0^2 \leq \beta_1^2 \leq \dots \leq \beta_{K-1}^2 < \infty$, are the orthogonal set of functions maximally robust to sampling noise.

D. Example: Resolvable expressive capacity and eigentasks for a classical optical learning system

Before presenting more involved examples of physical quantum systems, we discuss an example of the presented framework for noisy *classical* dynamical systems, within the popular PNN platform of photonic ELM [8,38] and reservoir computing (RC) [5,39]. The specific setup we consider is illustrated in Fig. 3(a), where computation of inputs \mathbf{u} is performed via the encoding, propagation, and measurement of propagating EM waves in a medium. Here, the entire 3D space is defined by coordinates (q^1, q^2, q^3) , and EM fields of wavelength λ propagate in the q^3 direction. The electric field distribution is then completely defined by the position vector \vec{d} defined in the plane orthogonal to the propagation direction, so that $\vec{d} = (q^1, q^2)$ [40].

The input embedding of \mathbf{u} is performed using a spatial light modulator that modulates the amplitude and/or phase of the electric field of the radiation as it passes through. We restrict this example to 1D inputs u that are uniformly distributed: $p(u) = \operatorname{Unif}[-1, 1]$. The scalar u is then mapped to all the pixels of the SLM through a specific mapping discussed in Appendix H 1. We consider this rather artificial input encoding for two reasons: for ease of visualization of the computed eigentasks (see Fig. 3) and to ensure the distribution is sufficiently sampled. In the simulation of the classical optical system we consider here, we choose $N = 300$. This is also the number of input samples used in our analysis of qubit-based quantum systems in Sec. IV.

The spatial profile of the electric field $E_0(u; \vec{d})$ following the SLM can be written generally in the form $E_0(u; \vec{d}) = A_0 \cos(\varphi_1(u; \vec{d})/2) \exp\{i(\varphi_1(u; \vec{d}) + 2\varphi_2(u; \vec{d})/2)\}$, where A_0 is the initial electric field amplitude and $\varphi_l(u; \vec{d})$ are input encoding functions [41] [cf. Eqs. (H2a) and (H2b)]. Following the input encoding, the radiation propagates through free space and then past a thin lens. The electric field in the focal plane of the lens, $E(u; \vec{d})$, can be shown to be related to the initial field $E_0(u; \vec{d})$ via a Fourier transform [42,43]: $E(u; \vec{d}) = \iint d^2\vec{d}' E_0(u; \vec{d}') \exp\{(i2\pi/\lambda f) \times (\vec{d} \cdot \vec{d}')\}$, where f is the focal length of the lens. The choice of the optical propagation medium as a lens is again for

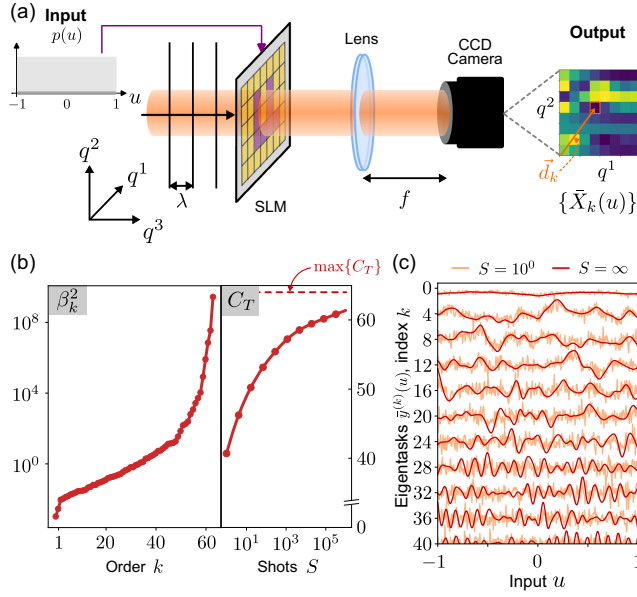


FIG. 3. (a) A photonic learning system we consider to demonstrate the REC analysis and eigentask construction in a classical optics setup. Inputs u are encoded into the amplitudes and/or phases of propagating light fields of wavelength λ via an SLM. The propagating fields are brought to focus via a converging lens of focal length f , and the fields are then imaged in the focal plane of the lens using a camera consisting of an array of photodetectors. Output features are given by the integrated stochastic photocurrents measured via these photodetectors. (b) Left: noise-to-signal spectrum β_k^2 as a function of k . Right: C_T vs S calculated from the left. (c) Eigentasks under infinite sampling ($S \rightarrow \infty$) and measured eigentasks constructed from features under finite sampling ($S = 10^0$).

convenience of analysis, not a limitation. More complex optical systems can be analyzed using the same techniques outlined here.

Finally, output features are extracted via photodetection (using a CCD camera) in the focal plane of this lens. Modeling this effectively requires us to address the important question of the measurement noise associated with photodetection. First, we consider the camera plane as being comprised of a discrete set of $K = P^2$ photodetectors (here, $P = 8$), arranged in a P -by- P square spatial grid, such that the k th photodetector is identified with coordinates $\vec{d}_k = (q_k^1, q_k^2)$ and $k \in [K]$ [see Fig. 3(a)]. This spatial grid ultimately defines the coarse-graining level at which the propagating fields can be probed and is set by the spatial resolution of the photodetection apparatus, as expected. Then the differential, stochastic photocurrent generated in a given photodetector in a single measurement—namely, the increment in photodetector counts in the time window $[t, t + dt]$, which we denote as $dN(\vec{d}_k, t)$, follows a Poisson point process (commonly referred to as shot noise) [44]. This instantaneous photocurrent $dN(\vec{d}_k, t)$ is often integrated over a finite time T_{int} in so-called integrate-and-dump

photodetectors. This defines the single-shot features of this RC scheme:

$$X_k^{(s)} = \int_0^{T_{\text{int}}} dN^{(s)}(\vec{d}_k, t) \in \mathbb{N}. \quad (11)$$

Note that $X_k^{(s)}$ are stochastic, integer quantities, simply counting the total number of photogenerated carriers in a time window T_{int} of a single measurement.

We now make the important physical connection between the measured photocurrents and the propagating fields reaching the photodetector. The power incident on the k th photodetector is simply set by the Poynting flux of the propagating fields and is proportional to the electric field intensity $\alpha|E(u; \vec{d}_k)|^2$, where α is a dimensionful constant that depends, for example, on the speed of light through the medium of traversal. Then, the expected value of the photocurrent in a time interval dt is simply proportional to the incident power, up to a factor η that encapsulates the efficiency of photodetection: $\mathbb{E}[dN(\vec{d}_k, t)] = \eta\alpha|E(u; \vec{d}_k)|^2 dt$.

The complete input-output map defined above fits within our very general framework. In particular, we define measured features $\bar{X}_k(u)$ as S -shot sample means of $\{X_k^{(s)}\}$, as in Eq. (1). Note that in most classical PNN schemes $S = 1$; here, we consider the shot number $S \geq 1$. We then express the measured features in terms of the decomposition in Eq. (3). First, using Eq. (11) and the definition of $x_k(u)$ in Eq. (2), we find

$$x_k(u) = \eta\alpha|E(u; \vec{d}_k)|^2 T_{\text{int}}. \quad (12)$$

Then, the remaining term in Eq. (3), $(1/\sqrt{S})\zeta_k(u)$, is a stochastic process with zero mean, and its second-order moment encodes the variance of the Poisson point process in one shot of the experiment, namely, that its variance is equal to its mean (see Appendix H for details):

$$\Sigma_{jk}(u) = \delta_{jk}x_k(u). \quad (13)$$

The form of the covariance matrix here is specific to the Poisson nature of the noise process inherited from the classical nature of the source (e.g., a coherent light source such as a laser) generating the beam of light. Other types of noise processes yield distinct covariance matrices, as we see in examples of quantum systems.

This is an appropriate place to remark that, for such a classical description of a physical system for learning, the stochastic photocurrent dN for any shot is specified by the *deterministic* electric field incident on the camera plane. In a fully quantum description, on the other hand, the power incident on the photodetectors in a given shot is determined by the expectation value of the field excitation number operator $\hat{\Psi}^\dagger \hat{\Psi}$ (in second-quantized notation) with respect

to the *conditional* density matrix $\hat{\rho}^{(c)}$, describing the measurement-conditioned state of the propagating radiation field for that shot. We are compelled to consider a description of this sort when describing physical *quantum* systems in Sec. III.

Equations (12) and (13) are sufficient to calculate the feature Gram and covariance matrices \mathbf{G} and \mathbf{V} , respectively, as per the discussion following Eq. (7). Therefore, we are set up to solve the eigenproblem of Eq. (8) and obtain the NSR spectrum and eigentasks for this toy model of a classical optical PNN. We first present the spectrum of NSR eigenvalues β_k^2 in Fig. 3(b). The NSR spectrum allows calculation of the REC as a function of S using Eq. (7); this is shown in Fig. 3(c). At finite S , we clearly observe that the REC remains below its upper bound of $K = 64$, approaching it only when S is increased, reducing the impact of sampling noise on measured features.

Finally, we discuss eigentask construction, also obtained by solving the eigenproblem of Eq. (8). In Fig. 3(c), we visualize as a function of u a selection of both the $S \rightarrow \infty$ eigentasks $y^{(k)}(u)$ defined in Eq. (9) and the measured eigentasks $\bar{y}^{(k)}(u)$ obtained from $S = 10^0$ sampled features (i.e., single-shot). We note that the sampled features are obtained by numerically integrating the stochastic differential equation defining independent measurements of the stochastic photocurrents in Eq. (11). These measured eigentasks exhibit sampling noise, which is evident when compared against the infinite shot eigentasks. We clearly see that eigentasks which are higher order in k are increasingly noisy across the input domain, as also encapsulated by the larger associated NSR eigenvalues. The ordered eigentasks, therefore, represent the functions that are optimally resolvable using this classical optical setup in the presence of the sampling noise that it is naturally, and unavoidably, subject to for finite S .

III. LEARNING WITH QUANTUM SYSTEMS

A. Sampling noise in quantum systems

Having developed our framework for REC in the most general context, in the remainder of this paper, we use it to analyze quantum systems in greater depth. The same quantitative metrics—REC, eigentasks, and NSR eigenvalues—now carry the significance of being determined by a parametrized quantum state. To be more specific, the classical data \mathbf{u} are now encoded through a quantum channel parametrized by $\boldsymbol{\theta}$ acting on a known initial state:

$$\hat{\rho}(\mathbf{u}; \boldsymbol{\theta}) = \mathcal{U}(\mathbf{u}; \boldsymbol{\theta})\hat{\rho}_0, \quad (14)$$

whose data dependence may be hard to model classically. The quantum channel \mathcal{U} includes all quantum operations applied to the input data; to obtain the computational output or perform further classical processing, one must extract information from the quantum system via a set of

measurements described most generally as a POVM. Specifically, we define a set of K POVM elements $\{\hat{M}_k\}$, each associated with a distinct measurement outcome indexed k , and constrained only by the normalization condition $\sum_{k=0}^{K-1} \hat{M}_k = \hat{\mathbf{I}}$ (and, hence, not necessarily commuting).

Each shot then yields a discrete index $k^{(s)}(\mathbf{u})$ specifying the observed outcome: for input \mathbf{u} , if outcome k is observed in shot s , then $k^{(s)}(\mathbf{u}) \leftarrow k$. In this case, the single-shot random-valued feature $X_k^{(s)}(\mathbf{u})$ is exactly the indicator $\delta(k^{(s)}(\mathbf{u}), k)$ of index k , so that the *measured features* are given by

$$\bar{X}_k(\mathbf{u}) = \frac{1}{S} \sum_s \delta(k^{(s)}(\mathbf{u}), k). \quad (15)$$

Hence, $\bar{X}_k(\mathbf{u})$ in this case is the empirical frequency of occurrence of the outcome k in S repetitions of the experiment with the same input \mathbf{u} . These measured features are formally random variables that are unbiased estimators of the expected value of the corresponding element \hat{M}_k as computed from $\hat{\rho}(\mathbf{u})$. Explicitly,

$$x_k(\mathbf{u}) = \text{Tr}\{\hat{M}_k \hat{\rho}(\mathbf{u}; \boldsymbol{\theta})\}, \quad (16)$$

so that x_k is the probability of occurrence of the k th outcome as specified by the quantum state. These probability amplitudes encompass the accessible information in $\hat{\rho}(\mathbf{u}; \boldsymbol{\theta})$: any observable under this set can be written as a linear combination of POVM elements $\hat{O}_W = \sum_k W_k \hat{M}_k$, such that $\langle \hat{O}_W \rangle = \mathbf{W}^T \mathbf{x}$.

In quantum machine learning (QML) theory, it is standard to consider the limit $S \rightarrow \infty$ and to therefore use expected features $\{x_k(\mathbf{u})\}$ for learning. In any actual implementation, however, measured features $\{\bar{X}_k(\mathbf{u})\}$ must be constructed under finite S , in which case their fundamentally quantum-stochastic nature can no longer be ignored. The decomposition Eq. (3) is still applicable $\bar{\mathbf{X}}(\mathbf{u}) = \mathbf{x}(\mathbf{u}) + \boldsymbol{\zeta}(\mathbf{u})/\sqrt{S}$, where now \mathbf{x} are the quantum-mechanical event probabilities and $\boldsymbol{\zeta}$ encodes the multinomial statistics of quantum sampling noise, whose covariance is explicitly

$$\boldsymbol{\Sigma}_{jk}(\mathbf{u}) = \delta_{jk} x_k(\mathbf{u}) - x_j(\mathbf{u}) x_k(\mathbf{u}). \quad (17)$$

This is simply the expression for the covariance of multinomial distribution with S trials and K mutually exclusive outcomes with probabilities $p_k = x_k$. For arbitrary orders of cumulants of multinomial statistics, we refer to Ref. [45]. For the quadratic loss function considered here, only the cumulants up to second order turn out to be sufficient.

One may wonder what specifically distinguishes a quantum system from a classical stochastic system that can generate a multinomial distribution in its output. First, certain $\{p_k(u)\}$ combinations can be generated efficiently

by only a quantum system. That is to say, given equal resources, a quantum system can access some $\{p_k\}$ that may be inaccessible to any classical stochastic system and hence, as discussed later, the accessible space of functions is far richer. However, we also find that resolvability of that function space in S measurements is the key determinant in learning. Note that all statistical properties of stochastic readout features $\bar{X}(\mathbf{u})$ —namely, first-order cumulants $\mathbf{x}(\mathbf{u})$, second-order cumulants $\Sigma(\mathbf{u})$, and all higher-order cumulants—are determined fully by the quantum state $\hat{\rho}(\mathbf{u})$, which itself may be hard to generate classically.

To proceed with our REC analysis in quantum systems, we write down the generalized eigenproblem Eq. (8) by computing $\mathbf{G} = \mathbb{E}_u[\mathbf{x}\mathbf{x}^T]$ and $\mathbf{V} = \mathbb{E}_u[\Sigma]$. Equation (17) enables us to simplify the exact form of \mathbf{V} , namely, $\mathbf{V} = \mathbf{D} - \mathbf{G}$, where $\mathbf{D} \in \mathbb{R}^{K \times K}$ is a diagonal matrix with elements $\mathbf{D}_{kk} = \mathbb{E}_u[x_k]$. Alternatively, for any encoding state ensemble $\{p(\mathbf{u})d\mathbf{u}, \hat{\rho}(\mathbf{u})\}$, the matrices \mathbf{D} and \mathbf{G} can be compactly expressed as (see Appendix C 1)

$$\mathbf{D}_{kk} = \text{Tr}\{\hat{M}_k \hat{\rho}^{(1)}\}, \quad (18)$$

$$\mathbf{G}_{jk} = \text{Tr}\{(\hat{M}_j \otimes \hat{M}_k) \hat{\rho}^{(2)}\} \quad (19)$$

by defining the t th-order *quantum ensemble moment* $\hat{\rho}^{(t)} = \int \hat{\rho}(\mathbf{u})^{\otimes t} p(\mathbf{u}) d\mathbf{u}$ in the t -copy space of the quantum state [46].

From Eq. (7), we have $\lim_{S \rightarrow \infty} C_T = \text{Rank}\{\mathbf{G}\}$, where $\text{Rank}\{\mathbf{G}\} = K$, the number of measured features, provided no special symmetries exist. This important result reveals that, in the absence of sampling noise, all quantum systems— independent of parametrization—have a capacity which is simply the number of independent accessible degrees of freedom [18,47]. The generic exponential scaling of measured degrees of freedom with the size of the quantum system (e.g., $K = 2^L$ for L -qubit systems subject to a computational basis measurement) is often cited as a motivator for studying ML with quantum systems [22,35,48]. However, as demonstrated shortly, the REC of quantum systems can be significantly reduced from this limit for finite S in a way that strongly depends on the encoding. By evaluating the ability of quantum systems to accurately express functions in the presence of quantum sampling noise, the capacity analysis above provides an important metric to assess the utility of quantum platforms for learning in practice.

B. Resolvable expressive capacity of quantum two-designs

We first consider the REC of quantum two-designs: systems with fixed θ that map inputs to a unitary ensemble $\{p(\mathbf{u})d\mathbf{u}, \hat{U}(\mathbf{u}; \theta)\}$ whose first and second moments agree with those from a uniform (Haar) distribution of unitaries. Quantum two-designs are important to recent QML studies [7,49] due to their role in defining and

studying “expressibility” [21,50]: a metric quantifying how close a parametrized quantum system is to such a two-design. The capacity eigenproblem Eq. (8) for any quantum two-design over K dimensions can be solved analytically (see Appendix F), yielding a flat spectrum of NSR eigenvalues $\beta_k^2 = K(1 - \delta_{k0})$. This results in an REC:

$$C_T = K \cdot \frac{S + 1}{S + K}, \quad (20)$$

which at finite S can be significantly lower than K . For quantum systems with $K = 2^L$, all $k \neq 0$ eigentasks have a noise strength $2^L/S$, requiring S to grow exponentially with qubit number L in order to extract useful features.

A quantum two-design is thought of as having maximal “expressibility”; however, we see that its REC always vanishes exponentially with system size for a fixed finite S . To emphasize the distinction with expressibility, we note that REC reflects how much classical information can be extracted from the entire “quantum computational stack” in practice: from an abstract algorithm, to the quantum hardware on which its implemented, and the classical electronics used for control and readout. REC requires only noisy computational outputs $\{\bar{X}_k(\mathbf{u})\}$ and is therefore efficiently computable in experiment—unlike more abstract metrics [21,50,51]—yielding a directly relevant metric for learning with quantum hardware.

IV. EXPERIMENTAL RESULTS IN QUANTUM SYSTEMS

In this section, we discuss the implementation of the eigentask construction in experiments we carry out on a seven-qubit IBMQ superconducting quantum processor *ibmq_perth*.

A. The quantum circuit *Ansatz* implemented in experiments

To demonstrate the practical utility of our framework, we now show how the spectrum $\{\beta_k^2\}$, the REC, and eigentasks can all be computed for real quantum devices in the presence of parameter fluctuations and device noise. We reiterate at the outset that our approach for quantifying the REC of a quantum system is very general and can be applied to a variety of quantum system models. For practical reasons, we perform experiments on L -qubit IBM Quantum (IBMQ) processors, whose dynamics is described by a parametrized quantum circuit containing single- and two-qubit gates. However, as an example of the broad applicability of our approach, in Appendix E, we compute the REC for L -qubit quantum annealers via numerical simulations, governed by the markedly different model of continuous-time Hamiltonian dynamics.

On IBMQ devices, each input u generates a quantum circuit; hence, the maximal number of distinct circuits places a resource constraint on input size. Especially, our

experiment and computation is limited to $N = 300$ 1D inputs u that are also uniformly distributed, $p(u) = \text{Unif}[-1, 1]$; see Fig. 4(a). A 1D distribution then ensures features $\{\bar{X}_k(u)\}$ are sufficiently densely sampled to approach the continuum limit and are also easy to visualize, as in the classical optical RC in Sec. IID. We emphasize that this analysis can be straightforwardly extended to multidimensional and arbitrarily distributed inputs given suitable hardware resources, without modifying the form of the Gram and covariance matrices.

We are finally required to specify the model of the quantum system and choose an *Ansatz* tailored to be natively implementable on IBMQ processors (see Appendix B). We fix $\hat{\rho}_0 = |0\rangle\langle 0|^{\otimes L}$; note, however, that any other initial state may be implemented via an additional unitary and absorbed

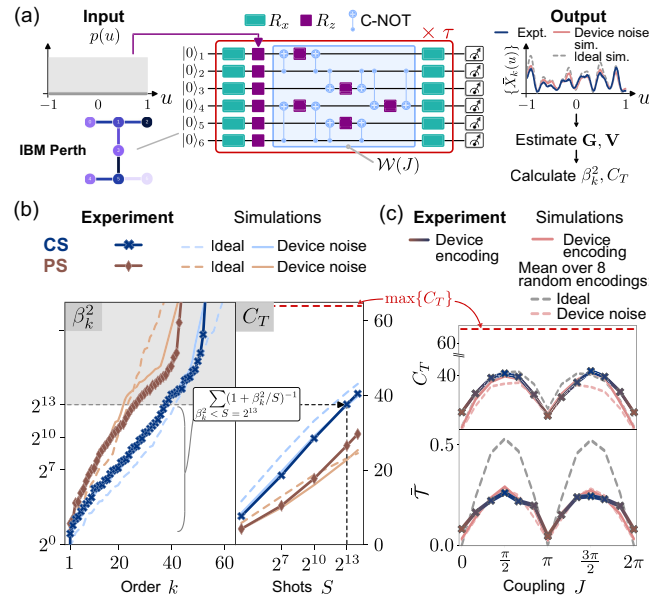


FIG. 4. (a) A representation of the REC analysis, featuring the IBMQ Perth device and a schematic of the quantum circuit considered in this section. On the right, the specific feature plotted is $\bar{X}_1(u)$ ($\mathbf{b}_1 = 000\,001$) with $S = 2^{14}$ shots. (b) Left: device noise-to-signal spectrum β_k^2 for a specific encoding as a correlated system (CS), $J = \pi/2$ (blue crosses) and product system (PS), $J = 0$ (brown diamonds). Ideal (solid line) and device noise (dashed line) simulations are also shown. Note the agreement between device and simulation, along with distortion from more direct exponential growth in β_k^2 with k in the ideal case, due to device errors. Right: C_T vs S calculated from the left. At a given S , the C_T can be approximated by performing the indicated sum over all $\beta_k^2 < S$, where $\bar{\beta}_k^2$ denotes the estimate of β_k^2 computed from noisy data. (c) Resolvable expressive capacity C_T (top) and expected total correlation \bar{T} (bottom) for the chosen encoding under $S = 2^{14}$ from the IBMQ device and device noise simulations (dashed peach line). Average metrics over eight random encodings for device noise (solid peach line) and ideal (solid gray line) simulations are also shown. The $S \rightarrow \infty$ resolvable expressive capacity of these encodings always attains the $\max\{C_T\} = 64$, indicated in dashed red.

into the encoding, i.e., the quantum channel $\mathcal{U}(u; \theta)$ in Eq. (14). In this way, the dependence of REC on initial states could be explored in future studies.

The circuit we choose consists of $\tau \in \mathbb{N}$ repetitions of the same input-dependent circuit block depicted in Fig. 4(a). The block itself is of the form $\mathcal{R}_x(\theta^x/2)\mathcal{W}(J) \times \mathcal{R}_z(\theta^z + \theta^l u)\mathcal{R}_x(\theta^x/2)$, where $\mathcal{R}_{x/z}$ are Pauli rotations applied qubitwise, e.g., $\mathcal{R}_z = \otimes_l \hat{R}_z(\theta_l^z + \theta^l u)$. A two-qubit coupling gate acts between physically connected qubits in the device and can be written as $\mathcal{W}(J) = \prod_{(l,l')} \mathcal{W}_{l,l'}(J) = \prod_{(l,l')} \exp\{-i(J/2)\hat{\sigma}_l^z \hat{\sigma}_{l'}^z\}$. Within the structure of this *Ansatz*, we choose all single-qubit rotation parameters randomly: $\theta_l^{x/z} \sim \text{Unif}[0, 2\pi]$ and $\theta_l^l \sim \text{Unif}[0, 10\pi]$, generally representing an untrained circuit. Each instance of random parameters, along with associated dissipative processes, specifies the quantum channel $\mathcal{U}(u; \theta)$ which we refer to as an encoding. We study the performance of an overall *Ansatz* by looking at the behavior averaged across encodings as hyperparameters such as J are varied. In this work, we also choose $\tau = 3$, which limits circuit depth and associated prevalence of gate errors while still generating a complex state with correlation generally distributed throughout all qubits.

Finally, we consider feature extraction via a computational basis measurement as is standard in quantum information processing: the POVM elements are the $K = 2^L$ projectors $\hat{M}_k = |\mathbf{b}_k\rangle\langle \mathbf{b}_k|$, where \mathbf{b}_k is the L -bit binary representation of the integer k . However, as with state preparation, measurements in any other basis can be (and, in practice, are) realized using an additional unitary prior to computational basis readout, whose effect can similarly be analyzed as part of the general encoding $\mathcal{U}(u; \theta)$.

Note that, for this *Ansatz*, the choice $J = 0 \pmod{\pi}$ yields either $\mathcal{W}_{l,l'}(J) = \hat{I}$ or $\hat{\sigma}^z \otimes \hat{\sigma}^z$, both of which ensure $\hat{\rho}(u)$ is a product state and measured features are simply products of uncorrelated individual qubit observables—equivalent to a noisy classical system. Starting from this *product system* (PS), tuning the coupling $J \neq 0 \pmod{\pi}$ provides a controllable parameter to realize a *quantum correlated system* (CS), for which the 2^L -dimensional multinomial distribution $\mathbf{x}(u)$ cannot be represented as a tensor product of L marginal binomial distributions on each qubit. In general, such nonproduct systems are intuitively expected to result in u -dependent quantum states which exhibit entanglement and can potentially be more difficult to describe classically. This control enables us to address a natural question regarding REC of quantum systems under finite S : what is the dependence of REC and realizable eigtasks on J and, hence, on quantum correlations?

B. Resolvable expressive capacity of quantum circuits

To perform the capacity analysis, one must extract measured features from the quantum system as the input u is varied, as exemplified in Fig. 4(a) for the IBMQ

ibmq_perth device. For comparison, we also show ideal-device simulations (unitary evolution, no device noise), where slight deviations are observed. The agreement with experimental results is improved when the effects of gate errors, readout errors, and qubit relaxation are included, hereafter referred to as “device noise” simulations, highlighting both the non-negligible role of device nonidealities and that our analysis incorporates them.

The measured features under finite S are used to estimate the Gram and covariance matrices and to, therefore, solve the eigenproblem Eq. (8) for NSR eigenvalues $\{\bar{\beta}_k^2\}$ and eigenvectors $\{\bar{r}^{(k)}\}$, as estimators of $\{\beta_k^2\}$ and $\{r^{(k)}\}$ [see Eqs. (D17) and (D18) in Appendix D 1 for detailed techniques]. Typical NSR spectra computed for a random encoding (i.e., set of rotation parameters) on the device are shown in Fig. 4(b), for $J = 0$ (PS) and $J = \pi/2$ (CS), together with corresponding spectra from device noise simulations, with which they agree well. We note that at lower k the device NSR eigenvalues are larger than those from ideal simulations and at larger k deviate from the direct exponential increase (with order) seen in ideal simulations. Both these effects are captured by device noise simulations as well and can, therefore, be attributed to device errors and dissipation. The NSR spectra, therefore, can serve as an effective diagnostic tool for quantum processors and encoding schemes.

The NSR spectra can be used to directly compute the REC of the corresponding quantum device for finite S , via Eq. (7). Practically, at a given S , only NSR eigenvalues $\bar{\beta}_k^2 \lesssim S$ contribute substantially to the REC. An NSR spectrum with a flatter slope, therefore, has more NSR eigenvalues below S , which gives rise to a higher capacity. Figure 4(b) shows that the CS generally exhibits an NSR spectrum with a flatter slope than the PS, yielding a larger capacity for function approximation across all sampled S .

To more precisely quantify the role of quantum correlations in REC, we introduce the *expected total correlation* (ETC) of the measured state over the input domain of u [52,53]:

$$\bar{T} = \mathbb{E}_u \left[\sum_{l=1}^L \mathcal{S}(\hat{\rho}_l^M(u)) - \mathcal{S}(\hat{\rho}^M(u)) \right], \quad (21)$$

where $\hat{\rho}^M(u) \equiv \sum_k \hat{\rho}_{kk}(u) |\mathbf{b}_k\rangle \langle \mathbf{b}_k|$ is the postmeasured state, $\mathcal{S}(\cdot)$ is the von Neumann entropy (see Appendix G), and $\hat{\rho}_l = \text{Tr}_{[L] \setminus \{l\}} \{\hat{\rho}\}$ is the reduced density matrix obtained by tracing over all qubits except qubit l . Therefore, nonzero ETC indicates the generation of quantum states over the input domain u that on average have nontrivial correlations amongst their constituents, including, for example, pure many-body states that are entangled.

We now compute REC and ETC using $S = 2^{14}$ in Fig. 4(c) as a function of J , for the same random encoding considered above on the device. We note that the

experimental results show excellent agreement in both cases with the corresponding device noise simulation. We also show average REC at $S = 2^{14}$ and ETC across eight random encodings in both ideal and device noise simulations. We find that the influence of individual encodings, i.e., random rotation parameters, leads only to small deviations from the overall REC trend when global hyperparameters are held fixed. This implies that no crucial features of the REC are missed by us foregoing fine-tuning (e.g., via gradient descent) of individual rotation parameters *in lieu* of sampling them from a given uniform probability distribution.

We note that product states by definition have $\bar{T} = 0$ [28]; this is seen in ideal simulations for $J = 0 \pmod{\pi}$. However, the actual device retains a small amount of correlation at this operating point, which is reproduced by device noise simulations. This can be attributed to gate or measurement errors as well as crosstalk, the latter being especially relevant for the transmon-based IBMQ platform with a parasitic always-on ZZ coupling [54]. With increasing J , \bar{T} increases and peaks around $J \approx \pi/2 \pmod{\pi}$; interestingly, C_T also peaks for the same coupling range. From the analogous plot of REC, we clearly see that, at finite S , increased ETC appears directly correlated with higher REC. We observe very similar behavior using completely different quantum system models (see Fig. 10 [55,56]). This indicates the utility of enhancing quantum correlations as a means of improving the general expressive capability of quantum systems.

We raise two notes of caution here. First, our analysis across different quantum system implementations often (though not always) finds that a certain threshold number of shots S is required before the finite- S capacity of a CS overtakes that of the corresponding PS (see Appendix E). This higher resolvability of functions using a PS under restricted shots may be due to the comparative ease of estimating probabilities from an effectively product distribution and merits further exploration. At a sufficiently large S , the increased complexity of the u dependence imposed by the input-output map of a CS results in an REC that eventually surpasses that of the PS.

Second, we caution that the connection between measurement correlations and REC is an observed trend rather than a law derived from first principles. One can come up with contrived situations where increasing correlation has no effect on REC: for example, appending a layer of CNOT gates directly prior to measurement generally increases the ideal ETC of any *Ansatz*. For measured features, however, this amounts to a simple shuffling of labels $x_k(\mathbf{u}) \leftrightarrow x_{k'}(\mathbf{u})$, thus yielding the same NSR spectrum and REC. The input, quantum state, and feature mapping ultimately governs REC: only increases in correlation that also increase the complexity of the measured features’ u dependence (as achieved via the intermediate \mathcal{W} gates here) are beneficial from the perspective of information processing.

As a final important point, note that at finite S , even with increased quantum correlations, the maximum REC is still substantially lower than the upper bound of $K = 64$. This remains true even for ideal simulations and over several random encodings, so the underperformance cannot be attributed to device noise or poor *Ansatz* choice, respectively. It is worth emphasizing that the impact of device noise is captured in the small REC gap between the ideal and noisy simulation curves, with the remainder of the reduction from $K = 64$ attributable to quantum sampling noise alone. These results clearly indicate that the resulting sampling noise at finite S is the fundamental limitation for QML applications on this particular IBM device rather than other types of noise sources and errors.

C. A robust approach to learning

While we demonstrate the REC as an efficiently computable metric of general expressive capability of a noisy quantum system, some important practical questions arise. First, does the general REC metric have implications for practical performance on *specific* ML tasks? Second, given the limiting—and unavoidable—nature of correlated sampling noise, does the REC provide any insights on optimal learning using a particular noisy quantum system and the associated encoding?

Our formulation addresses both these important questions naturally, as we now discuss. Recall that, beyond being a simple figure of merit, the REC is precisely the sum of capacities to approximate a particular set of orthogonal functions native to the given noisy quantum system: the eigentasks. Furthermore, these eigentasks $\bar{y}^{(k)}(u)$ can be directly estimated from a noisy quantum system via the generalized eigenvectors $\{\bar{r}^{(k)}\}$ and are ordered by their associated NSR eigenvalues $\{\bar{\beta}_k^2\}$. Figure 5(a) shows a selection of estimated eigentasks from the device for the CS ($J = \pi/2$) and PS ($J = 0$) encodings in Fig. 4(b). For both systems, the increase in noise with eigentask order is apparent when comparing two sampling values $S = 2^{10}$ and $S = 2^{14}$. Furthermore, for any order k , eigentasks for the PS are visibly noisier than the CS; this is consistent with NSR eigenvalues for PS being larger than those for CS [Fig. 4(b)]. The higher resolvable expressive capacity of the CS can be interpreted as the ability to accurately resolve more eigentasks at fixed S .

The resolvable eigentasks of a finitely sampled quantum system are intimately related to its performance at specific QML applications. To demonstrate this result, we consider a concrete application: a binary classification task that is not linearly separable. The domain $u \in [-1, 1]$ over which REC is evaluated is separated into two classes, as depicted in Fig. 5(b). A selection of $N_{\text{train}} = 150$ total samples—with equal numbers from each class—are input to the IBMQ device, and eigentasks $\{\bar{y}^{(k)}(u^{(n)})\}$ are estimated using $S = 2^{14}$ shots. A linear estimator applied to this set of

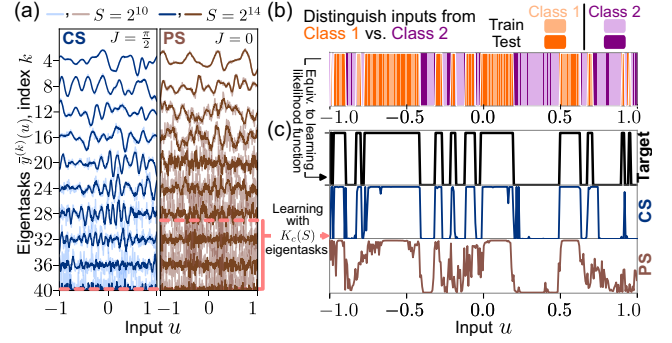


FIG. 5. (a) Device eigentasks for correlated system (CS, left) and product system (PS, right), constructed from noisy features at $S = 2^{10}$ and $S = 2^{14}$. (b) Classification demonstration on IBMQ Perth. Binary distributions to be classified over the input domain are shown. (c) The classification task can be cast as learning the likelihood function separating the two distributions; this target function is shown in the upper panel. Lower panels show the learned estimate of this target based on the $N_{\text{train}} = 150$ points shown in (b), using only $K_c(S)$ eigentasks for $S = 2^{14}$; this cutoff is indicated by the dashed red lines. For the correlated system $K_c(S) = 40$, while for the product system $K_c(S) = 29$.

eigentasks is then trained using logistic regression to learn the class label associated with each input. Finally, the trained IBMQ device is used to predict class labels of $N_{\text{test}} = 150$ distinct input samples for testing. Note that we use the *random* circuits of the previous section to draw more direct comparisons between REC and task performance. By training only external weights instead of internal parameters θ , we are employing the framework of quantum ELMs [6,10,22,57], which allows one to avoid the computational overhead and difficulty associated with training quantum systems while still achieving comparable performance.

This task can equivalently be cast as one of learning the likelihood function that discriminates the two input distributions, shown in Fig. 5(c), with minimum error. The set of up to K_L eigentasks $\{\bar{y}^{(k)}(u)\}_{k \in [K_L]}$, where $K_L \leq K$, serves as the native orthonormal basis of readout features used to approximate *any* target function using the quantum system. Importantly, the basis is *ordered*, with eigentasks at higher k contributing more noise, as dictated by the NSR eigenvalues $\bar{\beta}_k^2$. In particular, at any level of sampling S , there exists an eigentask order $K_c(S)$ after which the NSR eigenvalues $\bar{\beta}_k^2/S$ first drops below unity: $K_c(S) \equiv \max_k \{\bar{\beta}_k^2 < S\}$. Heuristically, including eigentasks $k > K_c(S)$ should contribute more “noise” to the function approximation task than “signal.” In Fig. 5(c), we plot the learned estimates of the likelihood function using $K_L = K_c(S)$ eigentasks for both the CS and PS. First, we note that K_c is lower for the PS than the CS; the former has fewer resolvable eigentasks at a given S . This limitation on resolvable features limits function approximation capacity: the learned estimate of the likelihood function using K_c eigentasks is visibly worse for the PS than the CS.

In this way, higher REC allows noisy quantum systems to better approximate more functions, which translates to improved learning performance—this result is explored systemically in Fig. 6(b). Of course, it is natural to ask whether using $K_c(S) \leq K$ eigentasks is optimal: exactly this question is investigated in Fig. 6(a), where we plot the training and test accuracy of both device encodings as a function of the number of measured eigentasks K_L . The performance on the specific training and test set shown in Fig. 5(b) is indicated with markers, and solid lines indicate the average performance over ten distinct divisions of the data into training and test sets. This permutation of the learning task is a standard technique to optimize hyperparameters in ML and is done here to eliminate the sensitivity of these results to the choice of training set. First note that, in all cases, using all eigentasks ($K_L = K$)—or, equivalently, all measured features $\{\bar{X}\}$ —leads to far lower test accuracy than is found in training. The observed deviation is a distinct signature of overfitting: the optimized estimator learns noise in the training set [comprised of noisy eigentask estimates $\bar{y}^{(k)}(u^{(n)})$] and, thus, loses generalizability to unseen samples in testing.

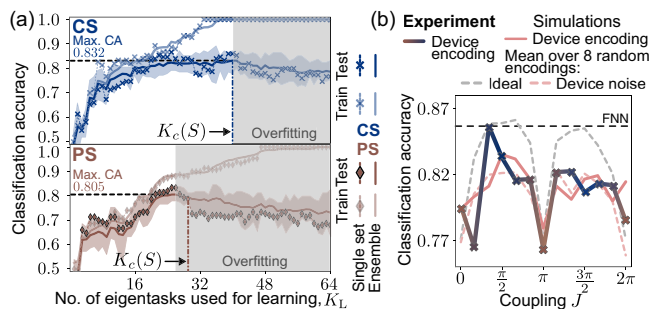


FIG. 6. (a) Training (light and upward trending curve) and testing (dark, peaking, then declining curve) accuracy for the device encodings in Fig. 5(a), as a function of the number of eigentasks used to approximate the target function. Markers indicate performance on the dataset shown in Fig. 5(b), and solid lines are the average over ten random selections of training and test sets. The shaded region denotes the maximum and minimum test accuracy observed. The optimal test set performance is found near the noise-to-signal cutoff $K_c(S = 2^{14})$ (dash-dotted lines) informed by the quantum system’s noise-to-signal spectra. (b) Testing set classification accuracy as a function of J for our optimal learning method. In all cases, the average performance over the ten task permutations is reported, using $K_c(S = 2^{14})$. Cross markers indicate device results for the chosen encoding, and the simulation result for this encoding is shown in solid peach. We further perform device noise simulations for a total of eight random encodings under finite S . The dashed peach line shows the average testing accuracy over these eight encodings. The dashed gray line, in contrast, is the average testing accuracy for ideal simulations and in the $S \rightarrow \infty$ limit, where all $K = 64$ features are used. The horizontal line denotes the performance of a software feed-forward neural network (FNN) with $K_L = 64$ nodes (and $1153 \gg K_c$ trained parameters) for comparison.

Improvements in model training performance with added features are meaningful only insofar as they also lead to better performance on new data: in both encodings, we see test set classification accuracy peaks near $K_c(S)$. This is particularly clear for the averaged results, but even for individual datasets the test accuracy at $K_c(S)$ is within approximately 2% of its maximum, thus confirming our heuristic reasoning that eigentasks beyond this order, with NSR eigenvalues < 1 , hinder learning. The eigentask-learning approach naturally allows one to decompose the outputs from quantum measurements into a compressed basis with known noise properties and then select the set of these which exactly captures the resolvable information at a given S . This robust approach to learning enabled by the capacity analysis maximizes the ability of a noisy quantum system to approximate functions without overfitting to noise, in this case fundamental quantum sampling noise.

Finally, Fig. 6(b) shows the classification accuracy for this device encoding as J is varied, where, following the above approach, the optimal $K_c(S)$ set of eigentasks is used for each encoding. We also show the performance of a similar-scale ($K_L = 64$ node) software neural network and ideal simulations in the $S \rightarrow \infty$ limit [$K_c(\infty) = 64$] for comparison. Note that only these infinite-shot results approach the classical neural network, with quantum sampling noise imposing a significant performance penalty even for $J \approx \pi/2 \pmod{\pi}$. We highlight the striking similarity with Fig. 4(c): encodings with larger quantum correlations and, thus, higher resolvable expressive capacity perform generically better on learning tasks in the presence of noise, because they generate a larger set of eigentasks that can be resolved at a given sampling S . Resolvable expressive capacity is *a priori* unaware of the specific problem considered here; this example therefore emphasizes its power as a general metric predictive of performance on arbitrary tasks.

V. DISCUSSION

We have developed a straightforward approach to quantify the resolvable expressive capacity of any physical system in the presence of fundamental sampling noise. Crucially, this analysis extends to physical quantum systems where sampling noise is fundamentally imposed by quantum measurement theory. Our analysis is built upon an underlying framework that determines the native function set that can be most robustly realized by a finitely sampled physical system: its eigentasks. We use this framework to introduce a methodology for optimal learning that we demonstrate using noisy quantum systems, which centers around identifying the minimal number of eigentasks required for a given learning task. The resulting learning methodology is resource efficient, and the empirical evidence we provide indicates that it is also robust to overfitting. We demonstrate that eigentasks can be efficiently estimated from experiments on real devices using a limited number of training points and finite shots. We also

demonstrate across two distinct qubit-based *Ansätze* that the presence of measured quantum correlations enhances resolvable expressive capacity.

We believe our work opens up several avenues of exploration in the field of learning with physical quantum systems in particular. First, our approach provides the tools to understand the limitations of sampling noise in noisy reservoir computing schemes (e.g., quantum reservoir computing [11,18,22,48,58,59]). In fact, during the final review of the present manuscript, work was posted to the arXiv [60] exploring limits to noisy reservoir computers using an approach closely aligned with our methods here. Second, our work has direct application to the design of circuits for learning with qubit-based systems. In particular, we propose the optimization of resolvable expressive capacity as a meaningful goal for the design of quantum circuits with finite measurement resources. This importantly includes the utilization of the eigentask formulation and eigentask learning as a useful tool for understanding the performance of physical quantum systems in practical learning tasks. Finally, the practical demonstration of our scheme under restrictions of finite input and output samples means that it can prove useful for studies on generalization

and training. For example, any difference in REC and eigentasks computed with optimal weights estimated using only a finite number of input samples—as opposed to the ideal but impractical infinite input sampling limit—would constitute a generalization error over the input domain, which one can seek to minimize for optimal learning in future work.

ACKNOWLEDGMENTS

We express our sincere gratitude to the anonymous reviewers for their invaluable guidance, which significantly contributed to the refinement and enhancement of the final manuscript. We thank Ronen Eldan, Fatih Dinç, Daniel Gauthier, Michael Hatridge, Benjamin Lienhard, Peter McMachon, Sridhar Prabhu, Shyam Shankar, Francesco Tacchino, Logan Wright, and Xun Gao for stimulating discussions about the work that went into this manuscript. This research was developed with funding from the DARPA Contract No. HR00112190072, AFOSR Grant No. FA9550-20-1-0177, and AFOSR MURI Grant No. FA9550-22-1-0203. The views, opinions, and findings expressed are solely the authors' and not the U.S. government's.

APPENDIX A: TABLE OF MAIN NOTATIONS

TABLE I. Table of abbreviations and symbols used in main text and appendixes.

Abbreviations	
REC	Resolvable expressive capacity, C_T
(Q)ML	(Quantum) machine learning
PNN	Physical neural network
POVM	Positive operator-valued measure
ELM	Extreme learning machine
RC	Reservoir computing
SLM	Spatial light modulator
NSR	Noise-to-signal ratio
PS	Product system
CS	Correlated system
ETC	Expected total correlation, \bar{T}
Symbols and notation	
S	Number of shots
N	Number of inputs; for each input, we obtain S output samples or shots
L	Number of qubits
K	Number of measured features; $K = 2^L$ for computational-basis projective measurement
\mathbf{u}	Input
p	Input distribution
$X_k^{(s)}$	Single-shot random-valued features in any physical system
$\mathcal{X}(\mathbf{u})$	Collection of k random-valued features for S shots, $\equiv \{(X_0^{(s)}(\mathbf{u}), X_1^{(s)}(\mathbf{u}), \dots, X_{K-1}^{(s)}(\mathbf{u}))\}_{s \in [S]}$
\mathcal{D}	Complete dataset, $\equiv \{(\mathbf{u}^{(n)}, \mathcal{X}(\mathbf{u}^{(n)}))\}_{n \in [N]}$
$\mathbb{E}_{\mathcal{X}}$	Expectation over the output samples, conditioned on some fixed \mathbf{u}
$\mathbb{E}_{\mathbf{u}}$	Expectation over the input, with underlying prior distribution $p(\mathbf{u})$, $\mathbb{E}_{\mathbf{u}}[f] \equiv \int d\mathbf{u} p(\mathbf{u}) f(\mathbf{u})$

(Table continued)

TABLE I. (*Continued*)

\bar{X}_k	Empirical observed features, $(1/S) \sum_s X_k^{(s)}$
x_k	Expected features, $\mathbb{E}_{\mathcal{X}}[\bar{X}_k]$
ζ_k	Noise component of \bar{X}_k
\mathbf{W}	General output weights
\mathbf{w}	Learned optimal output weights for finite- S features $\{\bar{X}_k\}$
\mathcal{L}	Loss function
\mathbf{G}	Gram matrix of expected features $\{x_k\}$
\mathbf{V}	Expected covariance matrix of random variables $X_k^{(s)}(\mathbf{u})$ over input distribution
\mathbf{D}	Expected second-order-moment matrix of random variable $X_k^{(s)}(\mathbf{u})$ over input distribution; it is diagonal if $X_k^{(s)}(\mathbf{u})$ obeys multinomial distribution
$y^{(k)}$	Eigentasks, $\sum_{k'} r_{k'}^{(k)} x_{k'}$
β_k^2	NSR eigenvalue associated with eigentask $y^{(k)}$
$\mathbf{r}^{(k)}$	Linear combination coefficients of expected features $\{x_{k'}\}$ forming $y^{(k)}$
$\tilde{\beta}_k^2$	Finite- S estimate of β_k^2
$\tilde{\mathbf{r}}^{(k)}$	Finite- S estimate of $\mathbf{r}^{(k)}$
$\bar{y}^{(k)}$	Finite- S estimate of eigentasks, $\sum_{k'} r_{k'}^{(k)} \bar{X}_{k'}$
$\boldsymbol{\theta}$	Quantum system parameters
$\hat{\rho}$	Generated quantum state
\mathcal{U}	Quantum channel
\hat{M}_k	POVM elements, $\equiv \mathbf{b}_k\rangle\langle\mathbf{b}_k $ for computational-basis projective measurement
\mathbf{b}_k	Computational basis eigenstate labels
$k^{(s)}$	Measurement outcome for shot s
$\hat{\rho}^M$	Diagonal post-measurement state, $\sum_k \hat{\rho}_{kk}(\mathbf{u}) \mathbf{b}_k\rangle\langle\mathbf{b}_k $
$K_c(S)$	Cutoff index where β_k^2 approaches S , $\max_k \{\beta_k^2 < S\}$

APPENDIX B: FEATURE MAPS GENERATED BY QUANTUM SYSTEMS

In the main text, we introduce the idea of encoding inputs into the state of a quantum system via a parametrized quantum channel, reproduced below:

$$\hat{\rho}(\mathbf{u}; \boldsymbol{\theta}) = \mathcal{U}(\mathbf{u}; \boldsymbol{\theta}) \hat{\rho}_0; \quad (\text{B1})$$

one then measures this state to approximate desired functions of the input. Figure 7 gives a simple example of this mapping from classical inputs \mathbf{u} (here, in a 2D compact domain) to a quantum state generated by a \mathbf{u} -dependent encoding and finally to the measured features in a two-qubit system undergoing commuting local measurements in the computational basis. The measurement outcomes are therefore bit strings, of which there are $K = 2^L = 4$, namely, $\mathbf{b}_k \in \{00, 01, 10, 11\}$. A given shot yields one of these possible bit strings.

On the right, we plot samples of S -shot features $\{\bar{X}_k\}$ constructed for different numbers of shots $S = 100, 1000, 10000$ (here, we enumerate the feature k with the associated bit string \mathbf{b}_k). The feature-space for a two-qubit system is four-dimensional. Owing to the normalization condition $\sum_k \bar{X}_k = 1$, only three of these dimensions are independent. For ease of visualization, we plot only a two-dimensional projection in the $\bar{X}_{00} - \bar{X}_{11}$ plane. Each dot in this plot is an average [cf. Eq. (1)] over the associated S

shots holding the input $\mathbf{u} = (u_1, u_2)$ identical over those experiments.

As expressed in Eq. (3), the structure of the noise and, thus, the correlations in the distribution is determined by the associated quantum state, subject to an overall scaling with S . It is important to notice here that as $S \rightarrow \infty$ this distribution collapses to a single deterministic point, the corresponding quantum probability $\mathbf{x}(\mathbf{u})$. It is also evident from this plot that the shape and orientation of these clusters depend on the underlying quantum state $\hat{\rho}(\mathbf{u}; \boldsymbol{\theta})$ and associated probabilities $\mathbf{x}(\mathbf{u})$ via Eq. (17). In the remainder of this section, we consider more complex quantum models, such that they generate mappings which can be useful for learning. A descriptive pseudoalgorithm for learning scheme-based circuit *Ansatz* can be found in Algorithm 1.

To describe these models, we begin by first limiting to 1D inputs \mathbf{u} as analyzed in the main text; generalizations to multidimensional inputs \mathbf{u} are straightforward. Then, we write Eq. (B1) in the form

$$\hat{\rho}(\mathbf{u}; \boldsymbol{\theta}) = \hat{U}(\mathbf{u}; \boldsymbol{\theta}) \hat{\rho}_0 \hat{U}^\dagger(\mathbf{u}; \boldsymbol{\theta}). \quad (\text{B2})$$

In the main text, we consider a model for dynamics of an L -qubit quantum system that is natively implementable on modern quantum computing platforms: namely, an *Ansatz* of quantum circuits with single- and two-qubit gates. We refer to this encoding as the *circuit Ansatz* (or *C Ansatz* for

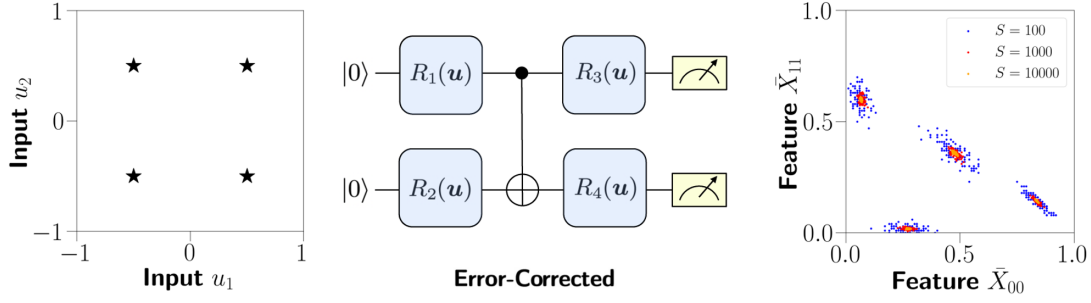


FIG. 7. Schematic of a simple $L = 2$ qubit circuit, comprised of a CNOT gate sandwiched by input-dependent local x -rotation gates $\{R_i(\mathbf{u})\}$. Different 2D inputs shown on the left are mapped to the finite- S feature space on the right via this circuit. Specifically, a 2D slice (\bar{X}_{00} and \bar{X}_{11}) of the 4D feature space is shown. Each point represents an individual sample or experiment, i.e., an output constructed with $S < \infty$ shots via Eq. (1). Distinct values of $S = 10^2, 10^3, 10^4$ are shown in different colors (blue, red, and green, respectively). For each input \mathbf{u} and shots S , the simulation is repeated 100 times, resulting in the distribution shown.

short), for which the operator $\hat{U}(u; \boldsymbol{\theta})$ takes the precise form

$$\hat{U}(u; \boldsymbol{\theta}) = \left[\mathcal{R}_x \left(\frac{\boldsymbol{\theta}^x}{2} \right) \mathcal{W}(J) \mathcal{R}_z(\boldsymbol{\theta}^z + \boldsymbol{\theta}^l u) \mathcal{R}_x \left(\frac{\boldsymbol{\theta}^x}{2} \right) \right]^\tau \quad (C \text{ Ansatz}). \quad (\text{B3})$$

For completeness, we recall that $\mathcal{R}_{x/z}$ are Pauli rotations applied qubitwise, e.g., $\mathcal{R}_z = \otimes_l \hat{R}_z(\theta_l^z + \theta_l^l u)$, while the coupling gate acts between physically connected qubits in the device and can be written as $\mathcal{W}(J) = \prod_{\langle l, l' \rangle} \exp\{-i(J/2)\hat{\sigma}_l^z \hat{\sigma}_{l'}^z\}$. We emphasize here again that $\tau \in \mathbb{N}^+$ is an integer, representing the number of repeated blocks in the *C-Ansatz* encoding. We note that the actual operations implemented on IBMQ processors also include dynamics due to noise, gate, and measurement errors and, thus, must be represented as a general quantum channel as in Eq. (B1). As discussed in the main text, the REC of a quantum system can be computed in the presence of these more general dynamics and is sensitive to the limitations introduced by them.

An alternative *Ansatz* analyzed here is the *Hamiltonian Ansatz* (or *H Ansatz* for short) where the operator $\hat{U}(u; \boldsymbol{\theta})$

describes continuous Hamiltonian dynamics. This *Ansatz* is relevant to computation with general quantum devices, such as quantum annealers and more generally quantum simulators:

$$\hat{U}(u; \boldsymbol{\theta}) = \exp\{-i\hat{H}(u)t\}, \quad \hat{H}(u) = \hat{H}_0 + u \cdot \hat{H}_1 \quad (H \text{ Ansatz}). \quad (\text{B4})$$

Here, t is a continuous parameter defining the evolution time, $\hat{H}_0 = \sum_{\langle l, l' \rangle} J_{l, l'} \hat{\sigma}_l^z \hat{\sigma}_{l'}^z + \sum_{l=1}^L h_l^x \hat{\sigma}_l^x + \sum_{l=1}^L h_l^z \hat{\sigma}_l^z$, and $\hat{H}_1 = \sum_{l=1}^L h_l^l \hat{\sigma}_l^z$. The transverse x -field strength $h_l^x = \bar{h}^x + \varepsilon_l^x$ and longitudinal z -drive strength $h_l^{z, l} = \bar{h}^{z, l} + \varepsilon_l^{z, l}$ are all randomly chosen and held fixed for a given realization of the quantum system:

$$\varepsilon_l^{x, z, l} \sim h_{\text{rms}}^{x, z, l} \mathcal{N}(0, 1), \quad (\text{B5})$$

where $\mathcal{N}(0, 1)$ defines the standard normal distribution with zero mean and unit variance. We consider nearest-neighbor interactions $J_{l, l'}$, which can be constant $J_{l, l'} \equiv J$, or drawn from $J_{l, l'} \sim \text{Unif}[0, J_{\text{max}}]$, where $\text{Unif}[a, b]$ is a uniform distribution with nonzero density within $[a, b]$.

Algorithm 1. Measured features under multinomial sampling in quantum system.

Input: $u \in [-1, +1]$

Output: $\bar{\mathbf{X}}(u)$, which approximates $x_k(u) := \text{Tr}\{\hat{\rho}(u)|\mathbf{b}_k\rangle\langle\mathbf{b}_k|\}$

For $s \leftarrow 1$ to S

Initialize overall state $\hat{\rho}_0 \leftarrow |0\rangle\langle 0|^{\otimes L}$;

Evolve under quantum channel $\mathcal{U}(u)$: $\hat{\rho}(u) \leftarrow \mathcal{U}(u)\hat{\rho}_0$;

Measure all L qubits: $\mathbf{b}^{(s)}(u) \leftarrow \mathbf{b}_k = (b_{k,1}, b_{k,2}, \dots, b_{k,L}) \in \{0, 1\}^L$;

EndFor

For $k \leftarrow 0$ to $K - 1$

Take the ensemble averages as readout features:

$$\bar{X}_k(u) \leftarrow \frac{1}{S} \sum_{s=1}^S \delta(\mathbf{b}_k, \mathbf{b}^{(s)}(u));$$

EndFor

/ Notice $x_k(u) := \text{Tr}\{\hat{\rho}(u)|\mathbf{b}_k\rangle\langle\mathbf{b}_k|\} = \lim_{S \rightarrow \infty} \bar{X}_k(u)$ */*

Algorithm 2. Training of output weights.

Input: $\{u^{(1)}, \dots, u^{(N)}\} \in [-1, +1]^N$

Output: $\tilde{\mathbf{w}}_N$, such that $y = \tilde{\mathbf{w}}_N \cdot \tilde{\mathbf{X}}(u)$ can approximate $f(u)$

For $n \leftarrow 1$ to N

 Generate features $\tilde{\mathbf{X}}(u^{(n)})$ through Algorithm 1

EndFor

Collect the features into a regression matrix $\tilde{\mathbf{F}}_N \in \mathbb{R}^{N \times K}$;

Compute empirical Gram matrix $\tilde{\mathbf{G}} \leftarrow \frac{1}{N} \tilde{\mathbf{F}}_N^T \tilde{\mathbf{F}}_N$;

/* For finite S , $\lim_{N \rightarrow \infty} \tilde{\mathbf{G}} = \tilde{\mathbf{G}} := \mathbf{G} + \frac{1}{S} \mathbf{V}$ */

Compute target vector $\mathbf{Y} \leftarrow (f(u^{(1)}), \dots, f(u^{(N)}))^T$;

$\mathbf{w}^* \leftarrow (\tilde{\mathbf{F}}_N^T \tilde{\mathbf{F}}_N)^{-1} \tilde{\mathbf{F}}_N^T \mathbf{Y}$;

/* For finite S , $\lim_{N \rightarrow \infty} \mathbf{w}^* = \mathbf{w} := \text{Eq. (C12)}$ */

As an aside, we note that the *C-Ansatz* quantum channel described by Eq. (B3) can be considered a Trotterization-inspired implementation of the *H Ansatz* in Eq. (B4). In particular, if we set $\theta^{x/z/I} = h^{x/z/I} \Delta \cdot \tau$, where $t = \Delta \cdot \tau$, and consider the limit $\Delta \rightarrow 0$ while keeping t fixed, Eq. (B3) corresponds to a Trotterized implementation of Eq. (B4). This correspondence is chosen for practical reasons but is not necessary in our analysis.

APPENDIX C: INFORMATION CAPACITY WITH SAMPLING NOISE

In this appendix, we provide a detailed *theoretical* construction of the REC analysis. Important related *numerical* techniques required for the *practical* calculation of REC are addressed in Appendix D.

We start building the theory by rewriting the function capacity Eq. (5) into a generalized Rayleigh quotient Eq. (C13) in Appendix C 1. By solving for the critical point of this generalized Rayleigh quotient, we prove in Appendix C 2 that it naturally defines a set of orthonormal functions—the *eigentasks* $\{y^{(k)}\}$. Each eigentask of the Rayleigh quotient is associated with a particular eigenvalue $\{\beta_k^2\}$, called eigen-noise-to-signal ratio (NSR), whose meaning is interpreted in Appendix C 3. In Appendix C 4, we derive a formula $C_T = \sum_k 1/(1 + \beta_k^2/S)$ to compute the resolvable expressive capacity, which reproduces the main result Eq. (7) in the main text. In many practical scenarios in machine learning, like classification problems, there is usually one more nonlinear postprocessing function σ_{NL} acting on $\mathbf{W} \cdot \tilde{\mathbf{X}}(u)$. We generalize our methodology to those cases in Appendix C 5, via an approximation truncating the third- and higher-order derivatives of σ_{NL} . Finally, we give a simplified, equivalent form of eigenproblem in Appendix C 6 specifically for systems obeying multinomial statistics, which allows solving for both the eigentasks and NSR eigenvalue more simply.

1. Definition of capacity for physical systems with sampling noise

Suppose an arbitrary probability distribution $p(u)$ for a random (scalar) variable u defined in $D \subseteq \mathbb{R}$. This naturally

defines a function space $L_p^2(D)$ containing all functions $f: D \rightarrow \mathbb{R}$ with $\int f^2(u)p(u)du < \infty$. The space is equipped with the inner product structure $\langle f_1, f_2 \rangle_p = \int f_1(u) \times f_2(u)p(u)du$. We first review the definition of the function approximation capacity known as the *information processing capacity* (IPC) introduced in Ref. [18]. This deterministic quantity is based on a metric quantifying the accuracy of a physical system to approximate one of the functions $f_\ell(u)$ of its input through a linear estimator based on its accessible (measurable) degrees of freedom $x_k(u)$:

$$C[f_\ell] = 1 - \min_{\mathbf{w}_\ell \in \mathbb{R}^K} \frac{\int (\sum_{k=0}^{K-1} W_{\ell k} x_k(u) - f_\ell(u))^2 p(u) du}{\int f_\ell(u)^2 p(u) du}, \quad (\text{C1})$$

where functions $f_\ell(u)$ are orthogonal target functions $\langle f_\ell, f_{\ell'} \rangle_p = \int f_\ell(u) f_{\ell'}(u) p(u) du = 0$ for $\ell \neq \ell'$. The IPC is defined as $C_T \equiv \sum_{\ell=0}^{\infty} C[f_\ell]$, capturing the ability of what type of function the linear combination of physical system readout features can produce. Reference [18] provides an upper bound for the IPC; the IPC of any generic dynamical system is bounded by the accessible degrees of freedom, $C_T \leq K$.

While this result is quite general, it neglects the limitations due to noise in readout features, which is unavoidable when using physical systems in the presence of finite computational and measurement resources. It is generally accepted that the capacity is reduced in the presence of additive noise, but there are no general results on how to quantify that reduction in the presence of *general physical noise*. This is our goal here, to arrive at an exact result for capacity reduction under well-defined but sufficiently general conditions given the physical system. Additional desideratum on this metric is that it provides a practical, calculable metric that can be calculated either (i) from numerical solution of the dynamics of the physical system incorporating a sufficiently accurate noise model or (ii) if desired from experimental data extracted from said physical system. In the main text, we provide a comparative analysis

of these two modalities for a seven-qubit superconducting quantum processor.

We start from considering the noisy readout features $\bar{X}(u)$ whose expectations and covariances are

$$\mathbb{E}_{\mathcal{X}}[\bar{X}(u)] \equiv \mathbf{x}(u), \quad (\text{C2})$$

$$\text{Cov}_{\mathcal{X}}[\bar{X}(u)] \equiv \frac{1}{S}\boldsymbol{\Sigma}(u), \quad (\text{C3})$$

where the expectation and covariance are evaluated under the product distribution of S -shot IID variable $\{X_k^{(s)}(u)\}$.

To determine the optimal capacity to compute an arbitrary normalized function $f(u)$ using the noisy readout features $\bar{X}(u)$ extracted from the physical system, we need to find an optimal \mathbf{W} such that

$$C[f] = 1 - \frac{\min_{\mathbf{W}} \int \mathbb{E}_{\mathcal{X}}[(\sum_{k=0}^{K-1} W_k \bar{X}_k(u) - f(u))^2] p(u) du}{\int f^2(u) p(u) du}. \quad (\text{C4})$$

By expanding the numerator of the right-hand side for a given, finite number of shots S , we find

$$\begin{aligned} & \int f^2(u) p(u) du - \int \mathbb{E}_{\mathcal{X}} \left[\left(\sum_{k=0}^{K-1} W_k \bar{X}_k(u) - f(u) \right)^2 \right] p(u) du \\ &= - \sum_{k_1=0}^{K-1} \sum_{k_2=0}^{K-1} W_{k_1} W_{k_2} \int \mathbb{E}_{\mathcal{X}}[\bar{X}_{k_1}(u) \bar{X}_{k_2}(u)] p(u) du + 2 \sum_{k=0}^{K-1} W_k \int \mathbb{E}_{\mathcal{X}}[\bar{X}_k(u)] f(u) p(u) du \\ &= - \sum_{k_1=0}^{K-1} \sum_{k_2=0}^{K-1} W_{k_1} W_{k_2} \int \left(x_{k_1}(u) x_{k_2}(u) + \frac{1}{S} \boldsymbol{\Sigma}(u)_{k_1 k_2} \right) p(u) du + 2 \sum_{k=0}^{K-1} W_k \int x_k(u) f(u) p(u) du, \end{aligned} \quad (\text{C5})$$

where the final line comes from the property of the covariance matrix $\mathbb{E}_{\mathcal{X}}[\bar{X}_j(u) \bar{X}_k(u)] - \mathbb{E}_{\mathcal{X}}[\bar{X}_j(u)] \mathbb{E}_{\mathcal{X}}[\bar{X}_k(u)] = (1/S) \boldsymbol{\Sigma}(u)_{jk}$.

The goal of the remaining part of this section is deducing a more compact generalized Rayleigh quotient form of function capacity. The dependence of readout features $x_k(u)$ on the input u can always be written in the form of a Taylor expansion:

$$x_k(u) = \sum_{j=0}^{\infty} (\mathbf{T})_{kj} u^j, \quad (\text{C6})$$

where we define the *transfer matrix* $\mathbf{T}(\boldsymbol{\theta}) \equiv \mathbf{T} \in \mathbb{R}^{K \times \infty}$ that depends on the density matrix $\hat{\rho}(u)$ and, in particular, on parameters $\boldsymbol{\theta}$ characterizing the physical system. The first term in Eq. (C5) does not depend explicitly on the function $f(u)$ being constructed and introduces quantities that are determined entirely by the response of the physical system of interest to inputs over the entire domain of u . In particular, we introduce the *Gram matrix* $\mathbf{G} \in \mathbb{R}^{K \times K}$ as

$$(\mathbf{G})_{k_1 k_2} = \int x_{k_1}(u) x_{k_2}(u) p(u) du = \sum_{j_1=0}^{\infty} \sum_{j_2=0}^{\infty} (\mathbf{T})_{k_1 j_1} \left(\int u^{j_1+j_2} p(u) du \right) (\mathbf{T})_{k_2 j_2} \equiv (\mathbf{T} \boldsymbol{\Lambda} \mathbf{T}^T)_{k_1 k_2}, \quad (\text{C7})$$

where we also introduce the *generalized Hilbert matrix* $\boldsymbol{\Lambda} \in \mathbb{R}^{\infty \times \infty}$ as

$$(\boldsymbol{\Lambda})_{j_1 j_2} = \int u^{j_1+j_2} p(u) du. \quad (\text{C8})$$

Second, we introduce the noise matrix $\mathbf{V} \in \mathbb{R}^{K \times K}$:

$$(\mathbf{V})_{k_1 k_2} = \int \boldsymbol{\Sigma}(u)_{k_1 k_2} p(u) du. \quad (\text{C9})$$

The second term in Eq. (C5) depends on $f(u)$ and can be simplified using the $\boldsymbol{\Lambda}$ matrix as well. Introducing the Taylor series expansion $f(u) = \sum_{j=0}^{\infty} (\mathbf{Y})_j u^j$,

$$\begin{aligned} \mathbb{E}_u[x_k f] &= \int x_k(u) f(u) p(u) du \\ &= \sum_{j_1=0}^{\infty} \sum_{j_2=0}^{\infty} (\mathbf{T})_{k j_1} \left(\int u^{j_1+j_2} p(u) du \right) (\mathbf{Y})_{j_2} \\ &= (\mathbf{T} \boldsymbol{\Lambda} \mathbf{Y})_k. \end{aligned} \quad (\text{C10})$$

With these definitions, Eq. (C4) can be compactly written in matrix form as a Tikhonov regularization problem:

$$C[f] = 1 - \min_{\mathbf{W}} \left(\frac{\|\boldsymbol{\Lambda}^{1/2} \mathbf{T}^T \mathbf{W} - \boldsymbol{\Lambda}^{1/2} \mathbf{Y}\|^2 + \frac{1}{S} \mathbf{W}^T \mathbf{V} \mathbf{W}}{\mathbf{Y}^T \boldsymbol{\Lambda} \mathbf{Y}} \right). \quad (\text{C11})$$

The least-squares form ensures that the optimal value (argmin) \mathbf{w} of \mathbf{W} has closed form:

$$\mathbf{w} = \left(\mathbf{G} + \frac{1}{S} \mathbf{V} \right)^{-1} \mathbb{E}_u[\mathbf{x}f] = \left(\mathbf{T}\mathbf{\Lambda}\mathbf{T}^T + \frac{1}{S} \mathbf{V} \right)^{-1} \mathbf{T}\mathbf{\Lambda}\mathbf{Y}. \quad (\text{C12})$$

Substituting \mathbf{w} into the expression for C , we obtain the optimal capacity with which a function f can be constructed, which takes the form of a *generalized Rayleigh quotient*:

$$\begin{aligned} C[f] &= \frac{\mathbb{E}_u[\mathbf{x}^T f] (\mathbf{G} + \frac{1}{S} \mathbf{V})^{-1} \mathbb{E}_u[\mathbf{x}f]}{\mathbb{E}_u[f^2]} \\ &= \frac{\mathbf{Y}^T \mathbf{\Lambda} \mathbf{T}^T (\mathbf{G} + \frac{1}{S} \mathbf{V})^{-1} \mathbf{T} \mathbf{\Lambda} \mathbf{Y}}{\mathbf{Y}^T \mathbf{\Lambda} \mathbf{Y}}. \end{aligned} \quad (\text{C13})$$

In the scenario of fundamental quantum measurement noise, or quantum sampling noise, given u and S , the quantum readout features $\bar{X}_k(u) = (1/S) \sum_{s=1}^S \delta(k^{(s)}(u), k)$ are stochastic variables. The expectation vector and

covariance matrix of $\bar{\mathbf{X}}(u)$ can be expressed in terms of $\hat{\rho}(u)$:

$$\mathbb{E}_{\mathcal{X}}[\bar{\mathbf{X}}(u)] \equiv \mathbf{x}(u) = \text{Tr}\{\hat{M}_k \hat{\rho}(u)\}, \quad (\text{C14})$$

$$\text{Cov}_{\mathcal{X}}[\bar{\mathbf{X}}(u)] \equiv \frac{1}{S} \mathbf{\Sigma}(u) = \frac{1}{S} (\text{diag}(\mathbf{x}) - \mathbf{x}\mathbf{x}^T). \quad (\text{C15})$$

To understand Eq. (C15), we provide brief proofs of some important identities involving the second-order statistics of multinomial distribution. In the case of quantum measurement noise, the single-shot random-valued feature is the indicator $X_k^{(s)}(u) = \delta(k^{(s)}(u), k)$ for shot s and input u . By definition, for any s , the expectation of indicator is always the probability of obtaining index k , given the input u : $\mathbb{E}_{\mathcal{X}}[X_k^{(s)}(u)] = x_k(u)$. If $s \neq s'$, the experiment for different shots must be independent, so we have $\mathbb{E}_{\mathcal{X}}[X_k^{(s)}(u)X_{k'}^{(s')}(u)] = x_k(u)x_{k'}(u)$; while if $s = s'$, the mutual exclusion for getting different indices in one shot implies that $\mathbb{E}_{\mathcal{X}}[X_k^{(s)}(u)X_{k'}^{(s)}(u)] = \delta_{kk'}x_k(u)$. Thus, we can unify them into one equation:

$$\mathbb{E}_{\mathcal{X}}[X_k^{(s)}(u)X_{k'}^{(s')}(u)] = \delta_{ss'}\delta_{kk'}x_k(u) + (1 - \delta_{ss'})x_k(u)x_{k'}(u) = x_k(u)x_{k'}(u) + \delta_{ss'}(\delta_{kk'}x_k(u) - x_k(u)x_{k'}(u)). \quad (\text{C16})$$

Then, the expectation of the product of \bar{X}_k and $\bar{X}_{k'}$ is

$$\begin{aligned} \mathbb{E}_{\mathcal{X}}[\bar{X}_k(u)\bar{X}_{k'}(u)] &= \frac{1}{S^2} \sum_{s,s'} \{x_k(u)x_{k'}(u) + \delta_{ss'}[\delta_{kk'}x_k(u) - x_k(u)x_{k'}(u)]\} \\ &= x_k(u)x_{k'}(u) + \frac{1}{S} [\delta_{kk'}x_k(u) - x_k(u)x_{k'}(u)], \end{aligned} \quad (\text{C17})$$

which can be directly used to derive Eq. (C15); namely, the element $\mathbf{\Sigma}_{kk'}(\mathbf{u})$ is

$$\begin{aligned} \mathbb{E}_{\mathcal{X}}[(\bar{X}_k(u) - x_k(u))(\bar{X}_{k'}(u) - x_{k'}(u))] &= \mathbb{E}_{\mathcal{X}}[\bar{X}_k(u)\bar{X}_{k'}(u)] - \mathbb{E}_{\mathcal{X}}[\bar{X}_k(u)x_{k'}(u)] - \mathbb{E}_{\mathcal{X}}[x_k(u)\bar{X}_{k'}(u)] + \mathbb{E}_{\mathcal{X}}[x_k(u)x_{k'}(u)] \\ &= x_k(u)x_{k'}(u) + \frac{1}{S} [\delta_{kk'}x_k(u) - x_k(u)x_{k'}(u)] - x_k(u)x_{k'}(u) \\ &= \frac{1}{S} [\delta_{kk'}x_k(u) - x_k(u)x_{k'}(u)]. \end{aligned} \quad (\text{C18})$$

Therefore, for quantum sampling noise, we can express the noise matrix \mathbf{V} more explicitly:

$$(\mathbf{V})_{k_1 k_2} = \int \mathbf{\Sigma}(u)_{k_1 k_2} p(u) du = \int (\delta_{k_1 k_2} x_{k_1}(u) - x_{k_1}(u)x_{k_2}(u)) p(u) du \equiv (\mathbf{D})_{k_1 k_2} - (\mathbf{G})_{k_1 k_2} \quad (\text{C19})$$

Here, we also introduce the *second-order-moment* matrix $\mathbf{D} \in \mathbb{R}^{K \times K}$ such that

$$(\mathbf{D})_{k_1 k_2} = \delta_{k_1 k_2} \sum_k \mathbf{G}_{kk_1} = \delta_{k_1 k_2} \int x_{k_1}(u) p(u) du. \quad (\text{C20})$$

Then, the noise matrix simply defines the covariance of readout features and is, therefore, given by $\mathbf{V} = \mathbf{D} - \mathbf{G}$.

In the study of quantum machine learning, it is convenient to define the t th-order *quantum ensemble moment* of ensemble $\mathcal{E} = \{p(\mathbf{u})d\mathbf{u}, \hat{\rho}(\mathbf{u})\}$ in a t -copy space [46]:

$$\hat{\rho}^{(t)} = \int \hat{\rho}(\mathbf{u})^{\otimes t} p(\mathbf{u}) d\mathbf{u}. \quad (\text{C21})$$

The Gram matrix and second-order-moment matrix can then be compactly expressed as

$$\mathbf{D}_{kk} = \int x_k(\mathbf{u}) p(\mathbf{u}) d\mathbf{u} = \int \text{Tr}\{\hat{M}_k \hat{\rho}(\mathbf{u})\} p(\mathbf{u}) d\mathbf{u} = \text{Tr}\left\{ \int \hat{M}_k \hat{\rho}(\mathbf{u}) p(\mathbf{u}) d\mathbf{u} \right\} = \text{Tr}\{\hat{M}_k \hat{\rho}^{(1)}\}, \quad (\text{C22})$$

$$\begin{aligned} \mathbf{G}_{kk'} &= \int x_k(\mathbf{u}) x_{k'}(\mathbf{u}) p(\mathbf{u}) d\mathbf{u} = \int \text{Tr}\{\hat{M}_k \hat{\rho}(\mathbf{u})\} \text{Tr}\{\hat{M}_{k'} \hat{\rho}(\mathbf{u})\} p(\mathbf{u}) d\mathbf{u} \\ &= \text{Tr}\left\{ (\hat{M}_k \otimes \hat{M}_{k'}) \left(\int \hat{\rho}(\mathbf{u}) \otimes \hat{\rho}(\mathbf{u}) p(\mathbf{u}) d\mathbf{u} \right) \right\} = \text{Tr}\{(\hat{M}_k \otimes \hat{M}_{k'}) \hat{\rho}^{(2)}\}. \end{aligned} \quad (\text{C23})$$

While most of the results in our paper do not utilize this representation, we do find that it provides a compact and natural representation, and we use this representation to derive the analytical results in Sec. III B; see Appendix F.

2. Eigentasks

Equation (C13) defines the optimal capacity of approximating an arbitrary function $f(\mathbf{u}) = \sum_{j=0}^{\infty} (\mathbf{Y})_j u^j$. We can, therefore, naturally ask which functions f maximize this optimal capacity. To this end, we first note that the denominator of Eq. (C13) is simply a normalization factor that can be absorbed into the definition of the function $f(\mathbf{u})$ being approximated, without loss of generality. More precisely, we consider

$$\langle f, f \rangle_p = 1 = (\Lambda^{1/2} \mathbf{Y})^T (\Lambda^{1/2} \mathbf{Y}) = \mathbf{Y}^T \Lambda \mathbf{Y}. \quad (\text{C24})$$

Then, we can rewrite the optimal capacity from Eq. (C13) as

$$C[f] = \mathbf{Y}^T \Lambda^{1/2} \mathbf{Q} \Lambda^{1/2} \mathbf{Y}. \quad (\text{C25})$$

Here, we define the matrix $\mathbf{Q} \in \mathbb{R}^{\infty \times \infty}$ as

$$\mathbf{Q} = \mathbf{B} \left(\mathbf{I} + \frac{1}{S} \mathbf{R} \right)^{-1} \mathbf{B}^T, \quad (\text{C26})$$

$$\mathbf{B} = \Lambda^{1/2} \mathbf{T}^T \mathbf{G}^{-1/2}, \quad (\text{C27})$$

$$\mathbf{R} = \mathbf{G}^{-1/2} \mathbf{V} \mathbf{G}^{-1/2} \quad (\text{C28})$$

by introducing the matrix square root of $\mathbf{G}^{1/2} \in \mathbb{R}^{K \times K}$ and the NSR matrix \mathbf{R} . The decomposition in Eq. (C26) may be verified by direct substitution into Eq. (C25). The ability to calculate matrix powers and, in particular, the inverse of \mathbf{G} requires constraints on its rank.

Before we analytically find the eigenvectors of \mathbf{Q} , we need to show, in general, that the number of linearly independent features always equals to the rank of the Gram matrix \mathbf{G} , no matter what symmetries the system is subject to. Let us consider any vector $\mathbf{c} \in \mathbb{R}^K$; the quadratic form

$$\sum_{k_1, k_2=0}^{K-1} c_{k_1} c_{k_2} (\mathbf{G})_{k_1, k_2} = \int \left(\sum_{k_1=1}^K c_{k_1} x_{k_1}(\mathbf{u}) \right) \left(\sum_{k_2=1}^K c_{k_2} x_{k_2}(\mathbf{u}) \right) p(\mathbf{u}) d\mathbf{u} = \left\langle \sum_{k=0}^{K-1} c_k x_k, \sum_{k=0}^{K-1} c_k x_k \right\rangle_p \quad (\text{C29})$$

gives the norm of function $\sum_{k=0}^{K-1} c_k x_k(\mathbf{u})$ in the rhs. The summation $\sum_{k_1, k_2=1}^K c_{k_1} c_{k_2} (\mathbf{G})_{k_1, k_2} = 0$ vanishes if and only if function $\sum_{k=0}^{K-1} c_k x_k(\mathbf{u})$ is a zero function; namely, \mathbf{c} is in the null space of \mathbf{G} . We conclude that the rank Gram matrix \mathbf{G} is equal to the number of linearly independent features. One way for the rank of \mathbf{G} to be deficient is through the use of the input encoding that leads to identical dependence of the features in a way that is not broken by

the rest of the system's interactions (dictated by $\boldsymbol{\theta}$). We do not consider such symmetries. All simulations utilize encodings that result in a full rank \mathbf{G} . In the case of a rank-deficient \mathbf{G} , one should replace all appearances of \mathbf{G}^{-1} above with the pseudoinverse \mathbf{G}^+ .

We now consider the measure-independent part of the eigenvectors of \mathbf{Q} , indexed $\mathbf{Y}^{(k)}$, satisfying the standard eigenvalue problem:

$$\mathbf{Q}\boldsymbol{\Lambda}^{1/2}\mathbf{Y}^{(k)} = C_k\boldsymbol{\Lambda}^{1/2}\mathbf{Y}^{(k)}, \quad (\text{C30})$$

where $k = 0, \dots, K-1$. From Eq. (C25), it is clear that these eigenvectors have a particular meaning. Consider the function $y^{(k)}(u)$ defined by the eigenvector $\mathbf{Y}^{(k)}$, namely,

$$y^{(k)}(u) = \sum_{j=0}^{\infty} \mathbf{Y}_j^{(k)} u^j, \quad (\text{C31})$$

which we refer to from now on as eigentasks. Suppose we wish to construct the function $y^{(k)}(u)$ using outputs obtained from the physical system defined by \mathbf{Q} in the $S \rightarrow \infty$ limit (namely, with *deterministic* outputs). At first glance, before we dive into solving the eigenproblem Eq. (C30), we do not know any relationship between $y^{(k)}$ and $\mathbf{x}(u)$. The rest of this subsection is aiming to prove that $y^{(k)}$ must be a specific linear combination of features $\mathbf{x}(u)$. Then, the physical system's capacity for this construction is simply given by the corresponding eigenvalue C_k , as may be seen by substituting Eq. (C30) into Eq. (C25). Formally, the $y^{(k)}(u)$ serves as the *critical point* (or *stationary point*) of the generalized Rayleigh quotient in Eq. (C13). Consequently, the function that is constructed with largest capacity then corresponds to the nontrivial eigenvector with largest eigenvalue.

To obtain these eigentasks, we must solve the eigenproblem defined by Eq. (C30). Here, the representation of \mathbf{Q} in Eq. (C26) becomes useful, as we see that the eigensystem of \mathbf{Q} is related closely to that of the NSR matrix \mathbf{R} . In particular, we first define the eigenproblem of \mathbf{R} :

$$\mathbf{R}\mathbf{G}^{1/2}\mathbf{r}^{(k)} = \beta_k^2\mathbf{G}^{1/2}\mathbf{r}^{(k)} \quad (\text{C32})$$

with NSR eigenvalues β_k^2 and corresponding eigenvectors $\mathbf{r}^{(k)}$, which satisfy the orthogonality relation $\mathbf{r}^{(k)T}\mathbf{G}\mathbf{r}^{(k)} = \delta_{k,k'}$. Here, the $\mathbf{r}^{(k)}$ can also be computed from the solution to a simpler generalized eigenproblem, where matrix square root operation $\mathbf{G}^{1/2}$ is not needed:

$$\mathbf{V}\mathbf{r}^{(k)} = \beta_k^2\mathbf{G}\mathbf{r}^{(k)}. \quad (\text{C33})$$

This is because $\mathbf{V}\mathbf{r}^{(k)} = \mathbf{G}^{1/2}\mathbf{R}\mathbf{G}^{1/2}\mathbf{r}^{(k)} = \beta_k^2\mathbf{G}^{1/2}\mathbf{G}^{1/2}\mathbf{r}^{(k)} = \beta_k^2\mathbf{G}\mathbf{r}^{(k)}$. The prefactor $\mathbf{G}^{1/2}$ is introduced for later convenience. Equation (C32) then allows us to define the related eigenproblem

$$\left(\mathbf{I} + \frac{1}{S}\mathbf{R}\right)^{-1}\mathbf{G}^{1/2}\mathbf{r}^{(k)} = \left(1 + \frac{\beta_k^2}{S}\right)^{-1}\mathbf{G}^{1/2}\mathbf{r}^{(k)}. \quad (\text{C34})$$

Next, we note that \mathbf{Q} is related to the matrix in brackets above via a *generalized* similarity transformation defined by \mathbf{B} , Eq. (C26). In particular, $\mathbf{B}^T\mathbf{B} = \mathbf{G}^{-1/2}\mathbf{G}\mathbf{G}^{-1/2} = \mathbf{I} \in \mathbb{R}^{K \times K}$,

while we remark that $\mathbf{B}\mathbf{B}^T \neq \mathbf{I}$ since it is in $\mathbb{R}^{\infty \times \infty}$. This connection allows us to show that

$$\begin{aligned} \mathbf{Q}\mathbf{B}\mathbf{G}^{1/2}\mathbf{r}^{(k)} &= \mathbf{B}\left(\mathbf{I} + \frac{1}{S}\mathbf{R}\right)^{-1}\mathbf{B}^T\mathbf{B}\mathbf{G}^{1/2}\mathbf{r}^{(k)} \\ &= \frac{1}{1 + \beta_k^2/S}\mathbf{B}\mathbf{G}^{1/2}\mathbf{r}^{(k)}. \end{aligned} \quad (\text{C35})$$

Comparing with Eq. (C30), we can now simply read off both the eigenvalues and eigenvectors of \mathbf{Q} :

$$\left. \begin{aligned} C_k &= \frac{1}{1 + \beta_k^2/S} \\ \boldsymbol{\Lambda}^{1/2}\mathbf{Y}^{(k)} &= \mathbf{B}\mathbf{G}^{1/2}\mathbf{r}^{(k)} \end{aligned} \right\} \Rightarrow \mathbf{Y}^{(k)} = \mathbf{T}\mathbf{r}^{(k)}, \quad (\text{C36})$$

where we use the definition of \mathbf{B} from Eq. (C27). The functions defined by the eigenvectors $\mathbf{Y}^{(k)}$ are automatically orthonormal:

$$\begin{aligned} \langle y^{(k_1)}, y^{(k_2)} \rangle_p &= (\boldsymbol{\Lambda}^{1/2}\mathbf{Y}^{(k_1)})^T (\boldsymbol{\Lambda}^{1/2}\mathbf{Y}^{(k_2)}) \\ &= \mathbf{r}^{(k_1)T}\mathbf{G}^{1/2}\mathbf{B}^T\mathbf{B}\mathbf{G}^{1/2}\mathbf{r}^{(k_2)} \\ &= \mathbf{r}^{(k_1)T}\mathbf{G}\mathbf{r}^{(k_2)} = \delta_{k_1 k_2}. \end{aligned} \quad (\text{C37})$$

3. Noisy eigentasks from readout features

We can now also discuss the interpretation of $\{\beta_k^2\}$ for a physical system for which $\{\mathbf{r}^{(k)}\}$ are known. Consider the evaluation by the physical system (for a given u) under finite shots S , which yields a single instance of the readout features $\bar{\mathbf{X}}(u)$. We can simply construct a *noisy* estimator of the k th eigentask $\bar{y}^{(k)}(u)$:

$$\bar{y}^{(k)}(u) = \sum_{k'=0}^{K-1} r_{k'}^{(k)} \bar{X}_{k'}(u), \quad (\text{C38})$$

which is equivalent to requiring the output weights $\mathbf{W} = \mathbf{r}^{(k)}$. The corresponding set of noisy function is also orthogonal; this is because $\mathbf{V}\mathbf{r}^{(k)} = \beta_k^2\mathbf{G}\mathbf{r}^{(k)}$ implies $\mathbf{r}^{(k)T}\mathbf{V}\mathbf{r}^{(k')} = \beta_k^2\delta_{k,k'}$ and, hence,

$$\begin{aligned} \mathbb{E}_u[\mathbb{E}_{\mathcal{X}}[\bar{y}^{(k_1)}\bar{y}^{(k_2)}]] &= \mathbf{r}^{(k_1)T}\left(\mathbf{G} + \frac{1}{S}\mathbf{V}\right)\mathbf{r}^{(k_2)} \\ &= \left(1 + \frac{\beta_k^2}{S}\right)\delta_{k_1 k_2}. \end{aligned} \quad (\text{C39})$$

Let us define $\bar{y}^{(k)}(u) = y^{(k)}(u) + \xi^{(k)}(u)$. It means, for each k , the noisy eigentask $\bar{y}^{(k)}(u)$ contains a signal part $y^{(k)}(u)$ and a noise part in $\xi^{(k)}(u)$, where the latter one is computed from the linear combination $\xi^{(k)}(u) = (1/\sqrt{S}) \times \sum_{k'=0}^{K-1} r_{k'}^{(k)} \zeta_{k'}(u)$. One can check $\mathbb{E}_u[\mathbb{E}_{\mathcal{X}}[y^{(k_1)}\xi^{(k_2)}]] = \mathbb{E}_u[\mathbb{E}_{\mathcal{X}}[\xi^{(k_1)}y^{(k_2)}]] = 0$, and

$$\mathbb{E}_u [y^{(k_1)} y^{(k_2)}] = \mathbf{r}^{(k_1)T} \mathbf{G} \mathbf{r}^{(k_2)} = \delta_{k_1 k_2}, \quad (\text{C40})$$

$$\mathbb{E}_u [\mathbb{E}_{\mathcal{X}} [\xi^{(k_1)} \xi^{(k_2)}]] = \frac{1}{S} \mathbf{r}^{(k_1)T} \mathbf{V} \mathbf{r}^{(k_2)} = \frac{\beta_{k_1}^2}{S} \delta_{k_1 k_2}. \quad (\text{C41})$$

It means that taking linear combinations of $\{x_k(u)\}$ and $\{\zeta_k(u)/\sqrt{S}\}$ with coefficients $\{\mathbf{r}^{(k)} \in \mathbb{R}^K\}_{k \in [K]}$ not only produces orthonormal eigentasks $\{y^{(k)}(u)\}$ for signal, but also induces a set of orthogonal noise functions $\{\xi^{(k)}(u)\}$.

If the physical system can be run multiple times for a given S , multiple instances of $\bar{\mathbf{X}}(u)$ can be obtained, from each of which an estimate of the k th eigentask $\bar{y}^{(k)}(u)$ can be constructed. The expectation value of these estimates then simply yields

$$\mathbb{E}_{\mathcal{X}} [\bar{y}^{(k)}(u)] = \sum_{k'=0}^{K-1} r_{k'}^{(k)} \mathbb{E}_{\mathcal{X}} [\bar{X}_{k'}(u)] = \sum_{k'=0}^{K-1} r_{k'}^{(k)} x_{k'}(u) = y^{(k)}(u). \quad (\text{C42})$$

If we have access to only a single instance of $\bar{\mathbf{X}}(u)$, however, and thus only one estimate $\bar{y}^{(k)}(u)$ [as $y^{(k)}(u)$ and $\bar{y}^{(k)}(u)$ depicted in Fig. 10], it is useful to know the expected error in this estimate. This error can be extracted from Eq. (C11). In particular, requiring $\mathbf{Y}^{(k)} = \mathbf{T}^T \mathbf{r}^{(k)}$, we have

$$\frac{\|\Lambda^{1/2} \mathbf{T}^T \mathbf{r}^{(k)} - \Lambda^{1/2} \mathbf{Y}^{(k)}\|^2 + \frac{1}{S} \mathbf{r}^{(k)T} \mathbf{V} \mathbf{r}^{(k)}}{\mathbf{Y}^{(k)T} \Lambda \mathbf{Y}^{(k)}} = \frac{1}{S} \mathbf{r}^{(k)T} \mathbf{V} \mathbf{r}^{(k)} = \frac{\beta_k^2}{S}. \quad (\text{C43})$$

This mean-squared error in using $\bar{y}^{(k)}(u)$ to estimate $y^{(k)}(u)$ over the domain of u decreases to zero for $S \rightarrow \infty$ as expected, since the noise in $\bar{\mathbf{X}}$ decreases with S . However, β_k^2 defines the S -independent contribution to the error. In particular, this indicates that, at a given S , certain functions with lower NSR eigenvalues β_k^2 may be better approximated using this physical system than others. We present in Fig. 10 the measured features $\bar{\mathbf{X}}$, the eigentasks \mathbf{y} , and their S -finite version $\bar{\mathbf{y}}$ in a six-qubit Hamiltonian-based system. The associated NSR spectrum, resolvable expressive capacity, and total correlations are also depicted for both CS ($J \neq 0$) and PS ($J = 0$) encodings.

4. S -shot resolvable expressive capacity: Derivation of the bound

Given an arbitrary set of complete orthonormal basis functions $f_\ell(u) = \sum_{j=0}^{\infty} (\mathbf{Y}_\ell)_j u^j$,

$$\langle f_\ell, f_{\ell'} \rangle_p = (\Lambda^{1/2} \mathbf{Y}_\ell)^T (\Lambda^{1/2} \mathbf{Y}_{\ell'}) = \delta_{\ell \ell'}. \quad (\text{C44})$$

The total capacity is independent of the basis choice

$$\begin{aligned} C_T(S) &= \sum_{\ell=0}^{\infty} C[f_\ell] = \sum_{\ell=0}^{\infty} \mathbf{Y}_\ell^T \Lambda^{1/2} \left(\Lambda^{1/2} \mathbf{T}^T \left(\mathbf{T} \Lambda \mathbf{T}^T + \frac{1}{S} \mathbf{V} \right)^{-1} \mathbf{T} \Lambda^{1/2} \right) \Lambda^{1/2} \mathbf{Y}_\ell \\ &= \text{Tr} \left(\Lambda^{1/2} \mathbf{T}^T \left(\mathbf{T} \Lambda \mathbf{T}^T + \frac{1}{S} \mathbf{V} \right)^{-1} \mathbf{T} \Lambda^{1/2} \right) = \text{Tr} \left(\left(\mathbf{G} + \frac{1}{S} \mathbf{V} \right)^{-1} \mathbf{G} \right) = \sum_{k=0}^{K-1} \frac{1}{1 + \frac{\beta_k^2}{S}}. \end{aligned} \quad (\text{C45})$$

5. Eigentask learning training procedure for a nonlinear postprocessing layer

The definition of REC metric and the training scheme underlying eigentask learning considered in the main text is based on a linear estimator $\mathbf{W} \cdot \bar{\mathbf{X}}$ fed into a quadratic loss function. Note that, in an experimental context, their calculations are performed on a classical processor after the measurement results are collected. This is what is done in the experiments on the superconducting quantum processor in Sec. IV. This choice of a linear estimator and a quadratic loss function may seem arbitrary, but the rationale behind it, as explained in the main text, is the desire to quantify the function expression capacity of solely the physical system itself rather than the classical postprocessing layer.

The eigentask learning training methodology introduced in the present work is, however, sufficiently general to be adapted to nonlinear postprocessing scenarios as well, which is the subject of this appendix. The central finding is that a cumulant expansion of the nonlinear loss function produces extra regularization terms whose magnitudes can be characterized by the NSR spectra $\{\beta_k^2\}$. The nonlinear training loss can then be well approximated using a truncated set of eigentasks. To be more specific, we demonstrate that those eigentasks $y^{(k)}$ whose corresponding β_k^2/S is larger should *qualitatively* contribute a larger penalty to the loss function.

The most general case of the output layer can involve a nonlinear activation function or kernel which may subsequently be fed into a nonlinear loss function. A unified

description of these two aspects can be achieved through the usage of a differentiable nonlinear function σ_{NL} :

$$\mathcal{L} = \mathbb{E}_u[\mathbb{E}_{\mathcal{X}}[\sigma_{\text{NL}}(\bar{\mathbf{X}})]] . \quad (\text{C46})$$

For our proof, we do not consider the most general form of σ_{NL} but employ some reasonable assumptions on it that allow us to qualitatively demonstrate the role of eigentasks in the presence of nonlinear postprocessing. First, we assume that it is legitimate to truncate a series expansion of σ_{NL} to second order; that is, all terms of the third- and higher-order derivatives of σ_{NL} are assumed to be much smaller in comparison. Second, we assume that the second-order derivative (namely, the *Hessian matrix*) $\nabla_{\mathbf{x}}\nabla_{\mathbf{x}}^T\sigma_{\text{NL}}(u)$ does not vary too strongly with respect to u . For mean-square loss in Eq. (5), the Hessian matrix of $\nabla_{\mathbf{x}}\nabla_{\mathbf{x}}^T\sigma_{\text{NL}}(u) = 2\mathbf{W}\mathbf{W}^T$ is a constant matrix. Another typical example of σ_{NL} is the cross-entropy loss function of logistic regression used in the toy binary classification problem considered in Sec. IV C. Here, the target function is the conditional probability distribution $f(u) := \text{Prob}[u \in C_1|u]$, where C_1 represents the class labeled by 1. The eventual loss function contains a softmax layer and a cross-entropy function $\mathcal{L} = \mathbb{E}_u[\mathbb{E}_{\mathcal{X}}[H(f(u), \sigma(\mathbf{W} \cdot \bar{\mathbf{X}}(u)))]]$, where σ is sigmoid

function [e.g., softmax function $\sigma(z) = 1/[1 + \exp(-z)]$] and $H(p, q) = -p \ln q - (1-p) \ln(1-q)$ is the cross-entropy. One can check that $\nabla_{\mathbf{x}}\nabla_{\mathbf{x}}^T\sigma_{\text{NL}} = \sigma(\mathbf{W} \cdot \mathbf{x}) \times [1 - \sigma(\mathbf{W} \cdot \mathbf{x})]\mathbf{W}\mathbf{W}^T$, where $\sigma(\mathbf{W} \cdot \mathbf{x})[1 - \sigma(\mathbf{W} \cdot \mathbf{x})] \in [0, 1/4]$ is a bounded function.

Suppose the eigentasks $\bar{\mathbf{y}}$ have been determined by solving the generalized eigenvalue problem Eq. (8). We proceed by expressing the nonlinear loss function Eq. (C46) in terms of $\bar{\mathbf{y}}$. This can be done by expressing $\bar{\mathbf{X}}$ in terms of $\bar{\mathbf{y}}$, $\bar{\mathbf{X}} = \mathbf{\Gamma}\bar{\mathbf{y}}$, with $\mathbf{\Gamma}^T = (\mathbf{r}^{(0)}, \dots, \mathbf{r}^{(K-1)})^{-1}$. This is possible by virtue of $\{\mathbf{r}^{(k)}\}$ being the eigenvectors of the problem Eq. (C33). All noisy measured features $\{\bar{X}_k\}$ can now be expressed in terms of the orthogonal signal basis $\{y^{(k)}\}$ and the noise basis $\{\xi^{(k)}\}$:

$$\bar{X}_{k'}(u) \equiv \sum_{k=0}^{K-1} \Gamma_{k'k} [y^{(k)}(u) + \xi^{(k)}(u)] . \quad (\text{C47})$$

Using a cumulant expansion for the nonlinear loss function and recalling that $\mathbb{E}_{\mathcal{X}}[\xi^{(k)}(u)] = 0$ and $\mathbb{E}_{\mathcal{X}}[\xi^{(k)}\xi^{(k')}] = \mathbf{r}^{(k)T}\mathbf{\Sigma}\mathbf{r}^{(k')}$, where $\mathbf{\Sigma}$ is the covariance of original sampling noise $\boldsymbol{\zeta}$,

$$\begin{aligned} \mathcal{L} &= \mathbb{E}_u[\mathbb{E}_{\mathcal{X}}[\sigma_{\text{NL}}(\bar{\mathbf{X}})]] = \mathbb{E}_u[\mathbb{E}_{\mathcal{X}}[\sigma_{\text{NL}}(\mathbf{\Gamma}\bar{\mathbf{y}})]] = \mathbb{E}_u \left[\mathbb{E}_{\mathcal{X}} \left[\sigma_{\text{NL}} \left(\sum_k \Gamma_{0,k} (y^{(k)} + \xi^{(k)}), \dots, \sum_k \Gamma_{K-1,k} (y^{(k)} + \xi^{(k)}) \right) \right] \right] \\ &= \mathbb{E}_u[\sigma_{\text{NL}}(\mathbf{\Gamma}\mathbf{y})] + \sum_{k=0}^{K-1} \mathbb{E}_u \left[\mathbb{E}_{\mathcal{X}} \left[\frac{\partial \sigma_{\text{NL}}}{\partial y^{(k)}} \xi^{(k)} \right] \right] + \frac{1}{2} \sum_{k_1=0}^{K-1} \sum_{k_2=0}^{K-1} \mathbb{E}_u \left[\mathbb{E}_{\mathcal{X}} \left[\frac{\partial^2 \sigma_{\text{NL}}}{\partial y^{(k_1)} \partial y^{(k_2)}} \xi^{(k_1)} \xi^{(k_2)} \right] \right] + \mathcal{O}\left(\frac{1}{S^2}\right) \\ &= \mathbb{E}_u[\sigma_{\text{NL}}(\mathbf{\Gamma}\mathbf{y})] + \sum_{k=0}^{K-1} \mathbb{E}_u \left[\frac{\partial \sigma_{\text{NL}}}{\partial y^{(k)}} \mathbb{E}_{\mathcal{X}}[\xi^{(k)}] \right] + \frac{1}{2} \sum_{k_1=0}^{K-1} \sum_{k_2=0}^{K-1} \mathbb{E}_u \left[\frac{\partial^2 \sigma_{\text{NL}}}{\partial y^{(k_1)} \partial y^{(k_2)}} \mathbb{E}_{\mathcal{X}}[\xi^{(k_1)} \xi^{(k_2)}] \right] + \mathcal{O}\left(\frac{1}{S^2}\right) \\ &= \mathbb{E}_u[\sigma_{\text{NL}}(\mathbf{\Gamma}\mathbf{y})] + \frac{1}{2} \sum_{k_1=0}^{K-1} \sum_{k_2=0}^{K-1} \mathbb{E}_u \left[\frac{\partial^2 \sigma_{\text{NL}}}{\partial y^{(k_1)} \partial y^{(k_2)}} \mathbf{r}^{(k_1)T} \mathbf{\Sigma} \mathbf{r}^{(k_2)} \right] + \mathcal{O}\left(\frac{1}{S^2}\right) . \end{aligned} \quad (\text{C48})$$

We see here that all terms including third- and higher-order derivatives of σ_{NL} are of $\mathcal{O}(1/S^2)$, which can be neglected in comparison to lower-order terms for large enough S , in accordance with the discussion following Eq. (C46). Second, the slow variation of the second-order derivative $\nabla_{\mathbf{y}}\nabla_{\mathbf{y}}^T\sigma_{\text{NL}}(u) = \mathbf{\Gamma}^T\nabla_{\mathbf{x}}\nabla_{\mathbf{x}}^T\sigma_{\text{NL}}(u)\mathbf{\Gamma}$ with respect to u allows us to make a further simplification of Eq. (C48) by taking the mean-value approximation

$$\mathbb{E}_u \left[\frac{\partial^2 \sigma_{\text{NL}}}{\partial y^{(k_1)} \partial y^{(k_2)}} \mathbf{r}^{(k_1)T} \mathbf{\Sigma} \mathbf{r}^{(k_2)} \right] \approx \mathbb{E}_u \left[\frac{\partial^2 \sigma_{\text{NL}}}{\partial y^{(k_1)} \partial y^{(k_2)}} \right] \mathbb{E}_u \left[\mathbf{r}^{(k_1)T} \mathbf{\Sigma} \mathbf{r}^{(k_2)} \right] \quad (\text{C49})$$

and using the equality $\mathbb{E}_u[\mathbf{r}^{(k_1)T} \mathbf{\Sigma} \mathbf{r}^{(k_2)}] = \delta_{k_1 k_2} \beta_{k_1}^2 / S$:

$$\mathcal{L} \approx \mathbb{E}_u[\sigma_{\text{NL}}(\mathbf{\Gamma}\mathbf{y})] + \sum_{k=0}^{K-1} \frac{\beta_k^2}{S} \cdot \mathbb{E}_u \left[\frac{\partial^2 \sigma_{\text{NL}}}{(\partial y^{(k)})^2} \right] = \mathbb{E}_u[\sigma_{\text{NL}}(\mathbf{\Gamma}\mathbf{y})] + \sum_k \frac{\beta_k^2}{S} \cdot (\mathbf{\Gamma}^T \mathbb{E}_u[\nabla_{\mathbf{x}}\nabla_{\mathbf{x}}^T\sigma_{\text{NL}}]\mathbf{\Gamma})_{kk} . \quad (\text{C50})$$

In typical scenarios, such as, for instance, the case of logistic regression, the loss function depends on the trainable parameters W_k through a linear combination $\mathbf{W} \cdot \bar{\mathbf{X}}(u)$. For such scenarios, it proves convenient to introduce $\mathbf{\Omega}$ such that $\mathbf{\Omega} = \mathbf{\Gamma}^T\mathbf{W}$:

$$\mathbf{W} \cdot \bar{\mathbf{X}}(u) \equiv \sum_{k=0}^{K-1} \Omega_k \cdot [y^{(k)}(u) + \xi^{(k)}(u)]. \quad (\text{C51})$$

For any loss function of the form $\sigma_{\text{NL}}(\mathbf{x}) = g(\mathbf{W} \cdot \mathbf{x})$, the second term in Eq. (C50) can then be expressed as $\mathbf{\Gamma}^T \nabla_{\mathbf{x}} \nabla_{\mathbf{x}}^T \sigma_{\text{NL}} \mathbf{\Gamma} = g''(\mathbf{\Omega} \cdot \mathbf{y}) \mathbf{\Omega} \mathbf{\Omega}^T$. In the case of logistic regression, $\sigma_{\text{NL}}(\mathbf{x}) = g(\mathbf{W} \cdot \mathbf{x}) = H(f, \sigma(\mathbf{W} \cdot \mathbf{x}))$. The final loss function can be further simplified to

$$\begin{aligned} \mathcal{L} \approx & \mathbb{E}_u [H(f, \sigma(\mathbf{\Omega} \cdot \mathbf{y}))] + \left(\sum_{k=0}^{K-1} \frac{\beta_k^2}{S} \Omega_k^2 \right) \\ & \cdot \mathbb{E}_u [\sigma(\mathbf{\Omega} \cdot \mathbf{y})(1 - \sigma(\mathbf{\Omega} \cdot \mathbf{y}))]. \end{aligned} \quad (\text{C52})$$

Here, the first term is the S -infinity value of the loss function, $\lim_{S \rightarrow \infty} \mathcal{L} = \mathbb{E}_u [H(f, \sigma(\mathbf{\Omega} \cdot \mathbf{y}))]$. For the softmax function, we have $\sigma(\mathbf{\Omega} \cdot \mathbf{y}(u))(1 - \sigma(\mathbf{\Omega} \cdot \mathbf{y}(u))) \in [0, 1/4]$, which is assumed not to fluctuate strongly as a function of u , compared to the fluctuation of $\mathbf{r}^{(k_1)T} \mathbf{\Sigma}(u) \mathbf{r}^{(k_2)}$.

While the final result Eq. (C52) is obtained under certain assumptions on σ_{NL} as detailed at the outset, the form of the last term suggests the interpretation of β_k^2/S as a natural regularization and the use of β_k^2/S as a metric for truncation. This truncation is successfully employed in calculations behind the generation of plots in Figs. 5 and 6 for the binary classification task, using the cross-entropy loss function of logistic regression. Equation (C52), therefore, provides some theoretical justification for the use of this truncation scheme when a nonlinear postprocessing layer is employed.

6. Simplifying the noise-to-signal matrix and its eigenproblem for quantum systems

We show that the problem of obtaining the eigentasks for a generic quantum system, and deducing its resolvable expressive capacity under finite measurement resources, can be reduced simply to solving the eigenproblem of its NSR matrix \mathbf{R} [Eq. (C32)]. Note that constructing $\mathbf{R} = \mathbf{G}^{-1/2} \mathbf{V} \mathbf{G}^{-1/2}$ requires computing the inverse of \mathbf{G} . However, \mathbf{G} can have small (although always nonzero) eigenvalues, especially for larger systems, rendering it ill conditioned and making the computation of \mathbf{R} numerically unstable. Fortunately, certain simplifications can be made to derive an equivalent eigenproblem that is much easier to solve. We begin by employing the fact that the second-order-moment matrix \mathbf{D} of multinomial sampling is diagonal. In particular,

$$(\mathbf{D})_{k_1 k_2} = \begin{cases} \sum_{k=0}^{K-1} (\mathbf{G})_{kk_1}, & \text{if } k_1 = k_2 \\ 0, & \text{if } k_1 \neq k_2. \end{cases} \quad (\text{C53})$$

Using $\mathbf{V} = \mathbf{D} - \mathbf{G}$, we can rewrite the eigenproblem for \mathbf{R} :

$$\begin{aligned} \mathbf{R}(\mathbf{G}^{1/2} \mathbf{r}^{(k)}) &= \beta_k^2 \mathbf{G}^{1/2} \mathbf{r}^{(k)} \\ \implies \mathbf{G}^{-1/2} (\mathbf{D} - \mathbf{G}) \mathbf{G}^{-1/2} (\mathbf{G}^{1/2} \mathbf{r}^{(k)}) &= \beta_k^2 \mathbf{G}^{1/2} \mathbf{r}^{(k)} \\ \implies \mathbf{G}^{-1} \mathbf{D} \mathbf{r}^{(k)} &= (1 + \beta_k^2) \mathbf{r}^{(k)}. \end{aligned} \quad (\text{C54})$$

Finally, considering the inverse of the matrix on the left-hand side, we obtain the simplified eigenproblem for the matrix $\mathbf{D}^{-1} \mathbf{G}$:

$$\mathbf{D}^{-1} \mathbf{G} \mathbf{r}^{(k)} = (1 + \beta_k^2)^{-1} \mathbf{r}^{(k)} \equiv \alpha_k \mathbf{r}^{(k)}, \quad (\text{C55})$$

which shares eigenvectors with \mathbf{R} and whose eigenvalues are a simple transformation of the NSR eigenvalues β_k^2 . Importantly, constructing $\mathbf{D}^{-1} \mathbf{G}$ no longer requires calculating any powers of \mathbf{G} , and it relies only on the straightforward inversion of a diagonal matrix \mathbf{D} .

APPENDIX D: SPECTRAL ANALYSIS BASED ON FINITE STATISTICS IN QUANTUM SYSTEMS

While Eq. (C55) is a numerically simpler eigenproblem to solve than Eq. (C32), it still requires the approximation of \mathbf{G} (recall that \mathbf{D} can be obtained from \mathbf{G}) from readout features $\bar{\mathbf{X}}(u)$ under finite sampling of the input (N) and finite shots (S). To be more precise, in experiment one has access to only measured features sampled at finite- S $\bar{\mathbf{X}}$ (indeed, this distinction is the underlying premise of this article). However, in Eq. (8), \mathbf{G} and \mathbf{V} are defined with respect to the ideal \mathbf{x} . Let $\tilde{\mathbf{G}} \equiv \mathbb{E}_u [\mathbb{E}_{\mathcal{X}} [\bar{\mathbf{X}} \bar{\mathbf{X}}^T]]$ and $\tilde{\mathbf{V}} \equiv \mathbb{E}_u [\mathbb{E}_{\mathcal{X}} [\text{diag}(\bar{\mathbf{X}}) - \bar{\mathbf{X}} \bar{\mathbf{X}}^T]]$. The objective of Appendix D 1 is showing that the eigenanalysis $\{\beta_k^2, \mathbf{r}^{(k)}\}$ can be accurately expressed with

$$\beta_k^2 = \frac{S \cdot \tilde{\beta}_k^2}{(S-1) - \tilde{\beta}_k^2}, \quad (\text{D1})$$

and $\mathbf{r}^{(k)} = \tilde{\mathbf{r}}^{(k)}$ from solving generalized eigenvalue problem $\tilde{\mathbf{V}} \tilde{\mathbf{r}}^{(k)} = \tilde{\beta}_k^2 \tilde{\mathbf{G}} \tilde{\mathbf{r}}^{(k)}$. In what follows, we show how an approximation $\tilde{\mathbf{G}}_N$ of \mathbf{G} can be constructed from finitely sampled readout features, as relevant for practical quantum devices. Second, we also describe an approach in Appendix D 2 to obtain the eigentasks $y^{(k)}(u)$ and corresponding NSR eigenvalues β_k^2 in a singular-value decomposition (SVD) problem that avoids explicit construction of the Gram matrix and is therefore numerically even more robust.

1. Approximating eigentasks and NSR eigenvalues under finite S and N

For practical computations, readout features $\bar{\mathbf{X}}(u)$ from the quantum system for finite S can be computed for a discrete set of $u^{(n)} \in [-1, 1]$ for $n = 1, \dots, N$. Labeling the corresponding readout features $\bar{\mathbf{X}}(u^{(n)})$, we can define the *regression matrix* constructed from these readout features:

$$\begin{aligned}\tilde{\mathbf{F}}_N &\equiv (\bar{X}(u^{(1)}), \bar{X}(u^{(2)}), \dots, \bar{X}(u^{(N)}))^T \\ &= \begin{pmatrix} \bar{X}_0(u^{(1)}) & \dots & \bar{X}_{K-1}(u^{(1)}) \\ \vdots & & \vdots \\ \bar{X}_0(u^{(N)}) & \dots & \bar{X}_{K-1}(u^{(N)}) \end{pmatrix}. \quad (\text{D2})\end{aligned}$$

Here, $\tilde{\mathbf{F}}_N \in \mathbb{R}^{N \times K}$, with subscript N indicating its construction from a finite set of N inputs, is a random matrix due to the stochasticity of readout features; in particular, it can be written as

$$\tilde{\mathbf{F}}_N = \mathbf{F}_N + \frac{1}{\sqrt{S}} \mathbf{Z}(\mathbf{F}_N). \quad (\text{D3})$$

where $(\mathbf{F}_N)_{nk} = \mathbb{E}_{\mathcal{X}}[\bar{X}_k(u^{(n)})] = x_k(u^{(n)})$ and \mathbf{Z} is the centered multinomial stochastic process, so that $\mathbb{E}_{\mathcal{X}}[\tilde{\mathbf{F}}_N] = \mathbf{F}_N$.

Using this regression matrix $\tilde{\mathbf{F}}_N$, we can obtain an estimation of the Gram matrix and second-order-moment matrix, which we denote $\tilde{\mathbf{G}}_N$ and $\tilde{\mathbf{D}}_N$, and whose matrix elements are defined via

$$\begin{aligned}(\tilde{\mathbf{G}}_N)_{k_1 k_2} &\equiv \frac{1}{N} \sum_{n=1}^N \bar{X}_{k_1}(u^{(n)}) \bar{X}_{k_2}(u^{(n)}) \\ &= \frac{1}{N} (\tilde{\mathbf{F}}_N^T \tilde{\mathbf{F}}_N)_{k_1 k_2} \approx \int \bar{X}_{k_1}(u) \bar{X}_{k_2}(u) p(u) du, \quad (\text{D4})\end{aligned}$$

$$(\tilde{\mathbf{D}}_N)_{k_1 k_2} \equiv \delta_{k_1, k_2} \frac{1}{N} \sum_{n=1}^N \bar{X}_{k_1}(u^{(n)}) \approx \delta_{k_1, k_2} \int \bar{X}_{k_1}(u) p(u) du. \quad (\text{D5})$$

While the quantities $\tilde{\mathbf{G}}_N$ and $\tilde{\mathbf{D}}_N$ are computed from stochastic readout features, their stochastic contributions are suppressed in the large- N limit by the Hoeffding inequality for sums of bounded stochastic variables.

In what follows, our goal is to prove that, by solving the eigenproblem $\tilde{\mathbf{D}}_N^{-1} \tilde{\mathbf{G}}_N \tilde{\mathbf{r}}_N^{(k)} = (1 + \tilde{\beta}_{N,k}^2)^{-1} \tilde{\mathbf{r}}_N^{(k)}$, the true NSR eigenvalue β_k^2 and eigentask coefficients $\mathbf{r}^{(k)}$ can be well approximated by $S \tilde{\beta}_{N,k}^2 / (S - \tilde{\beta}_{N,k}^2 - 1)$ and $\tilde{\mathbf{r}}_N^{(k)}$, respectively. In particular, to achieve the goal stated above, the first step is still taking the $N \rightarrow \infty$ limit and defining the deterministic limit of $\tilde{\mathbf{G}}_N$ and $\tilde{\mathbf{D}}_N$, according to Eq. (C17), as

$$\tilde{\mathbf{G}} \equiv \lim_{N \rightarrow \infty} \frac{1}{N} (\tilde{\mathbf{F}}_N^T \tilde{\mathbf{F}}_N) = \mathbf{G} + \frac{1}{S} \mathbf{V} = \mathbf{G} + \frac{1}{S} (\mathbf{D} - \mathbf{G}), \quad (\text{D6})$$

$$\tilde{\mathbf{D}} \equiv \lim_{N \rightarrow \infty} \tilde{\mathbf{D}}_N = \mathbf{D}. \quad (\text{D7})$$

In a generic physical system, the covariance matrix \mathbf{V} can be reconstructed by computing the empirical covariance of measured features and employing the well-known rule of

Bessel's correction, $\mathbb{E}_{\mathcal{X}}[(1/S) \sum_s (X_k^{(s)}(u) - \bar{X}_k(u)) \times (X_{k'}^{(s)}(u) - \bar{X}_{k'}(u))] = [(S-1)/S] \boldsymbol{\Sigma}(u)$. In this way, \mathbf{V} can be approximated from the whole record of single-shot random-valued features $\mathcal{X}(u) = \{X_k^{(s)}(u)\}_{k \in [K], s \in [S]}$:

$$\begin{aligned}\lim_{N \rightarrow \infty} \frac{1}{N(S-1)} \sum_{n=1}^N \sum_{s=1}^S (X_k^{(s)}(u^{(n)}) \\ - \bar{X}_k(u^{(n)}))(X_{k'}^{(s)}(u^{(n)}) - \bar{X}_{k'}(u^{(n)})) = \mathbf{V}_{kk'}. \quad (\text{D8})\end{aligned}$$

However, in the scenario of quantum sampling noise, the covariance matrix has a special structure. In particular, $\boldsymbol{\Sigma}$ is not independent of the first-order moments $\mathbf{x}(u)$. This allows us to invert Eqs. (D6) and (D7) to express the Gram matrix \mathbf{G} and second-order-moment matrix \mathbf{D} in terms of the estimators $\tilde{\mathbf{G}}$ and $\tilde{\mathbf{D}}$ computed using a finite number of shots S , in a numerically cheap way without using full readout record of $\mathcal{X}(u) = \{X_k^{(s)}(u)\}_{k \in [K], s \in [S]}$:

$$\mathbf{G} = \frac{S}{S-1} \tilde{\mathbf{G}} - \frac{1}{S-1} \tilde{\mathbf{D}}, \quad (\text{D9})$$

$$\mathbf{D} = \tilde{\mathbf{D}}. \quad (\text{D10})$$

We see that, to lowest order in $1/S$, $\mathbf{G} \approx \tilde{\mathbf{G}}$ and $\mathbf{D} \approx \tilde{\mathbf{D}}$, which is what one might expect naively. However, we clearly see that the estimation of \mathbf{G} can be improved by including a higher-order correction in $1/S$. This contribution arises due to the highly correlated nature of noise and signal for quantum systems: we are able to estimate the noise matrix $\tilde{\mathbf{G}}$ and $\tilde{\mathbf{D}}$ using knowledge of the readout features and correct for the contribution to $\tilde{\mathbf{G}}$ and $\tilde{\mathbf{D}}$ that arises from this noise matrix. We see that this contribution is important in more accurately approximating quantities of interest derived from \mathbf{G} , \mathbf{D} .

To this end, we recall that our ultimate aim is not just to estimate \mathbf{G} and \mathbf{D} but to solve the eigenproblem of Eq. (C55). Using the above relation, we can then establish $\tilde{\mathbf{D}}^{-1} \tilde{\mathbf{G}} = [(S-1)/S] \mathbf{D}^{-1} \mathbf{G} + (1/S) \mathbf{I}$ and write Eq. (C55) in a form entirely in terms of $\tilde{\mathbf{G}}$ and $\tilde{\mathbf{D}}$:

$$\begin{aligned}\mathbf{D}^{-1} \mathbf{G} \mathbf{r}^{(k)} &= (1 + \beta_k^2)^{-1} \mathbf{r}^{(k)} \\ \Rightarrow \tilde{\mathbf{D}}^{-1} \tilde{\mathbf{G}} \mathbf{r}^{(k)} &= \left[\frac{S-1}{S} (1 + \beta_k^2)^{-1} + \frac{1}{S} \right] \mathbf{r}^{(k)}. \quad (\text{D11})\end{aligned}$$

Note that the final form is conveniently another eigenproblem, now for the finite- S matrix $\tilde{\mathbf{D}}^{-1} \tilde{\mathbf{G}}$:

$$\tilde{\mathbf{D}}^{-1} \tilde{\mathbf{G}} \tilde{\mathbf{r}}^{(k)} = (1 + \tilde{\beta}_k^2)^{-1} \tilde{\mathbf{r}}^{(k)} \equiv \tilde{\alpha}_k \tilde{\mathbf{r}}^{(k)}, \quad (\text{D12})$$

whose eigenvalues and eigenvectors can be easily related to the desired eigenvalues β_k^2 and eigenvectors $\mathbf{r}^{(k)}$ in Eq. (C55). Following some algebra, we find




FIG. 8. Eigenanalysis in $L = 5$ H -Ansatz system by taking $S = 10^2$ shots on each of $N = 10^4$ input samples, with true NSR eigenvalues β_k^2 (black), S -finite sampled $\tilde{\beta}_{N,k}^2$ (blue), and corrected NSR $(S \cdot \tilde{\beta}_{N,k}^2) / ((S-1) - \tilde{\beta}_{N,k}^2)$ (purple). $\tilde{\beta}_k^2$, the large- N limit of $\tilde{\beta}_{N,k}^2$, is also plotted in red for comparison. The data correction is necessary, since all $\tilde{\beta}_{N,k}^2$ are below the gray dashed line (representing $S = 10^2$), and the corrected data (in purple) show much better performance even if $\beta_k^2 \gg S$. The corrected data (in purple) have a cutoff at $k = 25$, because in this example all sampled $\tilde{\beta}_{N,k}^2$ with $k > 25$ are larger than $S - 1$ and, hence, are not correctable.

$$\beta_k^2 = \frac{S}{(S-1) - \tilde{\beta}_k^2} \cdot \tilde{\beta}_k^2 = \tilde{\beta}_k^2 + \sum_{j=1}^{\infty} \tilde{\beta}_k^2 (1 + \tilde{\beta}_k^2)^j \left(\frac{1}{S}\right)^j, \quad (\text{D13})$$

$$\mathbf{r}^{(k)} = \tilde{\mathbf{r}}^{(k)}. \quad (\text{D14})$$

From Eq. (D13), we see that, to lowest order in $1/S$, $\beta_k^2 \approx \tilde{\beta}_k^2$. We note that $\tilde{\beta}_k^2$ are always smaller than $S - 1$, and the zeroth-order truncation of Taylor series expression for β_k^2 above is valid only for those k satisfying $\tilde{\beta}_k^2 \ll S - 1$ (see Fig. 8). However, this expression also supplies corrections to higher orders in $1/S$, which are non-negligible for $\tilde{\beta}_k^2 \lesssim S - 1$, as we see in the example in Fig. 8. In contrast, the estimated eigenvectors $\tilde{\mathbf{r}}^{(k)}$ to *any* order in $1/S$ equal the desired eigenvectors $\mathbf{r}^{(k)}$ without any corrections.

Of course, in practice, we do not have access to the matrices $\tilde{\mathbf{G}}$ and $\tilde{\mathbf{D}}$, as these are defined precisely only in the limit where $N \rightarrow \infty$. However, for sufficiently large N , we can approximate these matrices to lowest order by their finite N values, $\tilde{\mathbf{G}} = \tilde{\mathbf{G}}_N + \mathcal{O}(1/N)$ and $\tilde{\mathbf{D}} = \tilde{\mathbf{D}}_N + \mathcal{O}(1/N)$. Then, the eigenproblem in Eq. (D12) can be expressed in the final form

$$\tilde{\mathbf{D}}_N^{-1} \tilde{\mathbf{G}}_N \tilde{\mathbf{r}}_N^{(k)} = (1 + \tilde{\beta}_{N,k}^2)^{-1} \tilde{\mathbf{r}}_N^{(k)} \equiv \tilde{\alpha}_{N,k} \tilde{\mathbf{r}}_N^{(k)}, \quad (\text{D15})$$

where the eigenvalues $\tilde{\beta}_{N,k}^2$ and $\tilde{\alpha}_{N,k}$ and eigenvectors $\tilde{\mathbf{r}}_N^{(k)}$ in the large- N limit must satisfy

$$\lim_{N \rightarrow \infty} \tilde{\beta}_{N,k}^2 = \tilde{\beta}_k^2, \quad \lim_{N \rightarrow \infty} \tilde{\alpha}_{N,k} = \tilde{\alpha}_k, \quad \lim_{N \rightarrow \infty} \tilde{\mathbf{r}}_N^{(k)} = \tilde{\mathbf{r}}^{(k)} \equiv \mathbf{r}^{(k)}. \quad (\text{D16})$$

Here, the invertibility of the empirically computed matrix $\tilde{\mathbf{D}}_N$ required for Eq. (D15) is numerically checked, based on which we can establish a better numerical method in Appendix D 2.

Equation (D15) represents the eigenproblem whose eigenvalues $\tilde{\beta}_{N,k}^2$ and eigenvectors $\tilde{\mathbf{r}}_N^{(k)}$ we actually calculate. For large enough N and under finite S , we can use these as valid approximations to the eigenvalues and eigenvectors in Eq. (D12). According to using Eqs. (D13) and (D14), we are finally able to directly estimate the $N, S \rightarrow \infty$ quantities β_k^2 and $\mathbf{r}^{(k)}$ by the following two quantities:

$$\tilde{\beta}_k^2 \equiv \frac{S \cdot \tilde{\beta}_{N,k}^2}{(S-1) - \tilde{\beta}_{N,k}^2} = \frac{1 - \tilde{\alpha}_{N,k}}{\tilde{\alpha}_{N,k} - \frac{1}{S}}, \quad (\text{D17})$$

$$\tilde{\mathbf{r}}^{(k)} \equiv \tilde{\mathbf{r}}_N^{(k)}. \quad (\text{D18})$$

It is clear that the approximation of β_k^2 to lowest order is an underestimate, as the contribution of order $1/S$ is positive. In Fig. 8, we demonstrate a match between $\tilde{\beta}_k^2 = (S \cdot \tilde{\beta}_{N,k}^2) / ((S-1) - \tilde{\beta}_{N,k}^2)$ and β_k^2 for a wide range of eigentasks with lower order k . In Fig. 9, we plot the estimated eigenvectors $\tilde{\mathbf{r}}_N^{(k)}$ computed under finite statistics ($N = 300, S = 1000$, where these two numbers are relevant for IBM quantum processors) in H encoding, together with the $N, S \rightarrow \infty$ eigenvectors $\mathbf{r}^{(k)}$, and the estimated eigenvalues.

2. Gram matrix-free construction to approximate eigentasks and NSR eigenvalues

If we consider Eq. (D15) and multiply through by $\tilde{\mathbf{D}}_N^{-1/2}$, the resulting equation can be written as an equivalent eigenproblem:

$$\frac{1}{N} \tilde{\mathbf{D}}_N^{-1/2} \tilde{\mathbf{F}}_N^T \tilde{\mathbf{F}}_N \tilde{\mathbf{D}}_N^{-1/2} (\tilde{\mathbf{D}}_N^{1/2} \tilde{\mathbf{r}}_N^{(k)}) = \tilde{\alpha}_{N,k} (\tilde{\mathbf{D}}_N^{1/2} \tilde{\mathbf{r}}_N^{(k)}), \quad (\text{D19})$$

where we also write $\tilde{\mathbf{G}}_N = (1/N) \tilde{\mathbf{F}}_N^T \tilde{\mathbf{F}}_N$ as in the previous section. Note that, as written above, the eigenproblem is entirely equivalent to obtaining the singular-value decomposition of the matrix $(1/\sqrt{N}) \tilde{\mathbf{D}}_N^{-1/2} \tilde{\mathbf{F}}_N^T = U \Sigma V^T$, where $U \in \mathbb{R}^{K \times K}$ and $V \in \mathbb{R}^{N \times N}$ are unitary matrix and Σ is a non-negative diagonal matrix with nonincreasing diagonal entries:

$$\Sigma = \text{diag}(\tilde{\alpha}_{N,0}^{1/2}, \dots, \tilde{\alpha}_{N,K-1}^{1/2}) \approx \text{diag}(\tilde{\alpha}_0^{1/2}, \dots, \tilde{\alpha}_{K-1}^{1/2}). \quad (\text{D20})$$

To obtain the estimation of combination coefficients $\mathbf{r}^{(k)}$, let $\tilde{\mathbf{r}}^{(k)} \in \mathbb{R}^K$ be the normalized left singular vector of

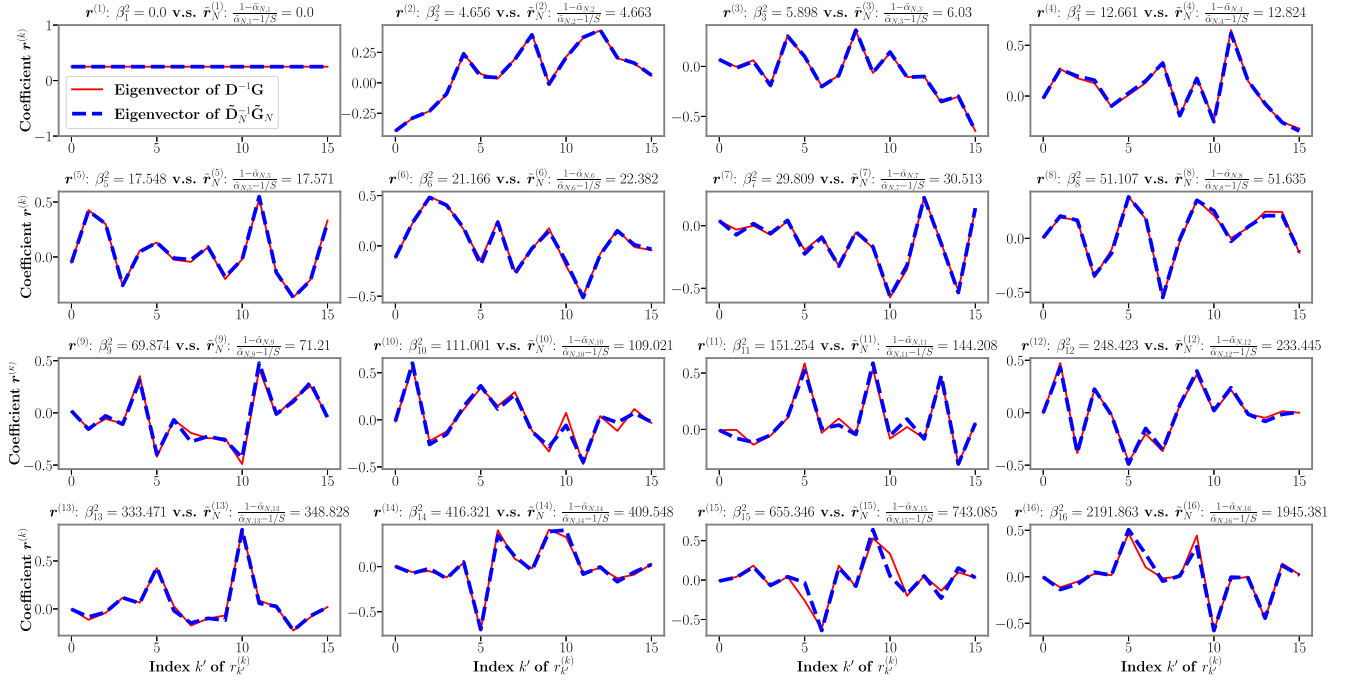


FIG. 9. Estimating noise-to-signal ratio eigenvalues and corresponding eigentask coefficients under finite statistics ($N = 300$, $S = 1000$) in a four-qubit H -encoding system and comparison with theoretical value for $N \rightarrow \infty$, $S \rightarrow \infty$.

$(1/\sqrt{N})\tilde{\mathbf{D}}_N^{-1/2}\tilde{\mathbf{F}}_N^T$ (which is also the eigenvector of $(1/N)\tilde{\mathbf{D}}_N^{-1/2}\tilde{\mathbf{F}}_N^T\tilde{\mathbf{F}}_N\tilde{\mathbf{D}}_N^{-1/2} \approx \mathbf{D}^{-1/2}\tilde{\mathbf{C}}\mathbf{D}^{-1/2}$ in the large- N limit). Then $\mathbf{r}^{(k)} \approx \tilde{\alpha}_{N,k}^{-1/2}\tilde{\mathbf{D}}_N^{-1/2}\tilde{\mathbf{t}}^{(k)} = \tilde{\mathbf{r}}^{(k)} \in \mathbb{R}^K$, and

$$U = (\tilde{\mathbf{t}}^{(0)}, \dots, \tilde{\mathbf{t}}^{(K-1)}) = \tilde{\mathbf{D}}_N^{1/2}(\tilde{\alpha}_{N,0}^{1/2}\tilde{\mathbf{r}}^{(0)}, \dots, \tilde{\alpha}_{N,K-1}^{1/2}\tilde{\mathbf{r}}^{(K-1)}). \quad (\text{D21})$$

Here, $\tilde{\mathbf{r}}^{(k)} = \tilde{\alpha}_{N,k}^{-1/2}\tilde{\mathbf{D}}_N^{-1/2}\tilde{\mathbf{t}}^{(k)}$ can be treated as the combination prefactor of \hat{M}_k , to obtain the observables which correspond to the eigentasks. The merit of an SVD analysis of $(1/\sqrt{N})\tilde{\mathbf{D}}_N^{-1/2}\tilde{\mathbf{F}}_N^T$ is that we need only to work with a K -by- N matrix of features $\tilde{\mathbf{F}}_N$, which is numerically cheaper than further constructing a Gram matrix $(1/N)\tilde{\mathbf{F}}_N^T\tilde{\mathbf{F}}_N$. Therefore,

$$\Sigma V^T = U^T \frac{1}{\sqrt{N}} \tilde{\mathbf{D}}_N^{-1/2} \tilde{\mathbf{F}}_N^T = \frac{1}{\sqrt{N}} \begin{pmatrix} \tilde{\alpha}_{N,0}^{1/2} \tilde{\mathbf{r}}^{(0)T} \\ \vdots \\ \tilde{\alpha}_{N,K-1}^{1/2} \tilde{\mathbf{r}}^{(K-1)T} \end{pmatrix} \tilde{\mathbf{D}}_N^{1/2} \tilde{\mathbf{D}}_N^{-1/2} \tilde{\mathbf{F}}_N^T = \frac{1}{\sqrt{N}} \begin{pmatrix} \tilde{\alpha}_{N,0}^{1/2} \tilde{\mathbf{r}}^{(0)T} \\ \vdots \\ \tilde{\alpha}_{N,K-1}^{1/2} \tilde{\mathbf{r}}^{(K-1)T} \end{pmatrix} \tilde{\mathbf{F}}_N^T. \quad (\text{D22})$$

The entries in Eq. (D22) are

$$(\Sigma V^T)_{k,n} = \frac{1}{\sqrt{N}} \tilde{\alpha}_{N,k}^{1/2} \sum_{n=1}^N \tilde{\mathbf{r}}_k^{(k)}(\tilde{\mathbf{F}}_N)_{n,k'} = \frac{1}{\sqrt{N}} \tilde{\alpha}_{N,k}^{1/2} \sum_{n=1}^N \tilde{\mathbf{r}}_k^{(k)} \bar{X}_{k'}(u^{(n)}) \propto \bar{y}^{(k)}(u^{(n)}). \quad (\text{D23})$$

This means that, for each data sample $u^{(n)}$, the value of the k th-order eigentask $\bar{y}^{(k)}(u^{(n)})$ is exactly the principal component coordinate of $(1/\sqrt{N})\tilde{\mathbf{D}}_N^{-1/2}\tilde{\mathbf{F}}_N^T$, up to a constant factor.

The appearance of the SVD above brings comparisons to a popular, powerful data-compressing tool: *principal component analysis*, or PCA, which is used to project a relatively high-dimensional dataset into a smaller space,

without losing much information. In standard PCA, the original dataset is cast into a feature matrix $F \in \mathbb{R}^{N \times K}$, representing N data samples and K features. Let $\mu_k = (1/N) \sum_n F_{n,k}$ and $\sigma_k^2 = (1/N) \sum_n (F_{n,k} - \mu_k)^2$. Then the standard-scored (or the z -scored) matrix F' is defined by

$$F'_{n,k} = \frac{F_{n,k} - \mu_k}{\sigma_k}. \quad (\text{D24})$$

The SVD of $F'^T = U\Sigma V^T$ gives the well-known principal component analysis. Each row of $\Sigma V^T \in \mathbb{R}^{K \times N}$ is called a *principal component* of the standard score dataset F' . To be more specific, for each data sample labeled by n , the coordinate of its k th principal component is $(\Sigma V^T)_{k,n}$.

While classical PCA focuses on how to *reconstruct* a dataset with minimal representative features, REC analysis focuses on *minimizing* the effect from quantum sampling noise. Since the aim of REC analysis is about the effect from quantum sampling noise, the normalization factor in REC is taking the reciprocal of $\sqrt{N}(\tilde{\mathbf{D}}_N^{1/2})_{kk} = (\sum_n \bar{X}_k(u^{(n)}))^2$, which is quite different from usual PCA, where one uses the reciprocal of the standard deviation σ_k of each feature over the whole dataset samples for normalization, as is described by the form of Eq. (D24).

APPENDIX E: H -ANSATZ QUANTUM SYSTEMS: NSR SPECTRA, RESOLVABLE EXPRESSIVE CAPACITY, AND EIGENTASKS

In this section, we evaluate the REC for quantum systems described by the H Ansatz introduced in Appendix B, as an example of how REC can be efficiently computed for a variety of general quantum systems, and is not just restricted to parametrized quantum circuits. The results of the analysis are compiled in Fig. 10 and discussed below.

Figure 10(a) presents the set of features $\{\bar{X}_k(u)\}$ for typical $L = 6$ qubit CS and PS at $S = 1000$ with randomly

chosen parameters (referred to as encodings; see the caption). The resultant noisy eigentasks $\{\bar{y}^{(k)}(u)\}$ and NSR spectra $\{\beta_k^2\}$ extracted via the eigenvalue analysis are shown in Figs. 10(b) and 10(c), respectively. In the side-by-side comparison in Fig. 10(b), we clearly see the $J = 0$ Ansatz transitioning to a regime with more noise at much lower k than the $J \neq 0$ Ansatz. This is reflected in Fig. 10(c), the β_k^2 spectrum, having a much flatter slope for larger k (note the plot is semilog). Finally, Fig. 10(d) shows the REC of both systems as a function of S . REC rapidly rises for small S for both systems, but the rise of the $J = 0$ system is steeper. After a certain threshold in S , however, the CS grows more rapidly, approaching the upper bound $2^6 = 64$ with $S \sim 10^8$; in contrast, the PS has a significantly lower C_T .

Just like the case of parametrized quantum circuits considered in the main text, we also explore how the REC C_T changes with J for the H Ansatz and compare it to the total correlation ETC \bar{T} , as shown in Fig. 10(f). For $J \rightarrow 0$ we have a PS with $\bar{T} = 0$, whereas in the $J \rightarrow \infty$ we also have $\bar{T} = 0$, because $\hat{\rho}_0 = |0\rangle\langle 0|^{\otimes L}$ is an eigenstate of the encoding $[\hat{\rho}(u) = \hat{\rho}_0]$. This implies there must be a peak at some intermediate J , which for both REC and ETC occurs when the coupling is proportional to the transverse field $J \sim h^x$. At finite S , increased ETC is directly related to a higher REC.

Another interesting aspect is the clear trend seen in the maximization of REC around $J \sim h_{\text{rms}}^x$ for various h_{rms}^x , possibly hinting at the role of increased correlation around

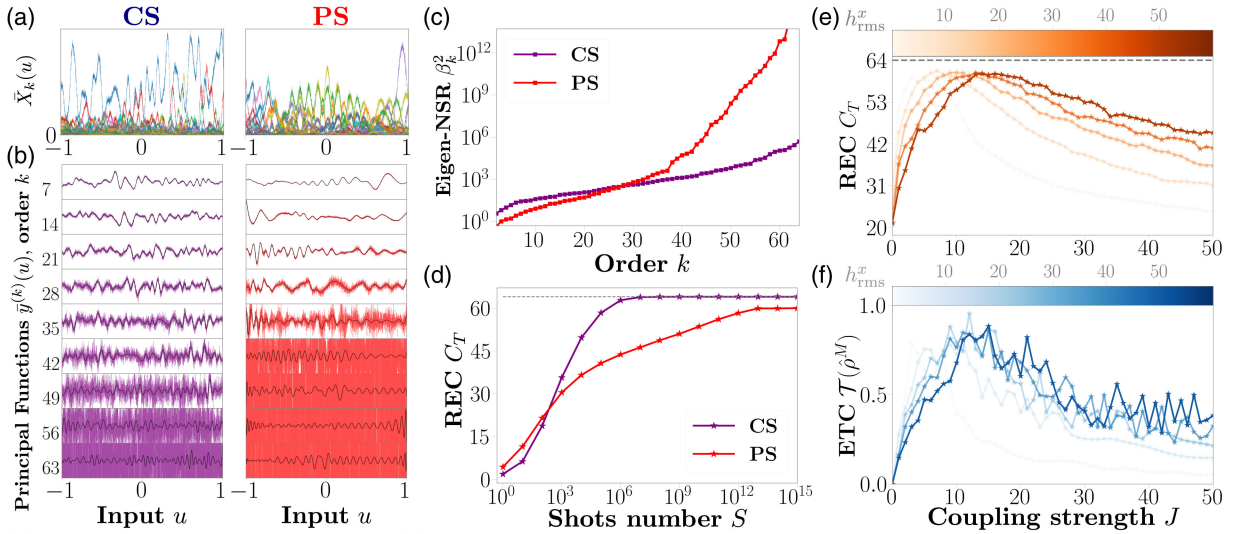


FIG. 10. Eigenanalysis in a six-qubit H -Ansatz system (with $N = 5000$ and $S = 1000$) forming a 1D ring. The Hamiltonian parameters are chosen randomly with zero mean and variance $(h_{\text{rms}}^x, h_{\text{rms}}^z, h_{\text{rms}}^l) = (20, 5, 5)$ and $t = 5$ (see Appendix B for details). Coupling strength is uniformly $J \neq 0$ (correlated system) or $J = 0$ (product system). (a) All $2^L = 64$ noisy features $\bar{X}_k(u)$ and (b) noisy eigentasks $\bar{y}^{(k)}(u) = \mathbf{r}^{(k)} \cdot \bar{\mathbf{X}}(u)$ for selected k from the features in (a), as well as their expected values $y^{(k)}(u) = \lim_{S \rightarrow \infty} \bar{y}^{(k)}(u) = \mathbf{r}^{(k)} \cdot \mathbf{x}(u)$ (black). (c) Noise-to-signal ratio spectrum β_k^2 and (d) C_T vs shots S for both correlated system and product system encodings. (e) C_T at $S = 10^5$ and (f) ETC $\bar{T}(\hat{\rho}^M)$ in a representative random six-qubit H Ansatz, as a function of coupling strength J . The peaks of capacity and correlation coincide, around $J \sim h_{\text{rms}}^x$.

the many-body localization phase transition in random spin systems [56]. This trend is consistent with results in quantum metrology—in general, the SNR obtained from averaging L uncorrelated probes scales as $1/\sqrt{L}$. This scaling can become favorable in the presence of quantum correlation and other nonclassical correlations, in which case the scaling of the SNR can show up as a quadratic improvement $1/L$ [55]. For even larger J , we find that $\hat{\rho}(u) \rightarrow \hat{\rho}_0 = |0\rangle\langle 0|^{\otimes L}$, which clearly reduces \bar{T} , but also C_T as the quantum system state becomes u independent.

APPENDIX F: ANALYTIC SOLUTION TO THE QUANTUM TWO-DESIGN RESOLVABLE EXPRESSIVE CAPACITY

There are many system-specific factors that can contribute to the scaling of resolvable expressive capacity with system size, making it challenging to create a general model that describes all systems. However, we can analytically solve for the REC of a class of quantum models for a specific system: *two-design parametric quantum circuits*

$\{p(\mathbf{u})d\mathbf{u}, \hat{U}(\boldsymbol{\theta}; \mathbf{u})\}$. We clarify that we are referring here to systems with specific parameters $\boldsymbol{\theta}$ which result in two-designs with respect to the input distribution $p(\mathbf{u})$; the ensemble average is taken with respect to inputs u . Quantum literature [49] often refers to general *Ansätze* which form two-designs with respect to parameters $\boldsymbol{\theta}$ instead, which is not what we are considering here.

To be more specific, an ensemble $\{p(\mathbf{u})d\mathbf{u}, \hat{U}(\boldsymbol{\theta}; \mathbf{u})\}$ is a two-design if the following two quantum channels, defined on *any* $2L$ -qubit state $\hat{\tau}$, are equal:

$$\begin{aligned} \mathcal{C}(\hat{\tau}) &= \int \hat{U}(\boldsymbol{\theta}; \mathbf{u})^{\otimes 2} \hat{\tau} (\hat{U}(\boldsymbol{\theta}; \mathbf{u})^\dagger)^{\otimes 2} p(\mathbf{u}) d\mathbf{u} \\ &= \int \hat{U}^{\otimes 2} \hat{\tau} (\hat{U}^\dagger)^{\otimes 2} d\mu_H(\hat{U}), \end{aligned} \quad (\text{F1})$$

where μ_H is the uniform (Haar) measure. We can verify that all information in the Gram matrix is explicitly contained in the elements of $\mathcal{C}(\hat{\rho}_0 \otimes \hat{\rho}_0)$. To be more specific,

$$\begin{aligned} &\langle \mathbf{b}_{k_1}, \mathbf{b}_{k_2} | \mathcal{C}(\hat{\rho}_0 \otimes \hat{\rho}_0) | \mathbf{b}_{k_1}, \mathbf{b}_{k_2} \rangle \\ &= \langle \mathbf{b}_{k_1}, \mathbf{b}_{k_2} | \left(\int (\hat{U}(\boldsymbol{\theta}; \mathbf{u}) \otimes \hat{U}(\boldsymbol{\theta}; \mathbf{u})) | \mathbf{b}_0, \mathbf{b}_0 \rangle \langle \mathbf{b}_0, \mathbf{b}_0 | (\hat{U}(\boldsymbol{\theta}; \mathbf{u})^\dagger \otimes \hat{U}(\boldsymbol{\theta}; \mathbf{u})^\dagger) p(\mathbf{u}) d\mathbf{u} \right) | \mathbf{b}_{k_1}, \mathbf{b}_{k_2} \rangle \\ &= \int |\langle \mathbf{b}_{k_1} | \hat{U}(\boldsymbol{\theta}; \mathbf{u}) | \mathbf{b}_0 \rangle|^2 \cdot |\langle \mathbf{b}_{k_2} | \hat{U}(\boldsymbol{\theta}; \mathbf{u}) | \mathbf{b}_0 \rangle|^2 p(\mathbf{u}) d\mathbf{u} \\ &= \int x_{k_1}(\mathbf{u}) x_{k_2}(\mathbf{u}) p(\mathbf{u}) d\mathbf{u} = (\mathbf{G})_{k_1 k_2}. \end{aligned} \quad (\text{F2})$$

However, $\mathcal{C}(\hat{\rho}_0 \otimes \hat{\rho}_0) = \int U^{\otimes 2} (\hat{\rho}_0 \otimes \hat{\rho}_0) (U^\dagger)^{\otimes 2} d\mu_H(U)$ implies that we can compute the Gram matrix by instead integrating over the Haar measure [61]:

$$(\mathbf{G})_{k_1 k_2} = \int |U_{0, k_1}|^2 |U_{0, k_2}|^2 d\mu_H(U) = \begin{cases} \frac{2}{K(K+1)}, & \text{if } k_1 = k_2, \\ \frac{1}{K(K+1)}, & \text{if } k_1 \neq k_2. \end{cases} \quad (\text{F3})$$

Then, the corresponding second-order-moment matrix \mathbf{D} is given by

$$(\mathbf{D})_{kk} = \frac{2}{K(K+1)} + (K-1) \times \frac{1}{K(K+1)} = \frac{1}{K}. \quad (\text{F4})$$

It is self-consistent that the matrix $\mathbf{D} = \text{diag}(1/K, 1/K, \dots, 1/K)$ obeys the normalization condition $\text{Tr}(\mathbf{D}) = K \cdot (1/K) = 1$. Then, we can solve the eigenvalues $\{\alpha_k\}_{k \in [K]}$ of random walk matrix $(\mathbf{D}^{-1} \mathbf{G})_{k_1 k_2} = [1/(K+1)](1 + \delta_{k_1 k_2})$. It gives

$$\alpha_k = \frac{1}{K+1} (1 + K\delta_{k0}). \quad (\text{F5})$$

Furthermore, we use $\alpha_k = 1/(1 + \beta_k^2)$ or $\beta_k^2 = (1/\alpha_k) - 1$ to compute the NSR eigenvalue:

$$(\beta_0^2, \beta_1^2, \beta_2^2, \dots, \beta_{K-2}^2, \beta_{K-1}^2) = (0, K, K, \dots, K, K). \quad (\text{F6})$$

Then, the resolvable expressive capacity of any two-design system is given by

$$C_T = 1 + \frac{K-1}{1 + \frac{K}{S}} = K \times \frac{1 + \frac{1}{S}}{1 + \frac{K}{S}} = 2^L \times \frac{S+1}{S+2^L}. \quad (\text{F7})$$

APPENDIX G: QUANTUM CORRELATION METRICS

There is no one standard metric to quantify correlation in a many-body state. The metric we would like to utilize here, the *quantum total correlation*, is a quantity inspired by the classical total correlation of L random variables

(b_1, \dots, b_L) , that is, $\sum_{l=1}^L \mathcal{H}(b_l) - \mathcal{H}(b_1, \dots, b_L)$. Using the chain rule of Shannon entropy $\mathcal{H}(b_1, b_2, \dots, b_L) = \mathcal{H}(b_1) + \mathcal{H}(b_2|b_1) + \dots + \mathcal{H}(b_L|b_1, b_2, \dots, b_{L-1})$,

$$\begin{aligned} & \sum_{l=2}^L \mathcal{H}(b_l) - \mathcal{H}(b_1, b_2, \dots, b_L) \\ &= \sum_{l=1}^L \mathcal{H}(b_l) - \sum_{l=1}^L \mathcal{H}(b_l|b_1, b_2, \dots, b_{l-1}) \\ &= \sum_{l=2}^L \mathcal{I}(b_1, \dots, b_{l-1}; b_l) \in [0, L-1], \end{aligned} \quad (\text{G1})$$

we can see that the classical total correlation tells us how a set of random variables reveals information about each other. Similarly, the quantum total correlation can be defined as [52,53]

$$\mathcal{T}(\hat{\rho}) = \sum_{l=1}^L \mathcal{S}(\hat{\rho}_l) - \mathcal{S}(\hat{\rho}), \quad (\text{G2})$$

where \mathcal{S} is von Neumann entropy and $\hat{\rho}_l := \text{Tr}_{[L]\setminus\{l\}}\{\hat{\rho}\}$ is the subsystem state at qubit l . Because of the subadditivity of von Neumann entropy $\sum_{l=1}^L \mathcal{S}(\hat{\rho}_l) \geq \mathcal{S}(\hat{\rho})$, the quantum total correlation is non-negative and is zero iff the state $\hat{\rho} = \otimes_{l=1}^L \hat{\rho}_l$ is a product state.

In this paper's measurement scheme, the specific readout POVMs are the projectors onto the computational states $\{|\mathbf{b}_k\rangle\langle\mathbf{b}_k|\}_{k \in [K]}$. Thus, we are, in particular, interested in analyzing the postmeasurement state $\hat{\rho}^M(u) = \sum_k \rho_{kk}(u) |\mathbf{b}_k\rangle\langle\mathbf{b}_k|$ whose subsystems are correspondingly in states $\hat{\rho}_l^M(u) = \text{Tr}_{[L]\setminus\{l\}}\{\hat{\rho}^M(u)\}$. We compute the average or *expected* quantum total correlation over the input domain u with respect to the input probability distribution $p(u)$:

$$\begin{aligned} \bar{\mathcal{T}}(\hat{\rho}^M) &= \mathbb{E}_u \left[\sum_{l=1}^L \mathcal{S}(\hat{\rho}_l^M(u)) - \mathcal{S}(\hat{\rho}^M(u)) \right] \\ &= \mathbb{E}_u \left[\sum_{l=1}^L \mathcal{H}(b_l(u)) - \mathcal{H}(b_1(u), \dots, b_L(u)) \right], \end{aligned} \quad (\text{G3})$$

where the second equality comes from the diagonal nature of the postmeasurement state which reduces the quantum total correlation to a normal classical total correlation.

The postmeasurement quantum total correlation always reaches its maximum $L-1$ when the postmeasurement state (which just constitutes the diagonal entries of the premeasurement state) is a Greenberger-Horne-Zeilinger-type state. As an additional example, for a W state $|W\rangle = (1/\sqrt{L})(|10\dots 0\rangle + |01\dots 0\rangle + \dots + |00\dots 1\rangle)$, the postmeasurement quantum total correlation $\mathcal{T}(|W\rangle\langle W|)$ is

$$\begin{aligned} & L \left(-\left(\frac{1}{L}\right) \log_2 \left(\frac{1}{L}\right) - \left(\frac{L-1}{L}\right) \log_2 \left(\frac{L-1}{L}\right) \right) \\ & - L \left(-\left(\frac{1}{L}\right) \log_2 \left(\frac{1}{L}\right) \right) = (L-1) \log_2 \left(\frac{L}{L-1}\right), \end{aligned} \quad (\text{G4})$$

which is upper bounded by $\lim_{L \rightarrow \infty} \mathcal{T}(|W\rangle\langle W|) = [1/\ln(2)] \approx 1.443$.

APPENDIX H: REC AND EIGENTASKS FOR CLASSICAL SYSTEMS: A BASIC OPTICAL PNN EXAMPLE

In this appendix, we present some additional details of the REC analysis of the optical setup considered in Sec. IID of the main text. Our focus here is (i) on the details of the sampling noise statistics, which are different when compared to quantum systems, and (ii) the form of the REC eigenproblem for this special case.

1. Sampling noise statistics for a classical optical system

For convenience, we recall the form of the electric field of propagating radiation presented in Sec. IID. The electric field after the SLM can be written generally in the form [41]

$$\begin{aligned} E_0(u; \vec{d}) &= A_0 \cos \left(\frac{\varphi_1(u; \vec{d})}{2} \right) \\ &\times \exp \left\{ i \left(\frac{\varphi_1(u; \vec{d}) + 2\varphi_2(u; \vec{d})}{2} \right) \right\}, \end{aligned} \quad (\text{H1})$$

where $\varphi_l(u; \vec{d})$ are input encoding functions and \vec{d} is the position vector describing the coordinates where the electric field is evaluated in the plane orthogonal to the propagation direction; in particular, $\vec{d} = (q^1, q^2)$. The specific form of the encoding functions is given by

$$\begin{aligned} \varphi_1(u; \vec{d}) &= B(\cos u[A_1(\vec{d}) \cos q^1 + A_2(\vec{d}) \sin q^2] \\ &+ \sin u[A_1(\vec{d}) \sin q^1 + A_2(\vec{d}) \cos q^2]), \end{aligned} \quad (\text{H2a})$$

$$\varphi_2(u; \vec{d}) = Bu(A_1(\vec{d})q^1 + A_2(\vec{d})q^2), \quad (\text{H2b})$$

where we set $B = 3.75$ and $A_{1,2}(\vec{d})$ are fixed *input-independent* spatial mask functions whose values are sampled from a normal distribution with zero mean and unit variance; more precisely, $A_{1,2}(\vec{d}) \sim \mathcal{N}(0, 1)$ for every $\vec{d} = (q^1, q^2)$.

Following the input encoding, the light propagates through a lens and is measured in the lens' focal plane. The electric field in the focal plane $E(u; \vec{d})$ can be shown to be related to the initial field $E_0(u; \vec{d})$ via a Fourier transform [42,43]:

$$E(u; \vec{d}) = \int \int d^2 \vec{d}' E_0(u; \vec{d}') \exp \left\{ \frac{i2\pi}{\lambda f} (\vec{d} \cdot \vec{d}') \right\}, \quad (\text{H3})$$

where $\lambda = (2\pi/k)$ is the wavelength of the propagating field with wave vector k and f is the focal length of the lens being used.

Finally, information must be extracted from this optical system via measurement for its use as a PNN, which also requires us to address the associated measurement noise in a classical setting. We consider photodetection using a camera in the focal plane of the lens. Furthermore, we consider the camera plane as being comprised of a discrete set of $K = P^2$ photodetectors, arranged in an P -by- P square spatial grid, such that the k th photodetector is identified with coordinates $\vec{d}_k = (q_k^1, q_k^2)$. This spatial grid

ultimately defines the coarse-graining level at which the propagating fields can be probed and is set by the spatial resolution of the photodetection apparatus, as expected. Then, it is known [44] that the differential, stochastic photocurrent generated in a given photodetector in a single measurement, which we name for reference $dI(\vec{d}_k, t)$, can be written as a Poisson point process:

$$dI(\vec{d}_k, t) = dN(\vec{d}_k, t), \quad (\text{H4})$$

where $dN(\vec{d}_k, t)$ describes the increment in photodetector counts in a time dt . The stochastic increments, which are independent at different time, has the specific statistical properties:

$$\text{Prob}[dN(\vec{d}_k, t) = 1] = \eta P(\vec{d}_k) dt + o(dt), \quad (\text{H5a})$$

$$\text{Prob}[dN(\vec{d}_k, t) \geq 2] = o(dt), \quad (\text{H5b})$$

$$\mathbb{E}[dN(\vec{d}_k, t) dN(\vec{d}_{k'}, t')] = \mathbb{E}[dN(\vec{d}_k, t)] \mathbb{E}[dN(\vec{d}_{k'}, t')], \quad \text{if } \vec{d}_k \neq \vec{d}_{k'}, \quad (\text{H5c})$$

where expectation values are formally computed over the distribution of Poisson point processes $dN(\vec{d}_k, t)$. Importantly, $P(\vec{d}_k)$ is the power incident on the photodetector with spatial coordinate \vec{d}_k . Then, the expectation value of the increment in counts is directly proportional to the intensity of the incident radiation and a measurement efficiency factor η . The second line defines the fact that the probability of more than a single increment in counts in the time interval dt is $o(dt)$ and, hence, higher order. Finally, the third line indicates that counts on spatially distinct photodetectors and at distinct times are uncorrelated.

The single-shot measured features of this photonic learning scheme become the integrated photocurrent values over an integration time T_{int} (the input u is not written explicitly for notational simplicity):

$$X_k^{(s)} \equiv I(\vec{d}_k) = \int_0^{T_{\text{int}}} dI(\vec{d}_k, t) \in \mathbb{N}, \quad (\text{H6})$$

which are once again stochastic quantities, as they vary from one measurement to the next. This also allows us to easily write down the statistical properties of the integrated photocurrent using Eqs. (H5a)–(H5c):

$$\mathbb{E}[I(\vec{d}_k)] = \eta P(\vec{d}_k) T_{\text{int}}, \quad (\text{H7a})$$

$$\mathbb{E}[I(\vec{d}_k) I(\vec{d}_{k'})] = \mathbb{E}[I(\vec{d}_k)] \mathbb{E}[I(\vec{d}_{k'})], \quad \text{if } \vec{d}_k \neq \vec{d}_{k'}, \quad (\text{H7b})$$

$$\mathbb{E}[I^2(\vec{d}_k)] = \eta P(\vec{d}_k) T_{\text{int}} + (\eta P(\vec{d}_k) T_{\text{int}})^2 = \eta P(\vec{d}_k) T_{\text{int}} + (\mathbb{E}[I(\vec{d}_k)])^2. \quad (\text{H7c})$$

The final expression then provides the variance of the integrated photocurrent for a given photodetector indexed by k :

$$\sigma_1^2(\vec{d}_k) = \mathbb{E}[I^2(\vec{d}_k)] - (\mathbb{E}[I(\vec{d}_k)])^2 = \eta P(\vec{d}_k) T_{\text{int}}. \quad (\text{H8})$$

As an aside, we calculate the measured power signal-to-noise ratio of the integrated photocurrent, SNR_I , which takes the form

$$\text{SNR}_I = \frac{(\mathbb{E}[I(\vec{d}_k)])^2}{\sigma_1^2(\vec{d}_k)} = \eta P(\vec{d}_k) T_{\text{int}} \quad (\text{H9})$$

and, therefore, grows with incident power $P(\vec{d}_k)$ and integration time T_{int} (assuming an unsaturated photodetector).

We now make the connection between the measured photocurrents and the propagating fields reaching the photodetector, described by $E(u; \vec{d}_k)$. More precisely, the power incident on the photodetector is simply set by the Poynting flux of the propagating fields and can be related to the electric field intensity, $P(\vec{d}_k) = \alpha |E(u; \vec{d}_k)|^2$, where α is as introduced in Sec. II D.

The complete input-output map defined above fits within our very general framework. To emphasize this, we now define the measured features $\bar{X}_k(u)$ extracted from this classical machine analogously to the case of measured features extracted from quantum systems, namely, Eq. (1). Precisely, we define $\bar{X}_k(u)$ as averages over individual shots s of integrated photocurrents for each photodetector:

$$\bar{X}_k(u) = \frac{1}{S} \sum_s X_k^{(s)}(u). \quad (\text{H10})$$

We here restore the u dependence of X_k , which arises via the encoded amplitudes and phases in the electric field $E(u; \vec{d}_k)$.

$\bar{X}_k(u)$ are therefore sums over IID random variables $X_k(u)$. Hence, we can directly write for the mean and covariance of these measured features calculated from infinitely many samples:

$$\mathbb{E}_{\mathcal{X}}[\bar{X}_k(u)] = \eta\alpha |E(u; \vec{d}_k)|^2 T_{\text{int}}, \quad (\text{H11a})$$

$$\text{Cov}_{\mathcal{X}}[\bar{X}_j, \bar{X}_k](u) = \frac{1}{S} \delta_{jk} \eta\alpha |E(u; \vec{d}_k)|^2 T_{\text{int}}. \quad (\text{H11b})$$

To connect with our prior notation, we further write

$$\bar{X}_k(u) = x_k(u) + \frac{1}{\sqrt{S}} \zeta_k(u), \quad (\text{H12})$$

which is simply Eq. (3) in the main text. Note that, for most classical machine learning schemes, $S = 1$; however, in our analysis we allow that the shot number S can be any integer. Here, $x_k(u)$ are deterministic quantities defined as

$$x_k(u) \equiv \mathbb{E}_{\mathcal{X}}[\bar{X}_k(u)] = \eta\alpha |E(u; \vec{d}_k)|^2 T_{\text{int}}. \quad (\text{H13})$$

Then, it follows that the remaining term $(1/\sqrt{S})\zeta_k(u)$ is a stochastic process with zero mean [as taking expectation values on both sides of Eq. (H12) demonstrates] and whose second-order moment (equivalent to the variance, as it has

zero mean) encodes the variance of the Poisson point process in one shot of experiment:

$$\begin{aligned} \text{Cov}_{\mathcal{X}}[\bar{X}_j, \bar{X}_k](u) &= \frac{1}{S} \text{Cov}_{\mathcal{X}}[\zeta_j, \zeta_k](u) = \frac{1}{S} \eta\alpha |E(u; \vec{d}_k)|^2 T_{\text{int}} \delta_{jk} \\ &= \frac{1}{S} \delta_{jk} x_k(u), \end{aligned} \quad (\text{H14})$$

where we use Eq. (H13). This finally yields

$$\text{Cov}_{\mathcal{X}}[\zeta_j, \zeta_k](u) \equiv \mathbf{\Sigma}_{jk}(u) = \delta_{jk} x_k(u) \quad (\text{H15})$$

as presented in Sec. II D.

2. REC analysis for a classical optical system

For the Poisson noise process of photodetection, the noise matrix \mathbf{V} is different compared to the case of the multinomial noise process. Fortunately, for the specific case we consider with spatially uncorrelated detectors, \mathbf{V} is, in fact, simpler. We note that

$$\mathbf{V}_{jk} = \int du p(u) \mathbf{\Sigma}_{jk}(u) = \int du p(u) \delta_{jk} x_k(u). \quad (\text{H16})$$

Hence, \mathbf{V} is itself now diagonal. The eigenproblem in question [Eq. (8)] can, therefore, be simplified to

$$\mathbf{V} \mathbf{r}^{(k)} = \beta_k^2 \mathbf{G} \mathbf{r}^{(k)} \implies \mathbf{V}^{-1} \mathbf{G} \mathbf{r}^{(k)} = \frac{1}{\beta_k^2} \mathbf{r}^{(k)}, \quad (\text{H17})$$

where we compute the inverse of \mathbf{V} , since it is a diagonal matrix; we also assume none of its diagonal entries vanish. This is not a strong constraint, since these entries are simply equal to the measured features, which are the sum of intensities incident on the photodetector, and typically are nonzero (assuming integrated photocurrents from any “dead” photodetectors or pixels are excluded from the measured features).

Solving Eq. (H17) allows us to calculate the infinite-shot eigentasks and the C_T as a function of S for the photonic learning system. We can also construct the Gram and covariance matrices using finitely sampled features $\bar{X}_k(u)$ over the input domain, as would be done in a real experiment. One does not need to employ the general Eq. (D8) to estimate the eigentasks. In fact, there is a simpler procedure which is similar to Eqs. (D6) and (D7). The following two quantities can be computed from the S -finite statistics:

$$(\tilde{\mathbf{G}}_N)_{k_1 k_2} \equiv \frac{1}{N} \sum_{n=1}^N \bar{X}_{k_1}(u^{(n)}) \bar{X}_{k_2}(u^{(n)}) = \frac{1}{N} (\tilde{\mathbf{F}}_N^T \tilde{\mathbf{F}}_N)_{k_1 k_2} \approx \int \bar{X}_{k_1}(u) \bar{X}_{k_2}(u) p(u) du, \quad (\text{H18})$$

$$(\tilde{\mathbf{V}}_N)_{k_1 k_2} \equiv \delta_{k_1, k_2} \frac{1}{N} \sum_{n=1}^N \bar{X}_{k_1}(u^{(n)}) \approx \delta_{k_1, k_2} \int \bar{X}_{k_1}(u) p(u) du. \quad (\text{H19})$$

Similar to the scenario in Appendix D 1, we can show that, by solving the eigenproblem $\tilde{\mathbf{V}}_N^{-1} \tilde{\mathbf{G}}_N \tilde{\mathbf{r}}_N^{(k)} = \tilde{\alpha}_{N,k} \tilde{\mathbf{r}}_N^{(k)}$, the true eigen-NSR β_k^2 and eigentask coefficients $\mathbf{r}^{(k)}$ can be well approximated by $1/(\tilde{\alpha}_{N,k} - 1/S)$ and $\tilde{\mathbf{r}}_N^{(k)}$, respectively. The derivation still starts from taking their limits for $N \rightarrow \infty$:

$$\tilde{\mathbf{G}} = \lim_{N \rightarrow \infty} \tilde{\mathbf{G}}_N = \mathbf{G} + \frac{1}{S} \mathbf{V}, \quad (\text{H20})$$

$$\tilde{\mathbf{V}} = \lim_{N \rightarrow \infty} \tilde{\mathbf{V}}_N = \mathbf{V}. \quad (\text{H21})$$

Therefore, $\tilde{\mathbf{G}}$ and $\tilde{\mathbf{V}}$ provide an eigenproblem:

$$\tilde{\mathbf{V}}^{-1} \tilde{\mathbf{G}} \mathbf{r}^{(k)} = \tilde{\alpha}_k \mathbf{r}^{(k)}, \quad (\text{H22})$$

where $\tilde{\alpha}_k = (1/\beta_k^2) + (1/S)$ or, equivalently, $\beta_k^2 = (1/\tilde{\alpha}_k - 1/S)$. It means that, by numerically solving $\tilde{\mathbf{V}}_N^{-1} \tilde{\mathbf{G}}_N \tilde{\mathbf{r}}_N^{(k)} = \tilde{\alpha}_{N,k} \tilde{\mathbf{r}}_N^{(k)}$, one can use

$$\bar{\beta}_k^2 \equiv \frac{1}{\tilde{\alpha}_{N,k} - 1/S}, \quad (\text{H23})$$

$$\bar{\mathbf{r}}^{(k)} \equiv \tilde{\mathbf{r}}_N^{(k)} \quad (\text{H24})$$

to accurately approximate to β_k^2 and $\mathbf{r}^{(k)}$, which finishes the proof of the desired statement.

-
- [1] S. Boyd and L. Chua, *Fading memory and the problem of approximating nonlinear operators with volterra series*, *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **32**, 1150 (1985).
- [2] L. G. Wright, T. Onodera, M. M. Stein, T. Wang, D. T. Schachter, Z. Hu, and P. L. McMahon, *Deep physical neural networks trained with backpropagation*, *Nature (London)* **601**, 549 (2022).
- [3] M. Nakajima, K. Inoue, K. Tanaka, Y. Kuniyoshi, T. Hashimoto, and K. Nakajima, *Physical deep learning with biologically inspired training method: Gradient-free approach for physical hardware*, *Nat. Commun.* **13**, 7847 (2022).
- [4] D. Marković, A. Mizrahi, D. Querlioz, and J. Grollier, *Physics for neuromorphic computing*, *Nat. Rev. Phys.* **2**, 499 (2020).
- [5] G. Tanaka, T. Yamane, J. B. Héroux, R. Nakane, N. Kanazawa, S. Takeda, H. Numata, D. Nakano, and A. Hirose, *Recent advances in physical reservoir computing: A review*, *Neural Netw.* **115**, 100 (2019).
- [6] P. Mujal, R. Martínez-Peña, J. Nokkala, J. García-Beni, G. L. Giorgi, M. C. Soriano, and R. Zambrini, *Opportunities in quantum reservoir computing and extreme learning machines*, *Adv. Quantum Technol.* **4**, 2100027 (2021).
- [7] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, *Variational quantum algorithms*, *Nat. Rev. Phys.* **3**, 625 (2021).
- [8] S. Ortín, M. C. Soriano, L. Pesquera, D. Brunner, D. San-Martín, I. Fischer, C. R. Mirasso, and J. M. Gutiérrez, *A unified framework for reservoir computing and extreme learning machines based on a single time-delayed neuron*, *Sci. Rep.* **5**, 14945 (2015).
- [9] V. Lopez-Pastor and F. Marquardt, *Self-learning machines based on hamiltonian echo backpropagation*, *Phys. Rev. X* **13**, 031020 (2023).
- [10] C. M. Wilson, J. S. Otterbach, N. Tezak, R. S. Smith, A. M. Polloreno, P. J. Karalekas, S. Heidel, M. S. Alam, G. E. Crooks, and M. P. da Silva, *Quantum kitchen sinks: An algorithm for machine learning on near-term quantum computers*, *arXiv:1806.08321*.
- [11] J. García-Beni, G. L. Giorgi, M. C. Soriano, and R. Zambrini, *Scalable photonic platform for real-time quantum reservoir computing*, *Phys. Rev. Appl.* **20**, 014051 (2023).
- [12] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, *Supervised learning with quantum-enhanced feature spaces*, *Nature (London)* **567**, 209 (2019).
- [13] G. E. Rowlands, M.-H. Nguyen, G. J. Ribeill, A. P. Wagner, L. C. G. Govia, W. A. S. Barbosa, D. J. Gauthier, and T. A. Ohki, *Reservoir computing with superconducting electronics*, *arXiv:2103.02522*.
- [14] D. Canaday, A. Griffith, and D. J. Gauthier, *Rapid time series prediction with a hardware-based reservoir computer*, *Chaos* **28**, 123119 (2018).
- [15] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, and M. Soljačić, *Deep learning with coherent nanophotonic circuits*, *Nat. Photonics* **11**, 441 (2017).
- [16] X. Lin, Y. Rivenson, N. T. Yardimci, M. Veli, Y. Luo, M. Jarrahi, and A. Ozcan, *All-optical machine learning using diffractive deep neural networks*, *Science* **361**, 1004 (2018).
- [17] S. Pai, Z. Sun, T. W. Hughes, T. Park, B. Bartlett, I. A. D. Williamson, M. Minkov, M. Milanizadeh, N. Abebe, F. Morichetti, A. Melloni, S. Fan, O. Solgaard, and D. A. B. Miller, *Experimentally realized in situ backpropagation for deep learning in photonic neural networks*, *Science* **380**, 398 (2023).
- [18] J. Dambre, D. Verstraeten, B. Schrauwen, and S. Massar, *Information processing capacity of dynamical systems*, *Sci. Rep.* **2**, 514 (2012).
- [19] F. C. Sheldon, A. Kolchinsky, and F. Caravelli, *Computational capacity of LRC, memristive and hybrid reservoirs*, *Phys. Rev. E* **106**, 045310 (2022).
- [20] M. Schuld, R. Sweke, and J. J. Meyer, *Effect of data encoding on the expressive power of variational quantum-machine-learning models*, *Phys. Rev. A* **103**, 032430 (2021).
- [21] Y. Wu, J. Yao, P. Zhang, and H. Zhai, *Expressivity of quantum neural networks*, *Phys. Rev. Res.* **3**, L032049 (2021).
- [22] L. G. Wright and P. L. McMahon, *The capacity of quantum neural networks*, *arXiv:1908.01364*.
- [23] C. M. Bishop, *Training with noise is equivalent to tikhonov regularization*, *Neural Comput.* **7**, 108 (1995).

- [24] A. Neelakantan, L. Vilnis, Q. V. Le, I. Sutskever, L. Kaiser, K. Kurach, and J. Martens, *Adding gradient noise improves learning for very deep networks*, [arXiv:1511.06807](#).
- [25] H. Noh, T. You, J. Mun, and B. Han, *Regularizing deep neural networks by noise: Its interpretation and optimization*, [arXiv:1710.05179](#).
- [26] O. I. Romyantsev, J. A. Lecoq, O. Hernandez, Y. Zhang, J. Savall, R. Chrapkiewicz, J. Li, H. Zeng, S. Ganguli, and M. J. Schnitzer, *Fundamental bounds on the fidelity of sensory cortical coding*, *Nature (London)* **580**, 100 (2020).
- [27] A few things to note here. (i) The initial state preparation is, in practice, often realized by an act of measurement as well. Then, the input-evolution-output sequence can be described as the sequence of measurement-evolution-measurement sequence. (ii) Some PNN realizations view input as provided through an input state $|\Psi(\mathbf{u})\rangle$. Within the framework we adopt, this can be described as a parametric evolution $\mathcal{U}(\mathbf{u})$ acting on an initial \mathbf{u} -independent state.
- [28] M. A. Nielsen and I. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, England, 2010).
- [29] H. Jaeger, *Short Term Memory in Echo State Networks* (Fraunhofer-Gesellschaft, Munich, 2001).
- [30] J. S. Montijn, G. T. Meijer, C. S. Lansink, and C. M. A. Pennartz, *Population-level neural codes are robust to single-neuron variability from a multidimensional coding perspective*, *Cell Rep.* **16**, 2486 (2016).
- [31] A. A. Faisal, L. P. J. Selen, and D. M. Wolpert, *Noise in the nervous system*, *Nat. Rev. Neurosci.* **9**, 292 (2008).
- [32] G.-B. Huang, Q.-Y. Zhu, and C.-K. Siew, *Extreme learning machine: A new learning scheme of feedforward neural networks*, in *Proceedings of the 2004 IEEE International Joint Conference on Neural Networks* (IEEE Catalog No. 04CH37541) (IEEE, New York, 2004).
- [33] C. Grienberger, A. Giovannucci, W. Zeiger, and C. Portera-Cailliau, *Two-photon calcium imaging of neuronal activity*, *Nat. Rev. Methods Primers* **2**, 67 (2022).
- [34] S. A. Khan, F. Hu, G. Angelatos, and H. E. Türeci, *Physical reservoir computing using finitely-sampled quantum systems*, [arXiv:2110.13849](#).
- [35] R. Martínez-Peña, J. Nokkala, G. L. Giorgi, R. Zambrini, and M. C. Soriano, *Information processing capacity of spin-based quantum reservoir computing systems*, *Cognitive Comput.* **15**, 1440 (2023).
- [36] H. S. Seung, H. Sompolinsky, and N. Tishby, *Statistical mechanics of learning from examples*, *Phys. Rev. A* **45**, 6056 (1992).
- [37] A. Canatar, B. Bordelon, and C. Pehlevan, *Spectral bias and task-model alignment explain generalization in kernel regression and infinitely wide neural networks*, *Nat. Commun.* **12**, 2914 (2021).
- [38] D. Pierangeli, G. Marcucci, and C. Conti, *Photonic extreme learning machine by free-space optical propagation*, *Photonics Res.* **9**, 1446 (2021).
- [39] J. Dong, M. Rafayelyan, F. Krzakala, and S. Gigan, *Optical reservoir computing using multiple light scattering for chaotic systems prediction*, *IEEE J. Sel. Topics Quantum Electron.* **26**, 1 (2020).
- [40] We assume the validity of the parabolic approximation here.
- [41] L. Zhu and J. Wang, *Arbitrary manipulation of spatial amplitude and phase using phase-only spatial light modulators*, *Sci. Rep.* **4**, 7441 (2014).
- [42] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics* (Wiley, New York, 1991).
- [43] A. Yariv and P. Yeh, *Photonics: Optical Electronics in Modern Communications*, sixth ed. (Oxford University, New York, 2007).
- [44] H. M. Wiseman and G. J. Milburn, *Quantum Measurement and Control* (Cambridge University Press, Cambridge, England, 2009).
- [45] J. Wishart, *Cumulants of multivariate multinomial distributions*, *Biometrika* **36**, 47 (1949).
- [46] A. W. Harrow and R. A. Low, *Random quantum circuits are approximate 2-designs*, *Commun. Math. Phys.* **291**, 257 (2009).
- [47] M. Hermans and B. Schrauwen, *Memory in linear recurrent neural networks in continuous time*, *Neural Netw.* **23**, 341 (2010).
- [48] W. D. Kalfus, G. J. Ribeill, G. E. Rowlands, H. K. Krovi, T. A. Ohki, and L. C. G. Govia, *Hilbert space as a computational resource in reservoir computing*, *Phys. Rev. Res.* **4**, 033007 (2022).
- [49] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, *Connecting ansatz expressibility to gradient magnitudes and barren plateaus*, *PRX Quantum* **3**, 010313 (2022).
- [50] S. Sim, P. D. Johnson, and A. Aspuru-Guzik, *Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms*, *Adv. Quantum Technol.* **2**, 1900070 (2019).
- [51] J. J. Meyer, *Fisher information in noisy intermediate-scale quantum applications*, *Quantum* **5**, 539 (2021).
- [52] V. Vedral, *The role of relative entropy in quantum information theory*, *Rev. Mod. Phys.* **74**, 197 (2002).
- [53] K. Modi, T. Paterek, W. Son, V. Vedral, and M. Williamson, *Unified view of quantum and classical correlations*, *Phys. Rev. Lett.* **104**, 080501 (2010).
- [54] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, *Procedure for systematically tuning up cross-talk in the cross-resonance gate*, *Phys. Rev. A* **93**, 060302(R) (2016).
- [55] V. Giovannetti, S. Lloyd, and L. Maccone, *Quantum metrology*, *Phys. Rev. Lett.* **96**, 010401 (2006).
- [56] R. Martínez-Peña, G. L. Giorgi, J. Nokkala, M. C. Soriano, and R. Zambrini, *Dynamical phase transitions in quantum reservoir computing*, *Phys. Rev. Lett.* **127**, 100502 (2021).
- [57] L. Innocenti, S. Lorenzo, I. Palmisano, A. Ferraro, M. Paternostro, and G. M. Palma, *Potential and limitations of quantum extreme learning machines*, *Commun. Phys.* **6**, 118 (2023).
- [58] K. Fujii and K. Nakajima, *Harnessing disordered-ensemble quantum dynamics for machine learning*, *Phys. Rev. Appl.* **8**, 024030 (2017).
- [59] J. Chen, H. I. Nurdin, and N. Yamamoto, *Temporal information processing on noisy quantum computers*, *Phys. Rev. Appl.* **14**, 024065 (2020).
- [60] A. M. Polloreno, *Limits to reservoir learning*, [arXiv:2307.14474](#).
- [61] Z. Puchała and J. Miszczyk, *Symbolic integration with respect to the Haar measure on the unitary groups*, *Bull. Pol. Acad. Sci.: Tech. Sci.* **65**, 21 (2017).