



# End-to-end variational quantum sensing



Benjamin MacLellan<sup>1,2,3,4</sup>✉, Piotr Roztock<sup>4</sup>, Stefanie Czischek<sup>5</sup> & Roger G. Melko<sup>1,3</sup>

Harnessing quantum correlations can enable sensing beyond classical precision limits, with the realization of such sensors poised for transformative impacts across science and engineering. Real devices, however, face the accumulated impacts of noise and architecture constraints, making the design and success of practical quantum sensors challenging. Numerical and theoretical frameworks to optimize and analyze sensing protocols in their entirety are thus crucial for translating quantum advantage into widespread practice. Here, we present an end-to-end variational framework for quantum sensing protocols, where parameterized quantum circuits and neural networks form trainable, adaptive models for quantum sensor dynamics and estimation, respectively. The framework is general and can be adapted towards arbitrary qubit architectures, as we demonstrate with experimentally-relevant ansätze for trapped-ion and photonic systems, and enables to directly quantify the impacts that noise and finite data sampling. End-to-end variational approaches can thus underpin powerful design and analysis tools for practical quantum sensing advantage.

By leveraging carefully prepared quantum correlations, sensing protocols that surpass the intrinsic precision bounds of classical approaches can be designed<sup>1</sup>. Such protocols comprise four steps (Fig. 1a): (i) the preparation of a metrologically useful probe state; (ii) the interaction of the probe with an unknown physical parameter of interest  $\phi$ ; (iii) the measurement of the perturbed probe state, producing a dataset of classical detection outcomes; and (iv) the estimation  $\hat{\phi}$  of the physical parameter from the classical dataset. This approach underpins the many applications of quantum sensing, spanning gravitational wave detection<sup>2</sup>, earth science<sup>3</sup>, illumination<sup>4</sup>, microscopy<sup>5</sup>, energy engineering<sup>6</sup>, and chemical dynamics<sup>7</sup>, among others.

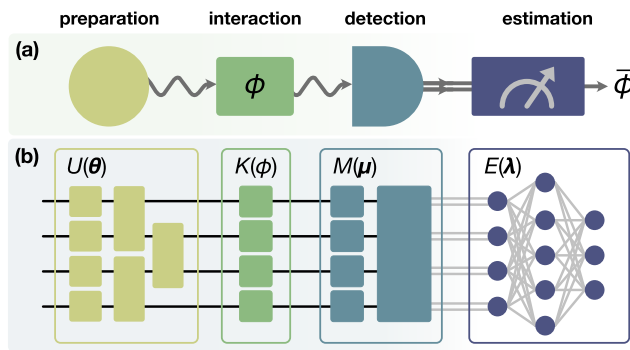
In recent years, approaches for the design and analysis of quantum sensing protocols have rapidly advanced<sup>8–10</sup>. Moreover, as variational and machine learning approaches become ubiquitous across the landscape of quantum technologies<sup>11–15</sup>, recent efforts have begun to explore their impact in the context of quantum sensing<sup>16</sup>. For example, variational quantum-classical algorithms have been utilized to drive programmable quantum devices into regimes suited for sensing tasks<sup>17–21</sup> and have been implemented on experimental architectures, e.g., atom arrays<sup>22</sup>. Numerical techniques including tensor networks<sup>21,23,24</sup>, and higher-order operations<sup>25</sup> have been used to represent quantum states, dynamics, and measurements in the design of optimal quantum sensing protocols; while learning and optimization techniques encompass see-saw methods<sup>26–28</sup>, semidefinite programming<sup>24,29</sup>, evolutionary algorithms<sup>20,30</sup>, global optimization<sup>21</sup>, gradient descent<sup>17,18</sup>, and reinforcement learning<sup>31</sup>. Further, classical machine learning methods, trained on real or simulated data from quantum devices, have also been proposed for performing the estimation step of sensing protocols in a Bayesian context<sup>32,33</sup>. Finally, adaptive quantum sensing

protocols—in which subsequent measurement settings are conditioned on previous outcomes—have also been studied in the context of learning and optimization methods, including with evolutionary algorithms<sup>30</sup>, tensor networks<sup>24</sup>, and reinforcement learning<sup>31</sup>.

Such approaches have helped identify the operational regimes where quantum sensing advantage may be found, but challenges remain in translating such insights into useful, near-term technologies. Noise sources, fabrication tolerances, constraints on quantum control (e.g., qubit connectivity, native operations), finite sampling rates, as well as unequal interaction of  $\phi$  across all parts of the probe state all arise in practical realizations of quantum sensing. Thus far, studies have largely focused on the impact of such effects on one or more protocol stages at a time by, e.g., quantifying the decrease in state preparation fidelity or quantum Fisher Information in the presence of noise or imperfect operations<sup>34</sup>. In practice, deviations from ideal behavior affect all protocol stages, from preparation through to estimation, impacting downstream stages of the protocol in non-trivial ways due to the tight interdependence of operational parameters. Importantly, if the impacts of experimental constraints and imperfections are not quantified in terms of final estimator performance, effective mitigation strategies that seek to maintain quantum advantage cannot be developed. This makes “end-to-end” approaches promising for the design and study of practical quantum sensing systems, i.e., numerical and theoretical frameworks that support simulation and optimization across all protocol stages, with direct and quantifiable impacts on estimator performance.

Here, we introduce a numerical workflow that unifies preparation, interaction, measurement, and estimator parameterization under one end-to-end variational framework (Fig. 1b). First, we examine how quantum

<sup>1</sup>University of Waterloo, Department of Physics & Astronomy, 200 University Ave., Waterloo, ON, Canada. <sup>2</sup>Institute for Quantum Computing, 200 University Ave., Waterloo, ON, Canada. <sup>3</sup>Perimeter Institute for Theoretical Physics, 31 Caroline St N., Waterloo, ON, Canada. <sup>4</sup>Ki3 Photonics Technologies, 2547 Rue Sicard, Montreal, QC, Canada. <sup>5</sup>University of Ottawa, Department of Physics, 75 Laurier Ave E, Ottawa, ON, Canada. ✉e-mail: [benjamin.maclellan@uwaterloo.ca](mailto:benjamin.maclellan@uwaterloo.ca)



**Fig. 1 | Building a quantum sensor.** **a** The value of an unknown parameter  $\phi$  is estimated from measurements on a probe state. Classical limits on sensing precision can be surpassed in certain operational regimes, pending the careful design and realization of all protocol steps. **b** Quantum sensing protocols lend themselves well to being framed as variational optimization problems, with end-to-end parameterized representations enabling powerful design and analysis tools. In this framing, the probe state is prepared via a parameterized unitary operation  $U(\theta)$ ; it interacts with the parameter of interest  $\phi$  through a quantum channel,  $K(\phi)$ ; and is read-out to classical information by a parameterized measurement  $M(\mu)$ . A classical neural network,  $E(\lambda)$ , then estimates the parameter value,  $\bar{\phi}$ , from a finite set of measurement data.

sensor dynamics can be modeled as parameterized quantum circuits. By computing and optimizing the Fisher Information as a function of circuit parameters, we converge the sensor towards regimes that support maximal achievable estimation precision. Next, we present the training of a classical neural network as a variational estimator of the unknown parameter  $\phi$  and, using a Bayesian approach to update the estimator as more data is observed, efficiently utilize the finite data sampling rates of quantum devices. We demonstrate the full end-to-end variational quantum sensing (VQS) protocol on probe state ansätze relevant to trapped-ion and photonic platforms, and benchmark the Fisher Information, estimator bias, and estimator variance of the learned protocols. Finally, we benchmark the scaling performance of the full framework and compare the effects of noisy probe state preparation for the GHZ and VQS protocols.

We find that the VQS approach enables the identification of optimal probe state and measurement regimes in the presence of realistic experimental constraints, as well as the direct study of how protocol imperfections and finite data effects impact estimation. These functionalities empower the development of strategies for realizing and maintaining quantum sensing advantages in practical, near-term devices.

## Results

### Variational optimization of quantum sensor dynamics

Parameterized quantum circuits are a useful model for studying and designing quantum systems in a unified mathematical framework. Applications of such circuits include variational quantum algorithms and learning ground states of many-body systems, among others<sup>15,35,36</sup>. For a quantum sensing task, the circuit representation of the sensor is divided into three parts. First, a parameterized unitary operation,  $U(\theta)$ , prepares a metrologically useful probe state. This probe state interacts dynamically with the parameter of interest,  $\phi$ , through a quantum channel,  $K(\phi)$ . Throughout this work, we consider the task of single-parameter phase estimation, such that  $K(\phi)$  is a local unitary rotation and  $\phi$  is a single scalar value. However, this method can be generalized to multiparameter estimation, where  $\phi$  is a real-valued vector, or to sensing tasks in which  $K$  represents an alternative dynamical process, e.g., a non-unitary channel. After the probe state interacts with and is perturbed by the parameter  $\phi$ , the probe state is measured with a parameterized positive operator-valued measurement (POVM),  $M(\mu) = \{\Pi_i(\mu) | \sum_i \Pi_i(\mu) = 1\}$ . Here, we consider the elements of the POVM to be projective measurements in locally-rotated bases, which is realized

via a layer of local rotation gates parameterized by  $\mu$ , followed by measuring in the computational basis. Such measurements are, in general, easily accessible and can be performed with high fidelity in many qubit architectures. This POVM maps the density operator to a likelihood probability distribution over measurement outcomes,

$$p(s_i|\phi) = \text{Tr} [\rho(\theta, \phi) \Pi_i(\mu)], \quad (1)$$

where  $s_i$  is the  $n$ -bit measurement outcome associated with the projector,  $\Pi_i$ . For a sensing protocol to realize the maximum precision that is possible, the state preparation and detection stages must be designed carefully around how the parameter of interest,  $\phi$ , interacts with the probe state.

In the variational quantum sensing framework, the device parameters,  $\theta$  and  $\mu$ , are tuned towards a regime where the probe state and its associated measurement can enable high-precision estimates of the value of  $\phi$ . Central to such a method is the choice of the objective function. For sensing tasks, the Fisher Information is a ubiquitous metric—often used for quantifying the performance of a sensing apparatus by providing a lower bound on the achievable variance of any unbiased estimator. The Fisher Information quantifies the amount of information that a probability distribution or quantum state carries about a random variable which parameterizes it. These quantities are referred to as the classical (CFI) or quantum (QFI) Fisher Information, respectively, and both are closely related to the susceptibility of the distribution or state to infinitesimal perturbations,  $\delta\phi$ .

The CFI of a discrete probability distribution  $p(s_i|\phi)$  that is parameterized by  $\phi$  and describes the probability of measurement outcomes,  $\{s_i\}$ , may be expressed as,

$$\mathcal{I}_\phi^{(C)} = \sum_i \frac{(\partial_\phi p(s_i|\phi))^2}{p(s_i|\phi)}. \quad (2)$$

Here,  $\partial_\phi$  denotes the first-order partial derivative with respect to the parameter of interest, i.e.,  $\partial/\partial\phi$ . The CFI can be viewed as the curvature (i.e., second-order derivative) of the Kullback–Leibler divergence between  $p(\phi)$  and  $p(\phi + \delta\phi)$ .

The QFI for an arbitrary,  $n$ -qubit quantum state, may be expressed as<sup>37</sup>,

$$\mathcal{I}_\phi^{(Q)} = \sum_{i, i' = 1}^{2^n} \frac{2 \text{Re}(\langle \lambda_i | \partial_\phi \rho | \lambda_{i'} \rangle)}{\lambda_i + \lambda_{i'}}, \quad (3)$$

$\lambda_i + \lambda_{i'} \neq 0$

where  $\{\lambda_i\}$  are the eigenvalues of the density operator. The QFI is an important quantity in quantum information, with close relationships to the Bures metric<sup>38</sup> and the quantum geometric tensor<sup>39</sup>, as well as a useful metric for detecting entanglement in quantum systems<sup>40–42</sup>. The QFI is equivalent to the optimal classical Fisher Information of a quantum state over all valid POVMs. As such, the probe state and measurement procedure must be tailored to each other and to the form of interaction with  $\phi$ .

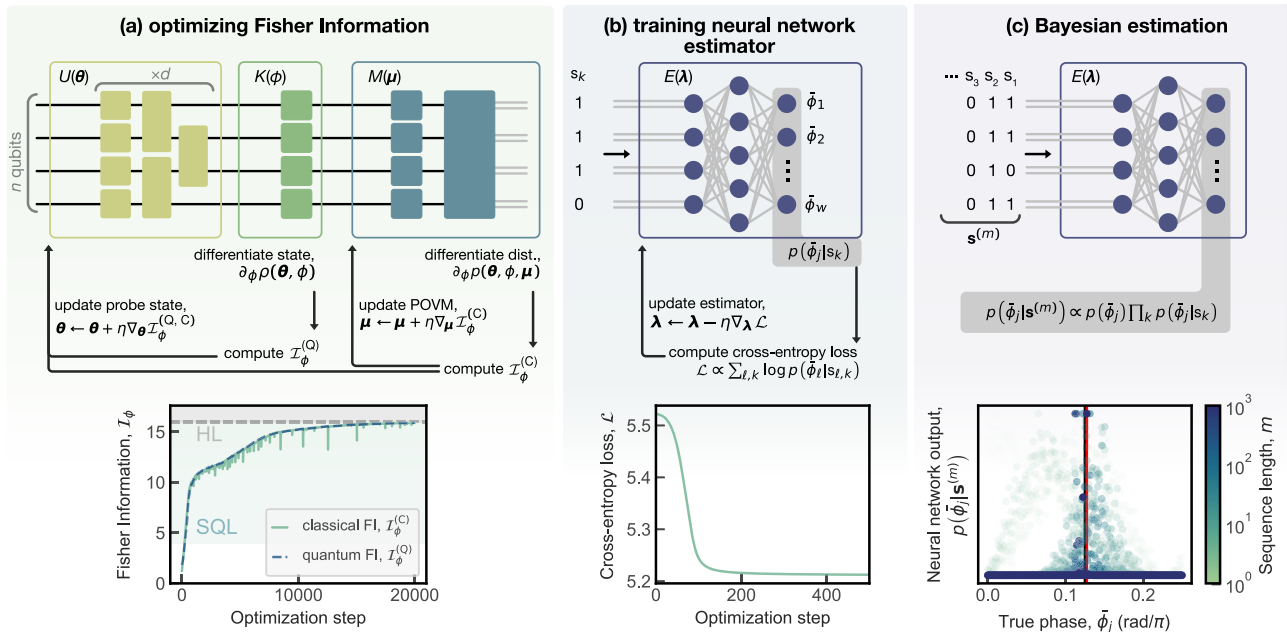
For closed systems, where the probe state is pure and  $K(\phi)$  is unitary, the QFI simplifies to<sup>37</sup>,

$$\mathcal{I}_\phi^{(Q)} = 4 \text{Re}(\langle \partial_\phi \rho | \partial_\phi \rho \rangle - |\langle \partial_\phi \rho | \rho \rangle|^2). \quad (4)$$

In the context of quantum sensing, the primary importance of the Fisher information is that it lowers the bounds of the variance of any unbiased estimator of the parameter  $\phi$ ,

$$\Delta^2 \bar{\phi} \geq \frac{1}{\mathcal{I}_\phi}. \quad (5)$$

This inequality is the famous Cramér–Rao bound, and holds in both the quantum and classical contexts<sup>43,44</sup>.



**Fig. 2 | End-to-end optimization and training of variational quantum sensing protocols.** **a**, top The quantum sensor dynamics are modeled as a parameterized circuit, where the probe state preparation and detection parameters ( $\theta$  and  $\mu$ , respectively) are tuned to optimize the quantum and classical Fisher Information with respect to  $\phi$ . (**a**, bottom) An example optimization curve using the CFI as the objective function and tracking the QFI throughout, which both converge to the HL. (**b**, top) A dense, feed-forward neural network is used as a variational ansatz for the estimator. The output layer of the network is interpreted as the posterior probability distribution of the phase value,  $\bar{\phi}$ , conditioned on measurement outcome  $s_i$ .<sup>33</sup> The

network is trained on measurement data from the quantum device, sampled across  $\phi$ , in a (**b**, bottom) supervised classification mode using the cross-entropy loss function. (**c**, top) Using a Bayesian update rule, the network is used to estimate  $\bar{\phi}$  from a sequence of measurements. This enables efficient training on individual samples and estimation on sequences of arbitrary length,  $m$ . (**c**, bottom) As more data were observed, the neural network posterior distribution converges around the true phase value,  $\phi$  (black vertical line), reducing the bias of the estimated phase value  $\bar{\phi}$  (black vertical line).

Using only classical resources (no entanglement, squeezing, etc.), the Fisher Information can scale, at most, linearly in the size of the probe state—e.g., for an  $n$ -qubit probe state  $\mathcal{I}_\phi \sim n$ . This scaling is known as the standard quantum limit (SQL) or shot-noise limit. However, if the probe state uses quantum resources such as entanglement, then the Fisher Information can scale with a quadratic improvement over the classical case,  $\mathcal{I}_\phi \sim n^2$ , which is known as the Heisenberg limit (HL).

As the first stage of the VQS protocol, we evaluate and optimize the Fisher Information of the parameterized quantum circuit (Fig. 2a). We use differentiable programming of the circuits to easily compute the quantum and classical Fisher Information with respect to  $\phi$ . First, a forward simulation pass of the circuit is performed to compute the perturbed quantum probe state,  $\rho(\phi)$ , and likelihood distribution  $p(s_i|\phi)$ . Using automatic differentiation, the derivative of the state and distribution with respect to the parameter of interest is computed,  $\partial_\phi \rho$  and  $\partial_\phi p(s_i|\phi)$ , respectively, and used to compute the QFI and/or CFI. As differentiable programming allows the arbitrary stacking of derivative passes, the gradient of the FI with respect to  $\theta$  and  $\mu$  can be computed, i.e.,  $\nabla_\theta \mathcal{I}_\phi^{(Q)}$ ,  $\nabla_\theta \mathcal{I}_\phi^{(C)}$ , and/or  $\nabla_\mu \mathcal{I}_\phi^{(C)}$ . By updating the values of the probe state and detection parameters via a gradient ascent optimization algorithm, the circuit can be tuned into a regime which has larger Fisher Information and, thus, maximum achievable estimation precision.

Once the circuit optimization has reached an appropriate stopping criterion (e.g., maximum number of iterations or plateau of the objective function), the parameters  $\theta$  and  $\mu$  are fixed. A dataset of measurement outcomes—classical bitstrings,  $s_i = \{0, 1\}^n$ —are collected for training and benchmarking the estimator. A sequence of measurement outcomes of length  $|D_\phi|$  are sampled at  $n_\phi$  distinct phase values across the range of interest,  $\phi \in [\phi_{\min}, \phi_{\max}]$ . Thus, the total number of bitstrings in the training dataset is  $|D| = n_\phi \times |D_\phi|$ . This dataset can be viewed as a three-dimensional binary array of shape  $(n_\phi, |D_\phi|, n)$ .

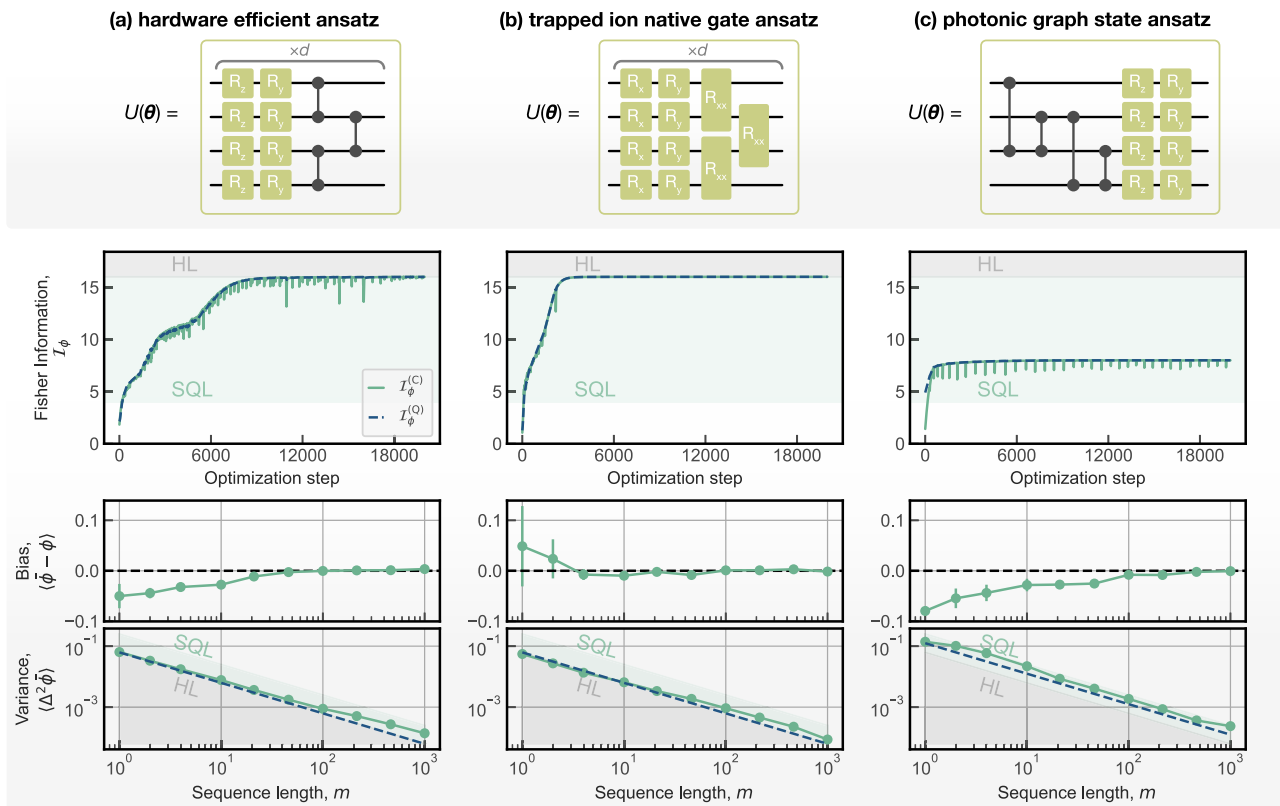
### Bayesian estimation via neural networks

With the quantum dynamics of the sensor optimized and datasets of measurement outcomes collected, we train a neural network ansatz as a variational estimation function. Here, we use the approach proposed by Nolan et al.<sup>33</sup>, in which a dense, feed-forward neural network,  $E(\lambda)$ , outputs an estimate of the phase value,  $\bar{\phi}$ , given a sequence of measurement shots,  $\mathbf{s}^{(m)}$ . The network is composed of  $n$  input nodes (corresponding to the projective measurement outcome on the  $n$ th qubit) and  $w$  output nodes, and is parameterized by weights and biases,  $\lambda$ . The output layer of the network is normalized to be a valid probability distribution (via the softmax function) and interpreted as the posterior distribution (via the softmax function) and interpreted as the posterior distribution,  $p(\bar{\phi}_j | s_i)$ , of estimating the phase value to be  $\bar{\phi}_j$  given the single measurement outcome,  $s_i$ . We note that in contrast to previous works using neural network estimators<sup>33</sup>, we use the raw bitstrings sampled from the quantum device as the input to the network. This format allows the network to more easily capture the relevant correlations between qubits. We find that this leads to improved convergence and estimator performance. As the network learns the correlations between  $\phi$  and the measurement outcomes directly from data, it can better compensate for noisy probe state preparation in the estimation, as we explore later.

The network is trained on single-shot measurements (encoded as one-hot vectors) via the cross-entropy loss function,

$$\mathcal{L} = \frac{1}{n_\phi |D_\phi|} \sum_{\ell=1}^{n_\phi} \sum_{k=1}^{|D_\phi|} \log p(\bar{\phi}_\ell | s_{k,\ell}). \quad (6)$$

Training the neural network can be viewed as indirectly performing Bayesian inversion, learning an approximate posterior distribution from sampled data. Once the network is trained, estimation is performed on sequences of arbitrary length shots via a Bayesian update rule. For a sequence of measurement outcomes of length  $m$ ,  $\mathbf{s}^{(m)} = \{s_1, s_2, \dots, s_m\}$ , each



**Fig. 3 | VQS with hardware-relevant ansätze.** We demonstrate the end-to-end framework with three unitary probe state preparation ansätze; including **a** a hardware-efficient ansatz (HEA) with controlled-phase interactions<sup>35</sup>, **b** trapped-ion native gate interactions (TIA)<sup>59</sup>, and **c** with states locally equivalent to a bundled graph state. The Fisher Information of the HEA and TIA plateau to the HL for an  $n$ -qubit system, the optimal value for entangled probe states, while the bundled graph

state plateaus to a FI of  $n^2/2$ , the theoretically known QFI value for such states (with  $n = 4$ ). The trained neural network estimators all converge towards being unbiased with increasing sequence length  $m$ , and partially saturate the CRB for the variance. For all results presented, the number of qubits is  $n = 4$  and the HEA and TIA circuit depth is  $d = 4$ .

individual shot is fed into the neural network, and the normalized product of the output layers is interpreted as the posterior distribution,

$$p(\bar{\phi}_j | s^{(m)}) \propto \prod_{k=1}^m p(\bar{\phi}_j | s_k), \quad (7)$$

where we assume conditional independence of the samples and a uniform prior. The estimated value is,

$$\bar{\phi} = \underset{\bar{\phi}_j}{\operatorname{argmax}} \left[ p(\bar{\phi}_j | s^{(m)}) \right], \quad (8)$$

and the bias and variance of this estimated value are computed as  $\langle \bar{\phi} - \phi \rangle$  and

$$\Delta^2 \bar{\phi} = \sum_{j=1}^w p(\bar{\phi}_j | s^{(m)}) (\bar{\phi} - \phi)^2, \quad (9)$$

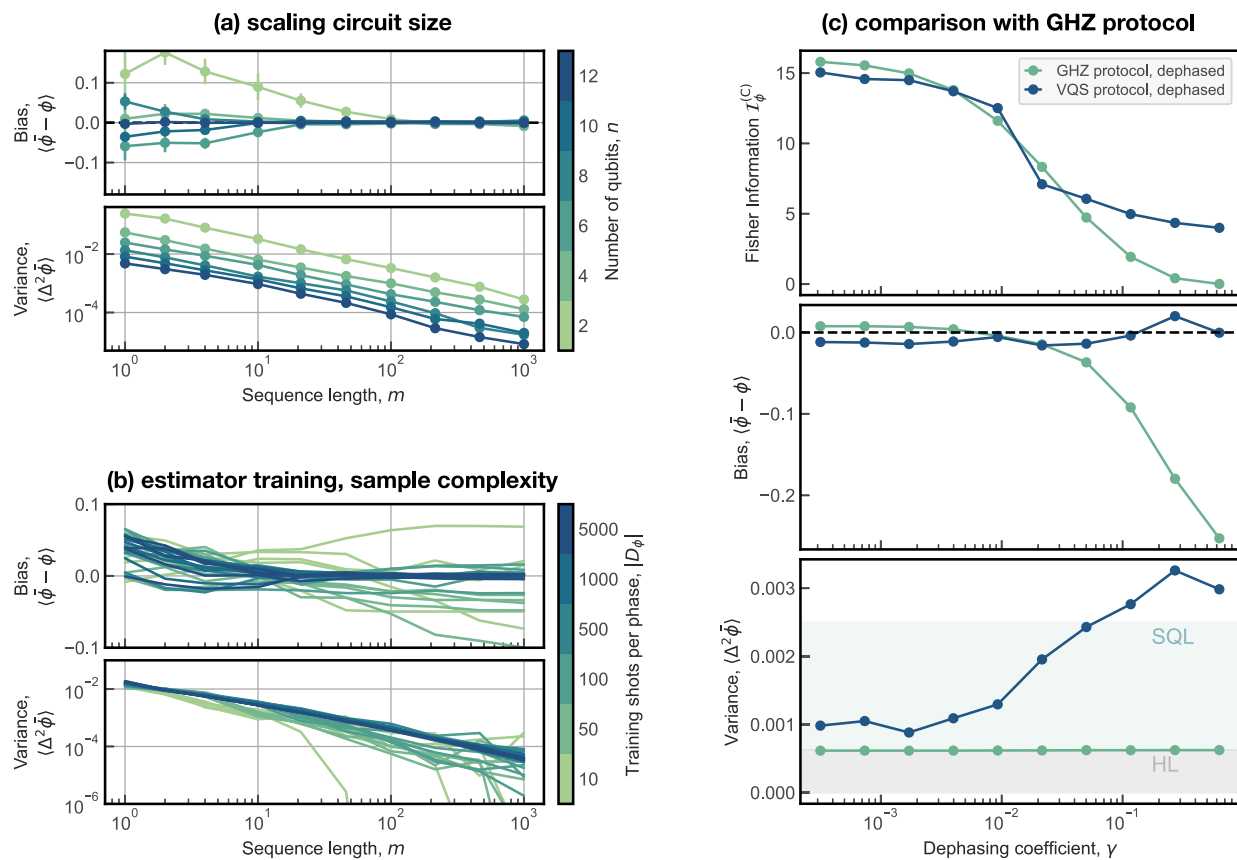
respectively<sup>33</sup>.

In the design and benchmarking of metrology schemes, it is desired for estimators to be both unbiased and have minimum variance; here, the Cramér–Rao bound provides an asymptotic lower bound of the variance for unbiased estimators. For an estimation procedures with repeated measurements, the Cramér–Rao bound is expressed as,  $\Delta^2 \bar{\phi} \geq (m \mathcal{I}_{\phi})^{-1}$ , such that  $(mn)^{-1}$  and  $(mn^2)^{-1}$  are the SQL and HL on estimating  $\phi$  from  $m$  measurements of an  $n$ -qubit system, respectively. Statements about estimator performance are often framed asymptotically—that is, in the assumption of an infinite number of measurement samples—an

assumption which often does not hold in real, experimentally-relevant contexts. Recent work has made significant progress in theoretical frameworks in limited-data regimes<sup>8,9</sup>; here, we explicitly benchmark VQS protocols in the regime of finite data using synthetic data sampled from differentiable tensor-network simulations of the quantum dynamics.

### Hardware-relevant ansätze

First, we demonstrate the VQS framework end-to-end with three ansätze that are relevant to current and near-term qubit architectures (Fig. 3). First, we consider a hardware-efficient ansatz (HEA)<sup>35</sup>, composed of local rotations interleaved with nearest-neighbor controlled-phase interactions (see Fig. 3a). Second, we consider a trapped-ion ansatz (TIA) composed of the native gates for ion-based quantum devices (Fig. 3b). The probe state is prepared by a unitary composed of interleaving local rotations around the X- and Y-axes, with Mølmer–Sørensen interactions which generate entanglement between adjacent qubits in the chain. In both the HEA and TIA cases, the “brickwork” circuit layer is repeated  $d$  times, and has open boundary conditions (i.e., the first and  $n$ th qubits do not interact). The final probe state ansatz we consider is a class of graph states amenable to quantum sensing<sup>45</sup>. Graph states are a class of entangled quantum states especially relevant for photonic quantum technologies<sup>46</sup>, as well as other architectures<sup>47</sup>. Certain families of graph states, including cluster and star states, have low QFI; however, Shettell et al. show that so-called bundled graph states have QFI which scales with the HL multiplied by a constant factor,  $1/k$ <sup>45</sup>. In Fig. 3c, we demonstrate the VQS on a probe state that is locally equivalent to a  $k = 2$  bundled graph state (corresponding to a complete bipartite graph). Figure 3 demonstrates the



**Fig. 4 | Scaling analysis of the VQS method.** **a** The VQS neural network estimator bias and variance for  $n$ -qubit probe state ansätze. **b** The required number of training shots per phase,  $|D_\phi|$ , to robustly train the neural network estimator. **c** Comparison between GHZ and VQS protocols with dephasing noise in the state preparation, including the classical Fisher Information of the quantum device and the neural network estimator bias and variance. The GHZ protocol estimator becomes biased as noise in the probe state preparation increases, though the variance remains

unchanged—due to the analytical form of the maximum likelihood estimator (see Methods). On the other hand, the VQS method can tune the device parameters toward probe states that are more robust to the particular form of noise. The VQS neural network estimator remains unbiased as the noise increases, as the network can learn the effect of noise from the training data, and instead trades-off for increased variance. For the results presented in (b, c), the number of qubits and circuit layers is  $n = d = 4$ .

quantum and classical Fisher Information of the sensing circuits during optimization, as well as the performance (i.e., bias and variance) of the neural network estimator after training.

In the Fisher Information optimization curves (Fig. 3, second row), we use the CFI as the objective function. We observe that the CFI is a more ‘rugged’ objective landscape compared to the QFI—the spurious decreases during optimization in the CFI are a result of gradient ascent update directions for  $\mu$ , which cause the projective measurement bases to non-optimally capture the large QFI of the parameterized probe state into classical information. Careful selection of the hyperparameters, e.g., the learning rate, is necessary for realizing optimization trajectories that do not get stuck in sub-optimal parameter regimes and converge in an efficient number of steps. We note that using the Fisher Information as an objective function demonstrates barren plateaus, a common and well-known challenge of parameterized quantum circuits and variational quantum algorithms.

The performance of the neural network estimator for all probe state ansätze (Fig. 3, bottom rows) tends towards being unbiased for longer sequence lengths. The variance of the trained VQS protocol reaches below the SQL in the finite data regime and partially saturates the HL. However, at longer sequence lengths tested ( $m \approx 1000$ ), the variance begins to increase and no longer reaches the HL. This is due to the combination of discretizing the posterior distribution and the Bayesian update, which leads to spurious noise in  $p(\phi_j | s^{(m)})$  that is amplified as  $m$  increases.

### Scaling of VQS protocols

After demonstrating the end-to-end variational optimization of the quantum sensing protocols, we turn our attention to exploring the scaling performance of the framework. In Fig. 4a we show the scaling in the number of qubits, from  $n = 2$  to  $n = 10$ , demonstrating for each that the method can realize estimation precision approaching the HL in noise-free environments. We also observe that as  $n$  increases, fewer measurement samples are required for the estimator to plateau towards being unbiased, as each individual measurement sample contains more information and correlations between qubits that are useful for the estimation. The exponential scaling of the Hilbert space limits the size of quantum sensors that can be simulated classically; however, variational and optimal quantum control algorithms for NISQ devices have been proposed and demonstrated for driving system dynamics towards states with large Fisher Information, and advances in this domain can be readily adopted for sensing protocols. Due to noisy circuits and limited shot rates, the capability to sample high-quality training datasets may be challenging for some current qubit architectures. However, as quantum sensing devices improve and scale, we expect such neural network approaches will be vital for sensing protocols. Such models, in general, provide an important tool to counter bias from limited training sets, especially when the state space dimension is much larger than the size of the available dataset.

In Fig. 4b, we characterize the number of measurement samples per phase needed to robustly train the neural network estimator. As



outlined in the section “Variational optimization of quantum sensor dynamics”, the total number of measurement samples in the training dataset is  $|\mathcal{D}| = |\mathcal{D}_\phi| \times n_\phi$ . For small training dataset sizes (light green lines), the estimator cannot be robustly trained, leading to the bias diverging away from zero and the variance collapsing to zero (due to the discretized posterior distribution collapsing towards a Kronecker delta function). We observe that training a modestly-sized network on datasets of size approximately  $|\mathcal{D}_\phi| \approx 1000$  and  $n_\phi \approx 100$  resulted in robust and high-performance estimation. However, this requirement is highly dependent on the number of qubits, sensing dynamics, and quality of the training data.

### Comparison to GHZ protocol

The GHZ protocol is a well-studied, paradigmatic metrology protocol that is optimal in noise-free environments. Here, we compare the effects of noisy probe state preparation on both a GHZ and VQS protocol. We use analogous noise models affecting the state preparation for the two protocols: following every two-qubit gate in the state preparation circuit, the control and target qubits are acted on by independent dephasing channels, each with coefficient  $\gamma \in [0, 1]$  (see Methods, section “Comparison with GHZ protocol”). For the GHZ state preparation, the circuit contains  $n - 1$  two-qubit gates, while the VQS using the HEA consists of  $(n - 1)d$  two-qubit interactions. For each value of  $\gamma$ , we optimize the circuit in the face of the noisy state preparation and sample simulated training and testing datasets from the optimized device. We sample a testing dataset of the GHZ and VQS circuits of the same size and values of  $\phi$ . For the estimation based on GHZ data, we use a maximum likelihood estimator (where the likelihood is under the assumption of noise-free evolution, see Methods); for the VQS, we use a trainable neural network ansatz as before. We then directly benchmark and compare the bias and variance of the two sensing protocols. In Fig. 4, we see that for low noise strengths, the GHZ and VQS have comparable Fisher Information, saturating to the optimal value of  $n^2$  (fixed to  $n = 4$ ). As the noise increases, the CFI of the GHZ probe decreases rapidly to 0 at large  $\gamma$ , such that the measurement outcomes are completely random and carry no information about the value of  $\phi$ . The VQS can, however, mitigate the effects of the noise by tuning  $\theta$  and  $\mu$  to remain above the SQL on the Fisher Information for larger amounts of environmental noise; in-line with recent demonstrations that parameterized circuits are able to mitigate the entropic effects of the environmental noise in shallow circuits<sup>48</sup>. Analyzing the final estimator performance (Fig. 4c, middle and bottom), the GHZ maximum likelihood estimator quickly becomes biased as the dephasing noise increases. The neural network estimator, on the other hand, compensates for the noise in its estimation as it learns the effect of the noise channels from the sampled data. As such, the neural network estimator remains unbiased as the noise increases, and instead trades-off for increased variance.

### Discussion

In this work, we introduced a variational framework for quantum sensing protocols, optimizing both the quantum evolution and the classical estimator. We model the protocol stages corresponding to probe state preparation, interaction, and measurement with parameterized quantum circuits, optimizing the Fisher Information towards regimes with maximal achievable precision. Subsequently, we use a neural network as a variational estimator ansatz, which is trained on single labeled measurement outcomes and used in a Bayesian context for performing estimation on arbitrarily long sequences of measurements.

The Fisher Information is a natural choice for the objective function, as it is a straightforward and interpretable way of quantifying the utility of a probe state for metrology. However, other important protocol considerations are not encapsulated by the Fisher Information, including, for example, the dynamic range of the sensor—i.e., the effective range over which the unknown parameter can be estimated. As such, continued investigations of new objective functions<sup>22</sup> and hybrid

control algorithms<sup>14</sup> may suggest novel approaches toward designing more useful and robust sensing devices. The end-to-end parameterization and differentiable programming of VQS protocols open promising avenues for using higher-order derivative quantities and probability distribution moments for improved objective functions<sup>49</sup>. In addition, the variational nature of such methods may enable new approaches for sensing based on adaptive feedback and active learning concepts<sup>31,50</sup>.

In parallel, as both variational algorithms and machine learning capabilities advance, with model architectures and compute infrastructures rapidly evolving, finding synergistic overlap with challenges in quantum sensing may prove to be a fruitful direction for development. In particular, future work should explore alternative neural network architectures and training methods towards variational quantum sensing protocols. Architectures that are implicitly designed for sequential data (e.g., recurrent neural networks, transformers<sup>51,52</sup>), as well as approaches developed for approximating complex, high-dimensional probability distributions (e.g., normalizing flows<sup>53</sup>) are likely well-suited towards quantum sensing protocols.

The insights provided by VQS frameworks can be used directly to identify improved operational settings for realistic, constrained qubit architectures, especially relevant for near-term programmable quantum devices<sup>22,54</sup>. Further, future work should investigate how such frameworks can be utilized for the automatic discovery of actionable and human-interpretable<sup>55</sup> mitigation strategies to realize and preserve quantum advantage for the next generation of sensing technologies.

## Methods

### Circuit simulations and training

Quantum circuit simulations were performed with differentiable tensor-network methods, using the TensorCircuit and JAX frameworks<sup>56,57</sup>. All circuits evolve a pure  $n$ -qubit initial state  $|\rho_0\rangle = |0\rangle^{\otimes n}$  through the parameterized state preparation,  $U(\theta)$ . The state preparation is unitary, except for in cases when noise is explicitly considered in the analysis of Fig. 4c. The probe state then transforms under local rotations parameterized by the parameter of interest,  $\phi$ . The local quantum channel which imprints  $\phi$  onto the state is  $K(\phi) = \otimes_{i=1}^n R_z(\phi)$ . Finally the perturbed probe state is measured by parameterized projective measurements in rotated bases, implemented as a layer of local rotations followed by projecting each qubit into the computational basis (i.e., eigenvectors of the Z operator). Optimization of the circuit parameters,  $\theta$  and  $\mu$ , is performed using the ADAM algorithm and the Fisher Information as the objective function (unless otherwise noted, we use the CFI as the loss function, but the QFI can also be used). The HEA, TIA, and bundled graph state circuits (including both  $\theta$  and  $\mu$ ) have  $2nd + 4n$ ,  $(3n - 1)d + 3n$ , and  $4n$  total variational parameters, respectively.

### Neural network training

In this work, the network architecture used as the estimator ansätze are fully-connected, feed-forward neural networks. The networks have dimension  $(n, h_1, \dots, h_{n_{\text{layer}}}, w)$ , where  $n$  is the number of qubits,  $h_i$  is the number of nodes in the  $i$ th hidden layer,  $n_{\text{layer}}$  is the number of hidden layers, and  $w$  is the number of output nodes (i.e., the number of phase discretization points). Each neuron passes information through a nonlinear activation function, here the ReLU function, defined as  $\max(x, 0)$ . The output layer is normalized to a valid probability distribution using the softmax function. The corresponding phase label,  $\phi_j$ , for a measurement sample is encoded as a one-hot encoded vector, i.e., a Kronecker delta function,  $\delta_{ij}$ . We use the cross-entropy function between the label one-hot vector and the output layer as the loss function, with an added L2 regularization term to improve training and estimator performance. The ADAM optimizer is used for stochastic gradient descent. Neural network estimators were implemented and trained using the JAX framework.

## Comparison with GHZ protocol

Here, the probe state is the highly entangled GHZ state,  $|\text{GHZ}\rangle = \frac{1}{\sqrt{2}}(|00\dots 0\rangle + |11\dots 1\rangle)$ , which interacts with the parameter of interest through a local rotation around the Z-axis of the Bloch sphere,  $R_z(\phi)$ . The GHZ state is prepared via the quantum circuit,

$$|\text{GHZ}\rangle = \left( \prod_{i=2}^n \text{CNOT}_{1,i} \right) H_1 |0\rangle^{\otimes n}. \quad (10)$$

A local dephasing channel is applied on both the control and target qubits for every CNOT, described by the Kraus operators,

$$A_0 = \begin{bmatrix} 1 & 0 \\ 0 & \sqrt{1-\gamma} \end{bmatrix}, \quad A_1 = \begin{bmatrix} 0 & 0 \\ 0 & \sqrt{\gamma} \end{bmatrix},$$

and the state evolves as,  $\rho \rightarrow A_0 \rho A_0^\dagger + A_1 \rho A_1^\dagger$ . We study this prototypical noise model as two-qubit interactions are the dominant source of noise in practically all quantum information processing platforms<sup>58</sup>. A parity measurement of the state in the X basis is performed,  $P(s_i)$ , i.e., calculating the sum modulo-two of the sampled bitstring,  $s_p$ , such that,

$$P(s_i) = \begin{cases} 0 & \text{if } s_i \text{ is even parity} \\ 1 & \text{if } s_i \text{ is odd parity} \end{cases}. \quad (11)$$

The likelihood distribution of measuring an even parity measurement for these sensing dynamics, in the noise-free case, is,

$$\begin{aligned} p(P(s_i) = 0|\phi) &= \cos^2\left(\frac{n\phi}{2}\right) \\ p(P(s_i) = 1|\phi) &= \sin^2\left(\frac{n\phi}{2}\right). \end{aligned} \quad (12)$$

Using a maximum likelihood estimator, i.e., maximum a posteriori estimation with a uniform prior distribution, our estimator for a sequence of  $m$  measurement samples is,

$$\begin{aligned} p(\bar{\phi}|s^{(m)}) &\propto \prod_{i=k}^m p(\bar{\phi}|s_i) \\ &= \prod_{k=1}^m P(s_i) \cos^2\left(\frac{n\phi}{2}\right) + (1 - P(s_i)) \sin^2\left(\frac{n\phi}{2}\right), \end{aligned} \quad (13)$$

with a variance which scales as  $\Delta^2 \bar{\phi} \sim n^{-2}$  and, as it is derived using the noise-free likelihood, is not dependent on  $\gamma$ . To compare the VQS and GHZ protocols, both are defined as quantum circuits, and projective measurement samples are collected from both over a range of phase values  $\phi_k$ . Training data were collected for the neural network estimator, and equivalent testing datasets were sampled for both the GHZ and VQS protocols.

## Data availability

All data related to this work and its results are available from the corresponding author on reasonable request.

## Code availability

The code supporting this work is released under the Apache 2.0 license at <https://github.com/benjimaclellan/queso>.

Received: 22 April 2024; Accepted: 25 October 2024;

Published online: 19 November 2024

## References

- Giovannetti, V., Lloyd, S. & Maccone, L. Quantum-enhanced measurements: beating the standard quantum limit. *Science* **306**, 1330–1336 (2004).
- Schnabel, R., Mavalvala, N., McClelland, D. E. & Lam, P. K. Quantum metrology for gravitational wave astronomy. *Nat. Commun.* **1**, 121 (2010).
- Stray, B. et al. Quantum sensing for gravity cartography. *Nature* **602**, 590–594 (2022).
- Lopaeva, E. D. et al. Experimental realization of quantum illumination. *Phys. Rev. Lett.* **110**, 153603 (2013).
- Taylor, M. A. et al. Biological measurement beyond the quantum limit. *Nat. Photonics* **7**, 229–233 (2013).
- Crawford, S. E. et al. Quantum sensing for energy applications: review and perspective. *Adv. Quantum Technol.* **4**, 2100049 (2021).
- Cimini, V. et al. Quantum sensing for dynamical tracking of chemical processes. *Phys. Rev. A* **99**, 053817 (2019).
- Jesus, R. J. Non-asymptotic quantum metrology: extracting maximum information from limited data. Univ. Sussex. Thesis. <https://hdl.handle.net/10779/uos.23474864.v1> (2020).
- Meyer, J. J., Khatri, S., França, D. S., Eisert, J. & Faist, P. Quantum metrology in the finite-sample regime. Preprint at arXiv:2307.06370 (2023).
- Zhou, S., Zhang, M., Preskill, J. & Jiang, L. Achieving the Heisenberg limit in quantum metrology using quantum error correction. *Nat. Commun.* **9**, 78 (2018).
- Torlai, G. et al. Neural-network quantum state tomography. *Nat. Phys.* **14**, 447–450 (2018).
- Metz, F. & Bukov, M. Self-correcting quantum many-body control using reinforcement learning with tensor networks. *Nat. Mach. Intell.* **5**, 780–791 (2023).
- Krenn, M., Landgraf, J., Foesele, T. & Marquardt, F. Artificial intelligence and machine learning for quantum technologies. *Phys. Rev. A* **107**, 010101 (2023).
- Youssry, A. et al. Experimental graybox quantum system identification and control. *npj Quantum Inf.* **10**, 1–9 (2024).
- Cerezo, M. et al. Variational quantum algorithms. *Nat. Rev. Phys.* **3**, 625–644 (2021).
- Huang, J., Zhuang, M., Zhou, J., Shen, Y. & Lee, C. Quantum metrology assisted by machine learning. *Adv. Quantum Technol.* 2300329 (2024).
- Meyer, J. J., Borregaard, J. & Eisert, J. A variational toolbox for quantum multi-parameter estimation. *npj Quantum Inf.* **7**, 1–5 (2021).
- Beckey, J. L., Cerezo, M., Sone, A. & Coles, P. J. Variational quantum algorithm for estimating the quantum Fisher information. *Phys. Rev. Res.* **4**, 013083 (2022).
- Yang, J. et al. Variational principle for optimal quantum controls in quantum metrology. *Phys. Rev. Lett.* **128**, 160505 (2022).
- Koczor, B., Ālrint, Endo, S., Jones, T., Matsuzaki, Y. & Benjamin, S. C. Variational-state quantum metrology. *N. J. Phys.* **22**, 083038 (2020).
- Kaubriegger, R. et al. Variational spin-squeezing algorithms on programmable quantum sensors. *Phys. Rev. Lett.* **123**, 260505 (2019).
- Marciniak, C. D. et al. Optimal metrology with programmable quantum sensors. *Nature* **603**, 604–609 (2022).
- Chabuda, K., Dziarmaga, J., Osborne, T. J. & Demkowicz-Dobrzański, Rafał. Tensor-network approach for quantum metrology in many-body quantum systems. *Nat. Commun.* **11**, 250 (2020).
- Kurdzialek, S., Dulian, P., Majsak, J., Chakraborty, S. & Demkowicz-Dobrzański, R. Quantum metrology using quantum combs and tensor network formalism. Preprint at arXiv:2403.04854 (2024).
- Bavaresco, J., Lipka-Bartosik, P., Sekatski, P. & Mehboudi, M. Designing optimal protocols in Bayesian quantum parameter estimation with higher-order operations. *Phys. Rev. Res.* **6**, 023305 (2024).
- Macieszczak, K., Fraas, M. & Demkowicz-Dobrzański, Rafał. Bayesian quantum frequency estimation in presence of collective dephasing. *N. J. Phys.* **16**, 113002 (2014).
- Tóth, G. éza, Vértesi, Tamás, Horodecki, Paweł. & Horodecki, R. Activating hidden metrological usefulness. *Phys. Rev. Lett.* **125**, 020402 (2020).

28. Trényi, R. et al. Activation of metrologically useful genuine multipartite entanglement. *N. J. Phys.* **26**, 023034 (2024).
29. Tóth, G. éza & Vértesi, Tamás Quantum states with a positive partial transpose are useful for metrology. *Phys. Rev. Lett.* **120**, 020506 (2018).
30. Hentschel, A. & Sanders, B. C. Efficient algorithm for optimizing adaptive quantum metrology processes. *Phys. Rev. Lett.* **107**, 233601 (2011).
31. Belliardo, F., Zoratti, F., Marquardt, F. & Vittorio, G. Model-aware reinforcement learning for high-performance Bayesian experimental design in quantum metrology. Preprint at <https://arxiv.org/abs/2312.16985> (2024).
32. Greplova, E., Andersen, C. K. & Mølmer, K. Quantum parameter estimation with a neural network. Preprint at arXiv:1711.05238 (2017).
33. Nolan, S., Smerzi, A. & Pezzè, L. A machine learning approach to Bayesian parameter estimation. *npj Quantum Inf.* **7**, 1–8 (2021).
34. Datta, A. et al. Quantum metrology with imperfect states and detectors. *Phys. Rev. A* **83**, 063836 (2011).
35. Kandala, A. et al. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246 (2017).
36. Benedetti, M., Lloyd, E., Sack, S. & Fiorentini, M. Parameterized quantum circuits as machine learning models. *Quantum Sci. Technol.* **4**, 043001 (2019).
37. Liu, J., Yuan, H., Lu, X.-M. & Wang, X. Quantum Fisher information matrix and multiparameter estimation. *J. Phys. A Math. Theor.* **53**, 023001 (2019).
38. Šafránek, D. Discontinuities of the quantum Fisher information and the Bures metric. *Phys. Rev. A* **95**, 052320 (2017).
39. Stokes, J., Izaac, J., Killoran, N. & Carleo, G. Quantum natural gradient. *Quantum* **4**, 269 (2020).
40. Hauke, P., Heyl, M., Tagliacozzo, L. & Zoller, P. Measuring multipartite entanglement through dynamic susceptibilities. *Nat. Phys.* **12**, 778–782 (2016).
41. Hyllus, P. et al. Fisher information and multiparticle entanglement. *Phys. Rev. A* **85**, 022321 (2012).
42. Toth, G. Multipartite entanglement and high precision metrology. *Phys. Rev. A* **85**, 022322 (2012).
43. Helstrom, C. W. Minimum mean-squared error of estimates in quantum statistics. *Phys. Lett. A* **25**, 101–102 (1967).
44. Tóth, G. & Apellaniz, I. Quantum metrology from a quantum information science perspective. *J. Phys. A Math. Theor.* **47**, 424006 (2014).
45. Shettell, N. & Markham, D. Graph states as a resource for quantum metrology. *Phys. Rev. Lett.* **124**, 110502 (2020).
46. Thomas, P., Ruscio, L., Morin, O. & Rempe, G. Efficient generation of entangled multiphoton graph states from a single atom. *Nature* **608**, 677–681 (2022).
47. Lanyon, B. P. et al. Measurement-based quantum computation with trapped ions. *Phys. Rev. Lett.* **111**, 210501 (2013).
48. Duschene, M., Carrasquilla, J. & Laflamme, R. Characterization of overparametrization in the simulation of realistic quantum systems. *Phys. Rev. A* **109**, 062607 (2024).
49. MacLellan, B. et al. Inverse design of photonic systems. *Laser Photon. Rev.* **18**, 2300500 (2024).
50. Lange, H. et al. Adaptive quantum state tomography with active learning. *Quantum* **7**, 1129 (2023).
51. Vaswani, A. et al. Attention is all you need. In *Advances in Neural Information Processing Systems* (Curran Associates, Inc., 2017).
52. Melko, R. G. & Carrasquilla, J. Language models for quantum simulation. *Nat. Comput. Sci.* **4**, 11–18 (2024).
53. Rezende, D. & Mohamed, S. Variational inference with normalizing flows. *Int Conf Mach Learn* **37**, 1530–1538 (2015).
54. Madsen, L. S. et al. Quantum computational advantage with a programmable photonic processor. *Nature* **606**, 75–81 (2022).
55. Krenn, M., Kottmann, J. S., Tischler, N. & Aspuru-Guzik, A. Conceptual understanding through efficient automated design of quantum optical experiments. *Phys. Rev. X* **11**, 031044 (2021).
56. Zhang, S.-X. et al. TensorCircuit: a quantum software framework for the NISQ Era. *Quantum* **7**, 912 (2023).
57. Bradbury, J. et al. JAX: composable transformations of Python + NumPy programs (2018).
58. De Leon, N. P. et al. Materials challenges and opportunities for quantum computing hardware. *Science* **372**, eabb2823 (2021).
59. Czischek, S., Torlai, G., Ray, S., Islam, R. & Melko, R. G. Simulating a measurement-induced phase transition for trapped ion circuits. *Phys. Rev. A* **104**, 062405 (2021).

## Acknowledgements

We thank Matthew Duschene, Yi Hong Teoh, Sebastian Wetzel, Logan Cooke, Johannes Jakob Meyer, and Augusto Smerzi for critically helpful discussions. This work was supported by NSERC and the Perimeter Institute for Theoretical Physics. B.M. acknowledges support from NSERC through the Vanier CGS program. R.G.M. acknowledges support from NSERC through the Discovery Grant program. The calculations in this work were enabled in part by support provided by SHARCNET and Compute Canada. Research at Perimeter Institute is supported in part by the Government of Canada through the Department of Innovation, Science and Economic Development Canada and by the Province of Ontario through the Ministry of Economic Development, Job Creation and Trade.

## Author contributions

B.M. and P.R. conceived the idea. B.M. developed the circuit optimization and machine learning algorithms, and performed the numerical analysis. P.R., S.C., and R.G.M. provided supervision and guidance during the project. All authors contributed to discussions and the writing and editing of the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to Benjamin MacLellan.

**Reprints and permissions 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-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, 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 licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence 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 licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024