

ARTICLE

OPEN



A machine learning approach to Bayesian parameter estimation

Samuel Nolan¹, Augusto Smerzi¹✉ and Luca Pezzè¹✉

Bayesian estimation is a powerful theoretical paradigm for the operation of the approach to parameter estimation. However, the Bayesian method for statistical inference generally suffers from demanding calibration requirements that have so far restricted its use to systems that can be explicitly modeled. In this theoretical study, we formulate parameter estimation as a classification task and use artificial neural networks to efficiently perform Bayesian estimation. We show that the network's posterior distribution is centered at the true (unknown) value of the parameter within an uncertainty given by the inverse Fisher information, representing the ultimate sensitivity limit for the given apparatus. When only a limited number of calibration measurements are available, our machine-learning-based procedure outperforms standard calibration methods. Our machine-learning-based procedure is model independent, and is thus well suited to "black-box sensors", which lack simple explicit fitting models. Thus, our work paves the way for Bayesian quantum sensors that can take advantage of complex nonclassical quantum states and/or adaptive protocols. These capabilities can significantly enhance the sensitivity of future devices.

npj Quantum Information (2021)7:169; <https://doi.org/10.1038/s41534-021-00497-w>

INTRODUCTION

Precise parameter estimation in quantum systems can revolutionize current technology and prompt scientific discoveries^{1,2}. Prominent examples include gravitational wave detection^{3–5}, time and frequency standards in atomic clocks⁶, field sensing in magnetometers⁷, inertial sensors^{8,9}, and biological imaging¹⁰. As such, improving the sensitivity of quantum sensors is currently an active area of research with most work focused on the control and reduction of noise and decoherence, and on the use of nonclassical probe states¹. Furthermore, the development of data analysis techniques to extract information encoded in complex quantum states^{11–19} is another crucial, yet often overlooked step toward ultra-precise quantum sensing.

Among different strategies^{20–22}, Bayesian parameter estimation (BPE) is known to be particularly efficient and versatile. The output of BPE is a conditional probability distribution $P(\theta|\mu)$ which is interpreted as a degree of belief that the parameter θ equals the true (unknown) value θ_{true} , given the sequence of m measurement results $\mu = \mu_1, \dots, \mu_m$ and any prior information about θ_{true} ^{16,17}. BPE is free of any assumption about the probability distribution of the measurement data and can meaningfully assign a confidence interval to any result, even a single detection event ($m = 1$). As m becomes large, $P(\theta|\mu)$ converges to a Gaussian centered at θ_{true} and with a width proportional to the inverse Fisher information, a result which crucially holds for any probability model and all values of the parameter θ_{true} ^{16,21,22}. Finally, BPE forms the basis of several adaptive protocols in parameter estimation^{23–30}. However, performing BPE necessitates a detailed characterization of the measurement apparatus, which typically requires either modeling the sensor explicitly, or else collecting a prohibitively large amount of calibration data. Although BPE has been demonstrated in single-qubit systems, such as NV center magnetometers^{25,29,31–33}, its demanding calibration requirements remain a major limitation when moving to more complex systems. For example, complex nonclassical states are now routinely generated in ensembles of ultra-cold atoms¹. BPE using entangled states has so far only

limited to some proof-of-principle investigations in few-particle systems^{12,14,15}. To employ BPE in systems that cannot be easily modeled, methods must be developed to efficiently calibrate the device given limited data.

In this manuscript, we provide a machine-learning approach to BPE. We propose that parameter estimation can be formulated as a classification task—similar to the identification of handwritten digits, see Fig. 1—able to be performed efficiently with supervised learning techniques based on artificial neural networks^{34–36}. Classification problems are naturally Bayesian: for instance, the output of the classification network in Fig. 1(a) is the probability that the handwriting is one of the digits 0, ..., 9, in this case, a well-trained network should assign the highest probability to the digit 2. Analogously, we design a neural network adapted for parameter estimation whose output is, naturally, a Bayesian parameter distribution. Based on this interpretation, we provide a theoretical framework that enables a network to be trained using the outcome of individual measurement results. This training provides a set of Bayesian distributions for each possible experimental outcome and a Bayesian prior that we unambiguously identify and directly link to the training of the network. These Bayesian distributions and prior are subsequently multiplied, depending on experimental outcomes, and used to perform BPE for the estimation of an arbitrary unknown parameter. We show that our BPE protocol is asymptotically unbiased and consistent: it obeys relevant Bayesian bounds¹⁷ dictated, in our examples, by quantum and statistical noise. Our method is tested on a variety of quantum states, demonstrating that classical sensitivity limits can be surpassed when using entangled states. Crucially, the neural network needs to be trained with a relatively small amount of data and thus provides a practical advantage over the standard calibration-based BPE.

Although there is a significant body of literature on the application of machine learning techniques to solve problems in quantum science^{37,38}, quantum sensing has received relatively little attention³⁹. Current studies have mainly focused on the

¹QSTAR, INO-CNR and LENS, Largo Enrico Fermi 2, 50125 Firenze, Italy. ✉email: augusto.smerzi@ino.it; luca.pezz@ino.it

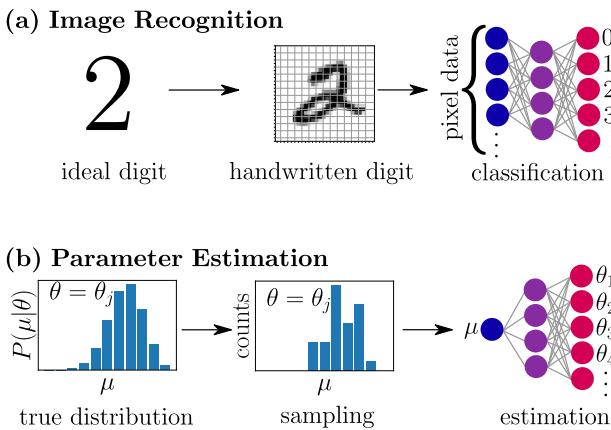


Fig. 1 Parameter estimation as a classification task. **a** By learning the characteristic features of ideal digits directly from training examples⁸⁰, the network can correctly classify handwritten digits with high accuracy. The network provides the conditional probability $P(\text{digit}|\text{data})$ that the image is assigned to a certain ideal digit (2 in this example) given the input pixel data. **b** In parameter estimation, the network input is the result μ of a measurement made at the output of a quantum sensor. In analogy with the digits 0–9, the true (but unknown) probability distribution $P(\mu|\theta)$ represents a category, labeled by the discrete parameter θ_j . A handwritten digit is analogous to a crude sampling from this distribution, used to train the network. Then, the output layer would assign a conditional probability $P(\theta_j|\mu)$ that a particular classification is correct, given the observed result μ .

optimization of adaptive estimation protocols^{40–49}, improved readout for magnetometry^{25,50} and state preparation⁵¹. Similar tasks such as tomography^{52–58}, learning quantum states^{59–64}, Hamiltonian estimation^{65–68}, and state discrimination^{69,70} have also been considered. Neural networks have been applied in the context of parameter estimation with the aim to infer/forecast noisy signals^{71–73}, and for the calibration of a frequentist estimator directly from training data⁷⁴. Unlike these approaches, we show here that a properly trained neural network naturally performs BPE without any assumptions about the system. The machine-learning-based parameter estimation illustrated in this manuscript can be readily applied for data analysis in current quantum sensors, providing all the important advantages of BPE, while enjoying less stringent calibration/training requirements. The method applies to any (mixed or pure) state and measurement observable. In practical applications, noise and decoherence that affect the apparatus are directly included (via the training process) in the Bayesian posterior distributions which therefore fully account for experimental imperfections.

RESULTS

In a general parameter estimation problem, a probe state ρ undergoes a transformation that depends on an unknown parameter θ_{true} . The goal is to estimate θ_{true} from measurements performed on the output state $\rho_{\theta_{\text{true}}}$. A detection event μ occurs with probability $P(\mu|\theta_{\text{true}}) = \text{Tr}[\rho_{\theta_{\text{true}}} E_\mu]$, where $\{E_\mu\}$ is a complete set of positive, $E_\mu \geq 0$, and complete, $\sum E_\mu = 1$ operators⁷⁵.

The parameter estimation discussed in this manuscript is divided in two parts: i) a neural network is trained and ii) Bayesian estimation performed on a test set, which we detail below. A test set refers to an arbitrary sequence of measurement results μ of length m , possibly different to the number of measurements found in the training set. To build intuition we first illustrate the theory with a pedagogical example consisting in the estimation of the rotation angle of a single-qubit state $\exp(-i\sigma_y\theta/2)|\uparrow\rangle$, ($\sigma_{x,y,z}$ are Pauli matrices and $|\uparrow\rangle, |\downarrow\rangle$ are eigenstates of σ_z). The rotation angle θ is

estimated by projecting the output state $|\psi(\theta)\rangle = \exp(-i\sigma_y\theta/2)|\uparrow\rangle$ on σ_z . The two possible output results, $\mu = \uparrow, \downarrow$, can occur with probability $P(\uparrow|\theta) = \cos^2(\theta/2)$ and $P(\downarrow|\theta) = 1 - P(\uparrow|\theta)$, respectively, which are monotonic over the interval $\theta \in [0, \pi]$. Aside of being purely pedagogical, such a system is relevant to NV center magnetometers^{25,29,31–33}. Later, we generalize to systems of many qubits, in separable and entangled states, eventually including noise during state preparation and/or in the output measurement.

Training of the neural network

First, the parameter domain is discretized to form a uniform grid of d points $\theta_1, \dots, \theta_d$ which are assumed to be perfectly known. The training set consists of m_θ measurements performed at each θ_j . For example, the training set for a single qubit would contain d tuples $\{m_{\uparrow,\theta}, m_{\downarrow,\theta}\}$, where $m_{\mu,\theta}$ is the number of times the result $\mu = \uparrow, \downarrow$ was observed at a particular θ . During training, the network is shown all $m_{\text{train}} = \sum_{j=1}^d m_{\theta_j}$ measurement results μ , along with the labels θ_j that are sampled from the (unknown) joint distribution³⁷,

$$P(\mu, \theta_j) = P(\mu|\theta_j)P(\theta_j). \quad (1)$$

Here, $P(\mu|\theta_j)$ is the probability to observe a measurement result μ when the parameter is set to θ_j . This distribution fully characterizes the experimental apparatus (including all sources of noise and decoherence). It is typically unknown to the experimentalist and is never seen by the network. Additionally, the probabilities $P(\mu|\theta_j)$ need not be sampled uniformly in θ_j , which may also have some distribution $P(\theta_j)$.

Via the optimization of weights and links of artificial neurons, the network attempts to learn the conditional probability $P_\Lambda(\theta_j|\mu)$ that gives the degree of certainty that θ_j is the correct label given the particular μ shown during training. This is the essential idea of supervised learning. Here, the subscript Λ denotes the dependence of the output on the randomly chosen initial network, the training algorithm, and the training data itself. In Fig. 2(a) we show the two possible outputs of the network for the single-qubit example: that is $P_\Lambda(\theta_j|\uparrow)$ and $P_\Lambda(\theta_j|\downarrow)$ (blue dots), as a function of the label set $\theta_1, \dots, \theta_d$ in $[0, \pi]$.

Bayesian inversion and prior distribution

Here, we recognize that the output of the neural network, $P_\Lambda(\theta_j|\mu)$, can be interpreted as a Bayesian posterior distribution. As we have discretized the continuous random variable θ , it is necessary to account for the grid spacing $\delta\theta = \theta_d/(d-1)$. We show that the posterior distribution is formally obtained from the Bayes rule,

$$P_\Lambda(\theta_j|\mu) = \frac{P_\Lambda(\mu|\theta_j)P_\Lambda(\theta_j)}{P_\Lambda(\mu)}. \quad (2)$$

We emphasize that the Bayesian inversion in Eq. (2) is performed indirectly by the network, which does not have access to any of the quantities on the right-hand side of Eq. (2). $P_\Lambda(\mu)$ normalizes the posterior distribution, $\sum_{j=1}^d P_\Lambda(\theta_j|\mu)\delta\theta = 1$ and $P_\Lambda(\theta_j)$ is called the prior, which plays a conceptual as well as a practical role. Throughout this manuscript, we are treating possible measurement results μ as a discrete random variable.

We calculate $P_\Lambda(\theta_j)$ from its definition as the marginal distribution, $P_\Lambda(\theta_j) = \sum_\mu P_\Lambda(\theta_j|\mu)P_\Lambda(\mu)$ with the sum extending over all possible measurement results μ . As $P_\Lambda(\mu)$ is also unknown, we can eliminate it by again inserting the marginal expression $P_\Lambda(\mu) = \sum_{k=1}^{d-1} P_\Lambda(\mu|\theta_k)P_\Lambda(\theta_k)\delta\theta$, which results in the implicit integral equation

$$P_\Lambda(\theta_j) = \sum_\mu P_\Lambda(\theta_j|\mu) \sum_{k=1}^d P_\Lambda(\mu|\theta_k)P_\Lambda(\theta_k)\delta\theta \quad (3)$$

Equation (3) is a consistency relation that can be solved for $P_\Lambda(\theta_j)$, given the network output $P_\Lambda(\theta_j|\mu)$ and the likelihood function

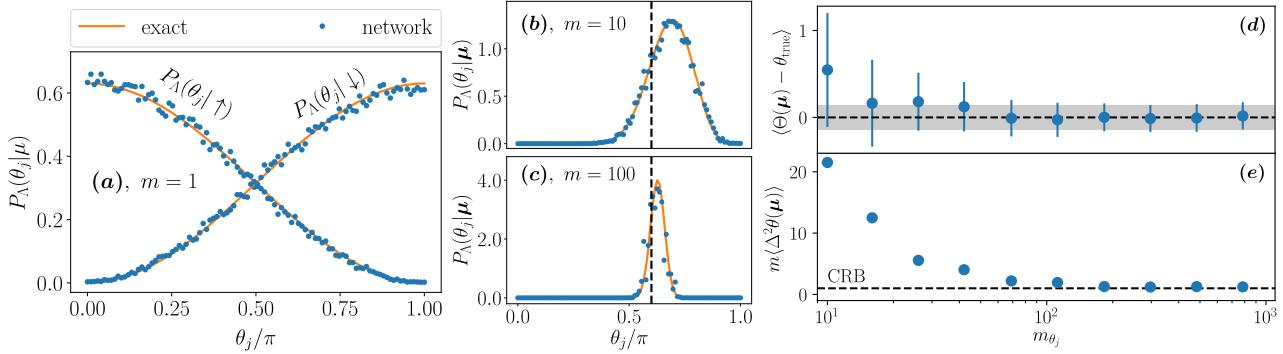


Fig. 2 Bayesian inference performed with a neural network. Here we show results of BPE for the pedagogical example of a single qubit (see Methods and text). The output layer $P_L(\theta_j|\mu)$ of a uniformly-trained network and finite training data (blue dots, details in Methods) compared to the exact Bayesian distribution $P(\theta|\mu) = P(\mu|\theta)P(\theta)/P(\mu)$ (orange line), where $P(\theta) = 1/\pi$ and $P(\mu)$ provides normalization. **a** Bayesian posterior probabilities corresponding to the single-measurement event \uparrow and \downarrow . Panels **(b)** and **(c)** show the Bayesian posterior distributions Eq. (5) for $m = 10$ and $m = 100$ repeated measurement events respectively. These are obtained from the $m = 1$ posterior distribution with Eq. (5). We set $\theta_{\text{true}} = 0.6\pi$ (black dashed vertical lines) and randomly generate a sequence of results: $\mu = \{m_{\uparrow, \theta_{\text{true}}}, m_{\downarrow, \theta_{\text{true}}}\} = \{3, 7\}$ in **(b)**, $\mu = \{29, 71\}$ in **(c)**. **d** Mean value of the maximum a-posterior estimators as a function of the training size m_{θ_j} . The shaded region is the CRB (here $\Delta^2\theta_{\text{CRB}} = 1/m$), and the error bars are the mean posterior variance, shown explicitly in **(e)**. In **(d, e)** we fixed $m = 50$.

$P_L(\mu|\theta_j)$. The relation Eq. (3) can be solved for $P_L(\theta) \equiv p_j$ by recasting it as an eigenvalue problem $\mathbf{A}\mathbf{p} = \mathbf{0}$, for the matrix

$$\mathbf{A}_{jk} = \delta_{jk} - \sum_{\mu} P_L(\theta_j|\mu) P_L(\mu|\theta_k) \delta\theta, \quad (4)$$

where δ_{jk} is the Kronecker delta. To evaluate Eq. (4) the likelihood $P_L(\mu|\theta_k)$ is needed; however, the network only provides $P_L(\theta_j|\mu)$. For a sufficiently well-trained network, we can approximate it with the ideal likelihood distribution, $P_L(\mu|\theta_k) \approx P(\mu|\theta_k)$, which is either known from theory, or else can be well approximated by the relative frequencies observed in the training data $P_L(\mu|\theta_j) \approx m_{\mu, \theta_j}/m_{\theta_j}$. We have found that the prior calculation in Eqs. (3) and (4) is robust to the choice of $P_L(\mu|\theta_k)$.

As shown in Fig. 3, the prior $P_L(\theta_j)$ is determined by the sampling of the training data. For instance, if the training data is distributed uniformly ($m_{\theta_j} = m$ independent of θ), then $P_L(\theta_j)$ is flat, as in Fig. 3 (a, b). A nonflat prior could be achieved by choosing a nonuniform distribution of training measurements. For instance, if m_{train} is the total number of measurements collected in the full training set, the number of measurements m_{θ_j} at each θ_j could be distributed according to $m_{\theta_j} = m_{\text{train}}q(\theta_j)$ where $q(\theta_j)$ is a positive function of θ_j with $\sum_{j=1}^d q(\theta_j) = 1$. In this case, a well-trained network will learn a prior well approximated by $P_L(\theta_j) \approx q(\theta_j)$. Two examples are shown in Fig. 3, panels (c, d) and (e, f). The grid itself could also be varied, resulting in a nonuniform grid spacing $\delta\theta_j = \theta_{j+1} - \theta_j$, which would also result in a nonflat prior. However, this is equivalent to a choice of $q(\theta_j)$ on a uniform grid. This is clearly illustrated by the step function example [Fig. 3(c, d)]. Rather than $q(\theta_j)$ itself being a step function, the same result could be achieved using a flat $q(\theta_j)$ over a grid spanning $[\pi/2, \pi]$ (rather than $[0, \pi]$) but sampled at twice the density. For this reason, we consider only uniform grid spacing throughout this manuscript. The prior thus retains the subjective nature that characterizes the Bayesian formalism: here, this subjectivity is associated with the arbitrariness in the collection of the training data.

Network-based BPE

The training of the network gives access to the single-measurement ($m = 1$) conditional probabilities $P_L(\theta_j|\mu)$ and the prior distribution $P_L(\theta_j)$. We thus proceed with the estimation of an unknown parameter θ_{true} (of course in the numerical experiment θ_{true} is known but this information is never used). Notice that θ_{true} does not need to coincide with one of the grid values θ_j . We sample m random measurement results $\mu = \mu_1, \dots, \mu_m$ from $P(\mu|\theta_{\text{true}})$. The

S. Nolan et al.

Bayesian posterior distribution corresponding to the sequence μ is

$$P_L(\theta_j|\mu) = \mathcal{N} P_L(\theta_j) \prod_{i=1}^m \tilde{P}_L(\theta_j|\mu_i), \quad (5)$$

where $\tilde{P}_L(\theta_j|\mu_i) = P_L(\theta_j|\mu_i)/P_L(\theta_j)$ and \mathcal{N} is the normalization factor. For concreteness, in the single-qubit example, if a sequence of m measurements gives m_{\uparrow} results \uparrow and $m_{\downarrow} = m - m_{\uparrow}$ results \downarrow , the corresponding Bayesian probability distribution is $P_L(\theta_j|\mu) = \mathcal{N} P_L(\theta_j) P_L(\theta_j|\downarrow)^{m_{\downarrow}} \tilde{P}_L(\theta_j|\uparrow)^{m_{\uparrow}}$, see Fig. 2(b, c). Equation (5) represents an update of knowledge about θ_{true} as measurements are collected. Such Bayesian update is based on single-measurement distributions $P_L(\theta_j|\mu)$ and the prior $P_L(\theta_j)$. Indeed, a key advantage of our method is that, while the network is trained with single ($m = 1$) measurement events, the Bayesian analysis can be performed, according to Eq. (5), for arbitrary large m . In other words, we do not need to train the network for each m : the network is trained for $m = 1$, which guarantees the optimal use of training data. We emphasize that the prior $P_L(\theta_j)$ in Eq. (5) is obtained by solving Eq. (3): even for a uniform training, the Eq. (3) gives a better results compared to $P(\theta) = 1/\pi$.

Given $P_L(\theta_j|\mu)$, we can estimate θ_{true} by, for instance,

$$\Theta(\mu) = \arg \max_{\theta_j} P_L(\theta_j|\mu), \quad (6)$$

where the corresponding parameter uncertainty is quantified by the posterior variance

$$\Delta^2\theta(\mu) = \sum_{j=1}^d P_L(\theta_j|\mu) [\Theta(\mu) - \theta_j]^2 \delta\theta, \quad (7)$$

which assigns a confidence interval to any measurement sequence μ . In a sufficiently well-trained network, as the number of measurements m increases, $P_L(\theta_j|\mu)$ converges to the Gaussian distribution^{16,21}

$$P_L(\theta_j|\mu) \approx \sqrt{\frac{m F(\theta_{\text{true}})}{2\pi}} e^{-m F(\theta_{\text{true}})(\theta_j - \theta_{\text{true}})^2/2}, \quad (8)$$

centered at the true value θ_{true} and with variance $1/[m F(\theta_{\text{true}})]$, where

$$F(\theta) = \sum_{\mu} \frac{1}{P(\mu|\theta)} \left(\frac{dP(\mu|\theta)}{d\theta} \right)^2 \quad (9)$$

is the Fisher information. $F(\theta)$ provides a frequentist bound on the precision of a generic estimator $\Delta^2\theta \geq \Delta^2\theta_{\text{CRB}} = 1/[m F(\theta_{\text{true}})]$, called the Cramér-Rao bound. This behavior is clearly exhibited by the

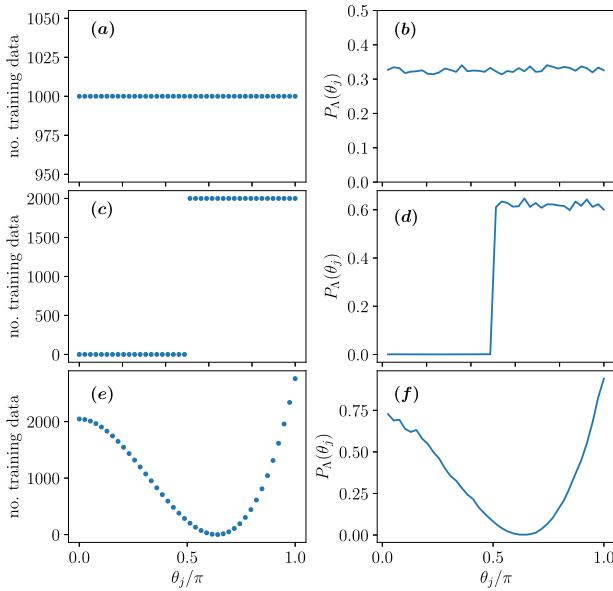


Fig. 3 Prior vs. training distribution. In the left column panels we show examples of distribution of training data, m_{θ_j} as a function of θ_j . The right column panels show the corresponding Bayesian prior distribution $P_A(\theta_j)$. In all three examples the total number of measurements are held fixed. Specifically, in (a, b) m_{θ_j} are distributed uniformly, resulting in a flat prior. In (c, d) data is distributed according to a step function, clearly resulting in a prior that is zero-valued over the part of the domain where no measurements were performed. Finally, in (e, f) a smooth distribution $q(\theta_j)$ is used, which is clearly reproduced by the network.

network in Fig. 2(b, c): the distribution narrows as a function of m and centers around θ_{true} . The result Eq. (8) is valid for a sufficiently dense grid (i.e. $\delta\theta \ll 1/\sqrt{mF(\theta_{\text{true}})}$) and in an appropriate phase interval around θ_{true} . Furthermore, Eq. (8) holds for any prior distribution $P(\theta_j)$, provided that $P(\theta_j)$ is non-vanishing around θ_{true} . By repeating the measurements and using Eq. (5), we can thus gain a factor \sqrt{m} in sensitivity, $\Delta\theta \sim 1/\sqrt{m}$, without requiring either additional training data or additional training for each m . In other words, a single network can be used to provide an estimate for any number of repeated measurements m , limited only by the grid size, meaningful for $\Delta\theta \gg \delta\theta$. In the opposite limit, and thus for $m \gg F(\theta_{\text{true}})/(\delta\theta)^2$, the estimation is biased, namely $|\langle\Theta(\mu) - \theta_{\text{true}}\rangle| \gtrsim \sqrt{\langle\Delta^2\theta\rangle}$. The brackets $\langle \dots \rangle$ denote the average over the likelihood function $P(\mu|\theta_{\text{true}})$. The presence of an asymptotic bias is intrinsic of Bayesian estimation on a finite grid, when θ_{true} does not coincide with one of the grid points. The effect is present also when using ideal probabilities (namely in the limit $m_{\text{train}} \rightarrow \infty$) and it is not associated with the neural network. Of course, insufficient training produces a network that poorly generalizes to larger m . Figure 2(d, e) shows convergence to the expected asymptotic result as a function of the number of training examples m_{θ_j} , for a fixed number of measurement events $m = 50$.

The strategy of classifying a sequence μ following training based on single-measurement results μ only ($\mu = \uparrow, \downarrow$ for the single-qubit case) is a key difference between this work and typical supervised learning problems such as image recognition^{34–36}. With image recognition there is a risk that during training a network will merely memorize the training images, and poorly generalize to unseen images (this is called overfitting). The single-measurement training that we use avoids this problem. Instead, our network is expected to generalize from the single-measurement results seen during training, to sequences with $m > 1$ via Eq. (5). Therefore, the network will never be asked to perform a prediction on an input μ not found in the training set (which will also only ever contain e.g., $\mu = \uparrow, \downarrow$, as

in the single-qubit example). Rather, if the machine-learned Bayesian posterior for the single-measurements μ is noisy or imperfect, this error will quickly compound when Eq. (5) is applied. Therefore, it is important to compute metrics relevant to parameter estimation such as the mean bias or posterior variance (as in Fig. 4).

Application to many-qubit states

In this section, we extend our procedure to systems of N qubits and demonstrate its effectiveness for both separable and entangled states. We introduce the collective spin operators $J_k = \sum_{i=1}^N \sigma_k^{(i)}/2$, where $\sigma_k^{(i)}$ is the k th Pauli matrix for the i th qubit. Making use of these observables, the generalization from a single qubit to many qubits is straightforward: the network is trained to recognize the result of a single J_z measurement with $N+1$ possible outcomes. The Bayesian posterior for many measurements is then obtained from Eq. (5). We consider phase-dependence encoded by a rotation about J_y , which is equivalent to a Mach-Zehnder interferometer¹. In Fig. 4 we apply our method to a coherent-spin state (CSS) $|\text{CSS}\rangle = |\downarrow\rangle^{\otimes N}$ (top panels), a twin-Fock state (TFS) given by the symmetrized combination of $N/2$ spin-up and $N/2$ spin-down particles $|\text{TFS}\rangle = \text{Sym}\{|\downarrow\rangle^{\otimes N/2}, |\downarrow\rangle^{\otimes N/2}\}$ (middle panels), and a depolarized TFS $\rho = (1-\epsilon)|\text{TFS}\rangle\langle\text{TFS}| + \epsilon I/(N+1)$, where I is the identity matrix (in the subspace of permutation-symmetric states) and $\epsilon = 0.1$ (bottom panels). We quantify the performance of the network by the mean posterior variance $\langle\Delta^2\theta(\mu)\rangle$ and bias $\langle\Theta(\mu) - \theta_{\text{true}}\rangle$, averaged over all possible measurement sequences μ . For all three states, Fig. 4 shows that our neural network-based BPE is asymptotically efficient and unbiased when tested on a θ not found in the training grid. As expected for the CSS, the posterior variance saturates the standard quantum limit on average (SQL, $\Delta^2\theta_{\text{SQL}} = 1/[mN]$). Similarly, the TFS posterior variance (7) overcomes the SQL and approaches, on average, the Cramér-Rao bound $\Delta^2\theta_{\text{CRB}} = 1/[mN(N/2 + 1)]$ in the limit of many repeated measurements m . The same is true for the depolarized TFS, demonstrating that our neural network-based BPE is also applicable to mixed states. Furthermore, on average, the estimator (6) gives the true value of the parameter, as expected—so long as the training set is sufficiently large relative to the desired number of measurements m . In particular, networks that are shown more measurements during training are better able to generalize to large m .

Comparison to calibration-based BPE

It is natural to ask how well the network compares to conventional (calibration-based) BPE^{12,14,15} making use of the same training data. Consider a training set where m_{θ_j} measurements are performed at each θ_j , with result μ occurring m_{μ} times at this θ_j . We assume a uniform distribution m_{θ_j} corresponding to a flat prior. The standard approach to either Bayesian or maximum likelihood estimation is to take this data set and estimate the likelihood functions $P(\mu|\theta_j)$ using the relative frequencies $P(\mu|\theta_j) \approx m_{\mu,\theta_j}/m_{\theta_j} \equiv f_{\mu,\theta_j}$, usually aided by some kind of fitting procedure^{12,14}. The posterior distribution $P(\theta_j|\mu)$ is then obtained by choosing a prior $P(\theta_j)$ and applying Bayes theorem $P(\theta_j|\mu) = P(\theta_j) \prod_{i=1}^m P(\mu_i|\theta_j)/P(\mu)$, where $P(\mu)$ provides normalization and $\mu = \mu_1, \dots, \mu_m$ is a measurement sequence. We call this a calibration-based Bayesian analysis. A drawback is that it generally requires collecting a large calibration data set, such that relative frequencies $f_{\mu,\theta}$ well approximate the corresponding probabilities. A further problem is that it is not possible to associate a Bayesian probability to (rare) detection events that did not appear during the calibration, unless the probability is inferred through an arbitrary fit or interpolation procedure. Both issues are overcome by our neural network-based BPE.

In Figure 5, we compare our network-based BPE to the calibration-based BPE. We consider a multipartite entangled

non-Gaussian state (ENG) of $N = 50$ qubits. Entanglement is generated using the one-axis twisting Hamiltonian $H_{\text{OAT}} = \hbar\chi J_z^{276}$, for $\chi = 0.3\pi$ which is in the over-squeezed regime²⁷. Being highly non-Gaussian, it is difficult to aid the calibration with parametric curve fitting. The network on the other hand, is well suited to learning arbitrary probability distributions. Figure 5(a) shows a typical example of a single-shot posterior distribution learned by the network, compared to the relative frequencies in Fig. 5(b). The relative frequencies are intrinsically coarse grained, e.g., in Fig. 5(b) the resolution limit $1/m_{\theta_j}$ is visible, unlike the network which is smooth. In Fig. 5(c, d) we compare the statistically-averaged posterior mean-square error (MSE),

$$\Delta^2\theta_{\text{MSE}}(\mu) = \sum_{j=1}^d P(\theta_j|\mu) (\theta_{\text{true}} - \theta_j)^2 \delta\theta, \quad (10)$$

which quantifies the fluctuations in the deviation of the Bayesian estimate from θ_{true} (see¹⁷ and refs. therein). The posterior MSE is a useful figure of merit in realistic models (either a network or a calibration attempt) because imperfections due to the unavoidable noise in training/calibration data can result in an individual estimate $\Theta(\mu)$ deviating from the true value θ_{true} , even asymptotically. Calibration/training noise can result in positively or negatively biased estimates with equal frequency, which can lead to a deceptively low bias on average (this explains the low bias in Fig. 4 when $m_{\theta_j} = 10$). Figure 5(c, d) clearly show that the neural network outperforms the calibration (see Methods for details), independently of the phase shift θ_{true} or the number of measurements m . As a sanity check, we have verified that the calibration and the network agree well when the training set is large enough. The solid orange curve is the exact result (as would be produced by a perfect calibration/network). This is clear evidence that with limited training/calibration data, our machine learning approach can provide an advantage over conventional

calibration techniques for states that are difficult or impossible to fit. Finally, in Fig. 5(e) we include the effects of finite detection resolution $\Delta\mu$, which is a major limitation in large N systems¹. Modeling of detection noise is discussed in Methods. Although the sensitivity is degraded, network-based BPE continues to outperform calibration-based BPE given equal training/calibration resources, see Methods for details.

DISCUSSION

By reformulating parameter estimation as a classification task, we have shown how to efficiently perform BPE using an artificial neural network with an optimal use of calibration data. The prior distribution—which is the characteristic trait of BPE—is directly linked to the training process: the subjectivity of prior knowledge is reflected by the subjective choice of the training strategy.

BPE offers important advantages, most notably the asymptotic saturation of the frequentist Cramér-Rao bound that holds regardless the statistical model. Indeed, we have demonstrated that our strategy is consistent and efficient for both separable and entangled states of many qubits. Compared to other BPE protocols based on calibration data, our method is the most effective for non-Gaussian states. We found that our neural network-based BPE procedure can outperform standard calibration-based BPE protocols when the training/calibration data is limited and in the absence of an obvious or simple fitting functions. This advantage persists in the presence of finite detection resolution and for noisy probe states. In fact, our approach is the most valuable when the quantum sensor is a *black box*, namely when conditional probabilities of possible measurement results lack an simple explicit model based on a few fitting parameters. In this case, our knowledge about the quantum sensor operations is limited to calibration data.

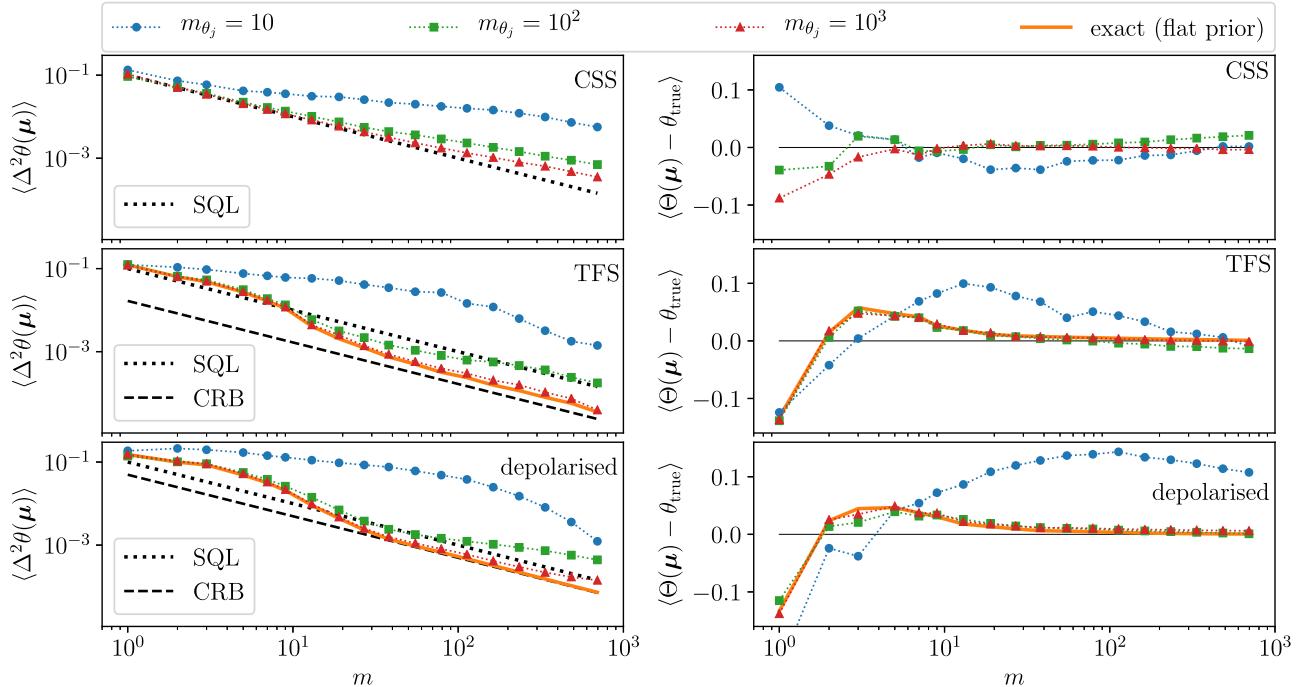


Fig. 4 Consistency and efficiency for many-qubit states. Here we plot the mean Bayesian posterior variance Eq. (7) (left panels) and the bias $\langle \Theta(\mu) - \theta_{\text{true}} \rangle$ (right) as a function of the number of repeated measurements m . Top panels consider a CSS, the middle panels a TFS, and the bottom panels a depolarized TFS, all with $N = 10$ qubits. For all states, three networks are trained with uniform data $m_{\theta_j} = 10$ (blue dots), $m_{\theta_j} = 10^2$ (green squares), $m_{\theta_j} = 10^3$ (red triangles). The solid orange lines are the exact result, obtained from Bayes rule using the true probabilities $P(\mu|\theta)$ and a flat prior. Dashed black lines are the standard quantum limit (SQL) and the frequentist Cramér-Rao bounds (CRB). Here $\theta_{\text{true}} = 0.3\pi$ (not found in the training grid) and all results are averaged over 10^3 randomly generated measurement sequences of length m . See Methods for details on the network parameters.

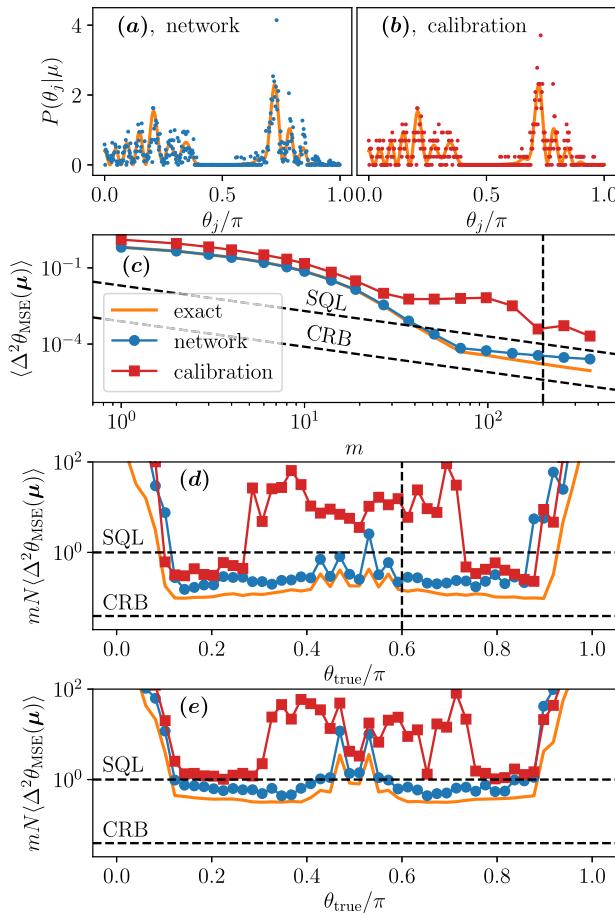


Fig. 5 Comparison with calibration-based BPE. Comparison between neural network-based BPE to calibration-based BPE using the same number of training/calibration measurements for an over-squeezed state of $N = 50$ qubits, as discussed in the main text. **a, b** Example of a single-shot posterior for $\mu = 15$, learned directly by the network **(a)** or inferred from the training data (assuming a flat prior) **(b)**, both from the same set of $m_{\theta_j} = 100$ training/calibration measurements at each phase. **(c, d, e)** The posterior MSE shows the advantage of the neural network procedure over the calibration, for $m_{\theta_j} = 500$. In **(c)** the network is shown to stay closer to the true posterior MSE (solid orange) over a much larger range of m values than the calibration, at a fixed value of the true phase $\theta_{\text{true}} = 0.6\pi$ [not an element of the training grid, vertical black dashed line in **(d)**]. In **(d, e)** we study the performance on a grid of θ_{true} values spanning the entire estimation domain $[0, \pi]$ that were not found in the training grid. In **(d)** the advantage is found to persist over many values of θ_{true} at $m = 200$ shots [vertical black dashed line in **(c)**]. Finally **(e)** includes the effects of finite detection resolution $\Delta\mu^2 = 0.25$, but otherwise parameters are the same as in **(d)**. The likelihood average is approximated by averaging over 10^4 randomly chosen measurement sequences μ . See Methods for details on numerical parameters.

Our neural network-based BPE is readily applicable to current optical and atomic experiments, and therefore could enable BPE with entangled non-Gaussian states in current high precision quantum sensors. Although we focus on single-parameter estimation, our result could also be extended to the simultaneous estimation of multiple parameters.

METHODS

Machine-learning methods

Throughout this manuscript, we employ densely connected, feed-forward neural networks. The networks are implemented and trained using

python-based, open-source package Keras⁷⁸. All hidden layers use ReLU neurons (rectified linear unit). All networks have a single input neuron, which accepts a single, real number μ . The number of hidden layers depends on the system, but for a single qubit a single layer of four neurons is sufficient (see Fig. 2). For larger and more complex states, more layers and neurons can help, as in Figs. 4 or 5. The output layer is d softmax neurons, one for each θ_j grid point, whose value is denoted a_j , which is normalized $\sum_j a_j = 1$ by construction. As we argue in the main text, the output of the network should be interpreted as a Bayesian posterior distribution,

$$a_j = P_A(\theta_j|\mu)\delta\theta. \quad (11)$$

The training process is described in depth elsewhere, see for instance refs. ^{34,36}. Briefly, the network is first initialized with random weights. For efficiency, the training set is randomly divided into subsets called mini-batches. The label θ_j is encoded as a d -dimensional vector whose k th element is a Kronecker delta function δ_{jk} . Each training element in the current mini-batch is fed into the network, and its label is used to evaluate a cost function C . We use the categorical cross-entropy, which for a μ with label θ_j is simply $C = -\log(a_j)$. C is then averaged over the whole mini-batch, and minimized using the ADAM algorithm⁷⁹. This is repeated until the entire training set is exhausted, which is called a training epoch. Typically many epochs are required to reach an optimal network.

Numerical details for figures

In Fig. 2, the network has a single input neuron (which takes as input the result of a single-measurement μ), a single hidden layer of 4 neurons and 100 output neurons (corresponding to a θ grid with 100 grid points). The training set contained $m_{\theta_j} = 10^3$ training measurements per grid point, evenly distributed (corresponding to a flat prior). The network was trained for five epochs with a mini-batch size of 128.

In Fig. 3, networks were trained to perform inference on a single qubit, and have 40 output neurons (corresponding to a θ -grid of 40 points), but otherwise have the same architecture as the network in Fig. 2. Training is performed for 10 epochs with a mini-batch size of 128. The training set contains total of $m_{\text{train}} = 40 \times 10^3$ measurement results.

In Fig. 4, the network trained for coherent-spin states had 1 input neuron, 1 hidden layer of 8 neurons, and 1000 output neurons between $0 \leq \theta_j \leq \pi$. The twin-Fock state network was more complex, 1 input neuron, 2 hidden layers with 32 neurons each, and 1000 output neurons uniformly distributed between $0 \leq \theta_j \leq \pi/2$. Training parameters are adapted to the size of the training set, which is uniform (corresponding to a flat prior). The coherent-spin state training parameters are for $m_{\theta_j} = 10, 100, 1000$: 60 epochs with a min-batch size of 8, 40 with 16, and 20 with 32, respectively. The twin-Fock state training parameters are for $m_{\theta_j} = 10, 100, 1000$: 60 epochs with a min-batch size of 8, 40 with 16, and 30 with 128, respectively.

In Fig. 5, the neural network had three hidden layers with 256 neurons in each, and an output grid with 2000 neurons between $0 \leq \theta_j \leq \pi$. The training was for 60 epochs with a mini-batch size of 1024. The calibration was performed by approximating the likelihood function $P(\mu|\theta)$ by the relative frequencies observed in the training data, smoothed with a cubic interpolation at twice the grid density. The interpolation was performed using `interp1d` from Python's `scipy` package.

Finite detection resolution

Figure 5(e) also includes the effects of finite detector resolution $\Delta\mu$. Following ref. ^{1,16}, detection resolution is modeled as Gaussian noise with variance $\Delta\mu^2$ and mean μ . The probability of measuring the correct result μ is given detector uncertainty $\Delta\mu$ is the convolution $P(\mu|\theta, \Delta\mu) = \sum_{\mu'} C_{\mu'} \exp[-(\mu - \mu')^2/2\Delta\mu^2] P(\mu'|\theta)$ where $C_{\mu'} = \left(\sum_{\mu} \exp[-(\mu - \mu')^2/2\Delta\mu^2]\right)^{-1}$ normalises $P(\mu|\theta, \Delta\mu)$.

DATA AVAILABILITY

The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

CODE AVAILABILITY

Any code used for the current study are available from the corresponding author on reasonable request.

Received: 8 September 2020; Accepted: 20 October 2021;
Published online: 10 December 2021

REFERENCES

1. Pezzè, L., Smerzi, A., Oberthaler, M. K., Schmied, R. & Treutlein, P. Quantum metrology with nonclassical states of atomic ensembles. *Rev. Mod. Phys.* **90**, 035005 (2018).
2. Degen, C. L., Reinhard, F. & Cappellaro, P. Quantum sensing. *Rev. Mod. Phys.* **89**, 035002 (2017).
3. Schnabel, R. Squeezed states of light and their applications in laser interferometers. *Phys. Rep.* **684**, 1–51 (2017).
4. Tse, M. et al. Quantum-enhanced advanced LIGO detectors in the era of gravitational-wave astronomy. *Phys. Rev. Lett.* **123**, 231107 (2019).
5. Acerne, F. Increasing the astrophysical reach of the advanced virgo detector via the application of squeezed vacuum states of light. *Phys. Rev. Lett.* **123**, 231108 (2019).
6. Ludlow, A. D., Boyd, M. M., Ye, J., Peik, E. & Schmidt, P. O. Optical atomic clocks. *Rev. Mod. Phys.* **87**, 637 (2015).
7. Rondin, L. et al. Magnetometry with nitrogen-vacancy defects in diamond. *Rep. Prog. Phys.* **77**, 056503 (2014).
8. Cronin, A. D., Schmiedmayer, J. & Pritchard, D. E. Optics and interferometry with atoms and molecules. *Rev. Mod. Phys.* **81**, 1051 (2009).
9. Barrett, B., Bertoldi, A. & Bouyer, P. Inertial quantum sensors using light and matter. *Phys. Scr.* **91**, 053006 (2016).
10. Taylor, M. & Bowen, W. Quantum metrology and its application in biology. *Phys. Rep.* **615**, 1–59 (2016).
11. Lane, A. S., Braunstein, S. L. & Caves, C. M. Maximum-likelihood statistics of multiple quantum phase measurements. *Phys. Rev. A* **47**, 1667 (1993).
12. Pezzè, L., Smerzi, A., Khoury, G., Hodelin, J. F. & Bouwmeester, D. Phase detection at the quantum limit with multiphoton Mach-Zehnder interferometry. *Phys. Rev. Lett.* **99**, 223602 (2007).
13. Olivares, S. & Paris, M. G. Bayesian estimation in homodyne interferometry. *J. Phys. B: Mol. Opt. Phys.* **42**, 055506 (2009).
14. Krischek, R. et al. Useful multiparticle entanglement and sub-shot-noise sensitivity in experimental phase estimation. *Phys. Rev. Lett.* **107**, 080504 (2011).
15. Xiang, G. Y., Higgins, B. L., Berry, D. W., Wiseman, H. M. & Pryde, G. J. Entanglement-enhanced measurement of a completely unknown optical phase. *Nat. Photonics* **5**, 43 (2011).
16. Pezzè, L. & Smerzi, A. *Quantum Theory of Phase Estimation, in Atom Interferometry, Proceedings of the International School of Physics "Enrico Fermi", Course 188, Varenna*, edited by G. M. Tino and M. A. Kasevich (IOS Press, Amsterdam, 2014), p. 691; arXiv:1411.5164.
17. Li, Y. et al. Frequentist and Bayesian Quantum Phase Estimation. *Entropy* **20**, 628 (2018).
18. Rubio, J., Knott, P. & Dunningham, J. Non-asymptotic analysis of quantum metrology protocols beyond the Cramér-Rao bound. *J. Phys. Commun.* **2**, 015027 (2018).
19. Cimini, V. et al. Diagnosing imperfections in quantum sensors via generalized Cramér-Rao bounds. *Phys. Rev. Appl.* **13**, 024048 (2020).
20. Kay, S. M. *Fundamentals of Statistical Signal Processing: Estimation Theory, Volume I*. (Prentice Hall, Upper Saddle River, NJ, USA, 1993).
21. Lehmann, E. L. & Casella, G. *Theory of Point Estimation*, Springer Texts in Statistics (Springer: New York, 1998).
22. Van Trees, H. L. & Bell, K. L. (eds.). *Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking* (Wiley, New York, NY, USA, 2007).
23. Wiebe, N. & Granade, C. Efficient Bayesian phase estimation. *Phys. Rev. Lett.* **117**, 010503 (2016).
24. Paesani, S. et al. Experimental Bayesian quantum phase estimation on a silicon photonic chip. *Phys. Rev. Lett.* **118**, 100503 (2017).
25. Santagati, R. et al. Magnetic-field learning using a single electronic spin in diamond with one-photon Readout at room temperature. *Phys. Rev. X* **9**, 021019 (2019).
26. Berry, D. W. & Wiseman, H. M. Optimal states and almost optimal adaptive measurements for quantum interferometry. *Phys. Rev. Lett.* **85**, 5098 (2000).
27. Higgins, B. L., Berry, D. W., Bartlett, S. D., Wiseman, H. M. & Pryde, G. J. Entanglement-free Heisenberg-limited phase estimation. *Nature* **450**, 393–396 (2007).
28. Berni, A. A. et al. Ab initio quantum-enhanced optical phase estimation using real-time feedback control. *Nat. Photonics* **9**, 577–581 (2015).
29. Bonato, C. et al. Optimized quantum sensing with a single electron spin using real-time adaptive measurements. *Nat. Nanotechnol.* **11**, 247–252 (2015).
30. Vodola, D. & Müller, M. Adaptive Bayesian phase estimation for quantum error correcting codes. *N. J. Phys.* **21**, 123027 (2019).
31. Hincks, I., Granade, C. & Cory, D. G. Statistical inference with quantum measurements: methodologies for nitrogen vacancy centers in diamond. *N. J. Phys.* **20**, 013022 (2012).
32. Aharon, N. et al. NV center based nano-NMR enhanced by deep learning. *Sci. Rep.* **9**, 17802 (2019).
33. Schwartz, L. et al. Blueprint for nanoscale NMR. *Sci. Rep.* **9**, 6938 (2019).
34. Nielsen, M. A. *Neural Networks and Deep Learning* (Determination Press, 2015), available at <http://neuralnetworksanddeeplearning.com>.
35. Murphy, K. P. *Machine Learning: A Probabilistic Perspective*. (MIT Press, Cambridge, MA, 2012).
36. Metha, P. et al. High-bias, low-variance introduction to Machine Learning for physicists. *Phys. Rep.* **810**, 1–124 (2019).
37. Dunjko, V. & Briegel, H. J. Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Rep. Prog. Phys.* **81**, 074001 (2018).
38. Carleo, G. et al. Machine learning and the physical sciences. *Rev. Mod. Phys.* **91**, 045002 (2019).
39. Polino, E., Valeri, M., Spagnolo, N. & Sciarrino, F. Photonic quantum metrology. *AVS Quantum Sci.* **2**, 024703 (2020).
40. Hentschel, A. & Sanders, B. C. Machine learning for precise quantum measurements. *Phys. Rev. Lett.* **104**, 063603 (2010).
41. Hentschel, A. & Sanders, B. C. Efficient algorithm for optimizing adaptive quantum metrology process. *Phys. Rev. Lett.* **107**, 233601 (2011).
42. Lovett, N. B., Crosnier, C., Perarnau-Llobet, M. & Sanders, B. C. Differential evolution for many-particle adaptive quantum metrology. *Phys. Rev. Lett.* **110**, 220501 (2013).
43. Lumino, A. et al. Experimental phase estimation enhanced by machine learning. *Phys. Rev. Appl.* **10**, 044033 (2018).
44. Xiao, T., Huang, J., Fan, J. & Zeng, G. Continuous-variable quantum phase estimation based on machine learning. *Sci. Rep.* **9**, 12410 (2019).
45. Xu, H. et al. Generalizable control for quantum parameter estimation through reinforcement learning. *npj Quantum Inf.* **9**, 82 (2019).
46. Palittapongarnpim, P. & Sanders, B. Robustness of quantum-enhanced adaptive phase estimation. *Phys. Rev. A* **100**, 012106 (2019).
47. Peng, Y. & Fan, H. Feedback ansatz for adaptive-feedback quantum metrology training with machine learning. *Phys. Rev. A* **101**, 022107 (2020).
48. Schuff, J., Fiderer, L. J. & Braun, D. Improving the dynamics of quantum sensors with reinforcement learning. *N. J. Phys.* **22**, 035001 (2020).
49. Fiderer, L. J., Schuff, J. & Braun, D. Neural-Network Heuristics for adaptive Bayesian quantum estimation. *PRX Quantum* **2**, 020303 (2021).
50. Qian, P. et al. Machine-learning-assisted electron-spin readout of nitrogen-vacancy center in diamond. *Appl. Phys. Lett.* **118**, 084001 (2021).
51. Haine, S. & Hope, J. A Machine-Designed Sensor to Make Optimal Use of Entanglement-Generating Dynamics for Quantum Sensing. *Phys. Rev. Lett.* **124**, 060402 (2020).
52. Gross, D., Liu, Y. K., Flammia, S. T., Becker, S. & Eisert, J. Quantum State Tomography via Compressed Sensing. *Phys. Rev. Lett.* **105**, 150401 (2010).
53. Xu, Q. & Xu, S. Neural network state estimation for full quantum state tomography. arXiv:1811.06654 (2018).
54. Torlai, G. et al. Neural-network quantum state tomography. *Nat. Phys.* **14**, 447–450 (2018).
55. Quek, Y., Fort, S. & Ng, H. K. Adaptive quantum state tomography with neural networks. *npj Quantum Inf.* **7**, 105 (2021).
56. Xin, T. et al. Local-measurement-based quantum state tomography via neural networks. *npj Quantum Inf.* **14**, 109 (2019).
57. Carrasquilla, J., Torlai, G., Melko, R. G. & Aolita, L. Reconstructing quantum states with generative models. *Nat. Mach. Intell.* **1**, 155 (2019).
58. Macarone Palmieri, A. et al. Experimental neural network enhanced quantum tomography. *npj Quantum Inf.* **6**, 20 (2020).
59. Spagnolo, N. et al. Learning an unknown transformation via a genetic approach. *Sci. Rep.* **7**, 14316 (2017).
60. Rocchetto, A. et al. Experimental learning of quantum states. *Sci. Adv.* **5**, 1946 (2019).
61. Yu, S. et al. Reconstruction of a photonic qubit state with reinforcement learning. *Adv. Q. Tech.* **2**, 1800074 (2019).
62. Aaronson, S. The learnability of quantum states. *Proc. R. Soc. A* **463**, 3089 (2007).
63. Torlai, G., Mazzola, G., Carleo, G. & Mezzacapo, A. Precise measurement of quantum observables with neural-network estimators. *Phys. Rev. Research* **2**, 022060(R) (2020).
64. Flurin, E., Martin, L. S., Hacohen-Gourgy, S. & Siddiqi, I. Using a Recurrent Neural Network to Reconstruct Quantum Dynamics of a Superconducting Qubit from Physical Observations. *Phys. Rev. X* **10**, 011006 (2020).
65. Granade, C. E., Ferrie, C., Wiebe, N. & Cory, D. G. Robust online Hamiltonian learning. *N. J. Phys.* **14**, 103013 (2012).
66. Wang, J. et al. Experimental quantum Hamiltonian learning. *Nat. Phys.* **13**, 551 (2017).

67. Wang, D. et al. Machine learning magnetic parameters from spin configurations. *Adv. Sci.* **7**, 2000566 (2020).
68. Wozniakowski, A., Thompson, J., Gu, M. & Binder, F. Boosting on the shoulders of giants in quantum device calibration, *arXiv:2005.06194* (2020).
69. You, C. et al. Identification of light sources using artificial neural networks. *Appl. Phys. Rev.* **7**, 021404 (2020).
70. Gebhart, V. & Bohmann, M. Neural-network approach for identifying non-classicality from click-counting data. *Phys. Rev. Res.* **2**, 023150 (2020).
71. Greplova, E., Andersen, C. K. & Mølmer, K. Quantum parameter estimation with a neural network, *arXiv:1711.05238* (2017).
72. Liu, W. et al. Parameter estimation via weak measurement with machine learning. *J. Phys. B: At., Mol. Optical Phys.* **52**, 045504 (2019).
73. Khanahmadi, M. & Mølmer, K. Time-dependent atomic magnetometry with a recurrent neural network. *Phys. Rev. A* **103**, 032406 (2021).
74. Cimini, V. et al. Calibration of quantum sensors by neural networks. *Phys. Rev. Lett.* **123**, 230502 (2019).
75. Braunstein, S. L. & Caves, C. M. Statistical distance and the geometry of quantum states. *Phys. Rev. Lett.* **72**, 3439 (1994).
76. Kitagawa, M. & Ueda, M. Squeezed spin states. *Phys. Rev. A* **47**, 5138 (1993).
77. Pezzè, L. & Smerzi, A. Entanglement, nonlinear dynamics, and the Heisenberg limit. *Phys. Rev. Lett.* **102**, 100401 (2009).
78. Chollet, F. et al. Keras (2015), available at <http://keras.io>.
79. Kingma, D. P. & Ba, J. Adam: a method for Stochastic optimization, *arXiv:1412.6980* (2014).
80. LeCun, Y., Cortes, C. & Burges, C. (ATT Labs [Online], 2010), available at <http://yann.lecun.com/exdb/mnist>.

ACKNOWLEDGEMENTS

We would like to thank V. Gebhart for useful discussions. We acknowledge funding from the project EMPIR-USOQS, EMPIR projects are co-funded by the European Unions Horizon2020 research and innovation program and the EMPIR Participating States. We also acknowledge financial support from the European Union's Horizon 2020 research and innovation program - Qombs Project, FET Flagship on Quantum Technologies grant no. 820419, and from the H2020 QuantERA ERA-NET Cofund in Quantum Technologies projects QCLOCKS and CEBBEC.

AUTHOR CONTRIBUTIONS

L.P. and A.S. were responsible for the inception of the project, and all authors contributed to its ongoing design and development. S.P.N. wrote the code and performed the numerical analysis presented in this manuscript. All authors contributed to the writing of the manuscript.

COMPETING INTERESTS

The authors declare no competing interests.

ADDITIONAL INFORMATION

Correspondence and requests for materials should be addressed to Augusto Smerzi or Luca Pezzè.

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

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



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

© The Author(s) 2021