

Multidimensional Quantum Estimation and Model Learning Framework Based on Variational Bayesian Inference

Federico Belliardo^{1,2,*}, Erik M. Gauger¹, Mohamed H. Abobeih^{3,4}, Tim H. Taminiau^{3,4},
Yoann Altmann^{5,†} and Cristian Bonato^{1,‡}

¹*Institute of Photonics and Quantum Sciences, SUPA, School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, United Kingdom*

²*Pritzker School of Molecular Engineering, University of Chicago, Chicago, Illinois 60637, USA*

³*QuTech, Delft University of Technology, PO Box 5046, Delft 2600 GA, The Netherlands*

⁴*Kavli Institute of Nanoscience Delft, Delft University of Technology, PO Box 5046, Delft 2600 GA, The Netherlands*

⁵*Institute of Signals, Sensors and Systems, School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh EH14 4AS, United Kingdom*



(Received 9 August 2025; accepted 16 April 2026; published 15 June 2026)

The advancement and scaling of quantum technology has made the learning and identification of quantum systems and devices in highly multidimensional parameter spaces a pressing task for a variety of applications. In many cases, the integration of real-time feedback control and adaptive choice of measurement settings places strict demands on the speed of this task. Here we present a joint model selection and parameter estimation algorithm that is fast and operable on a large number of model parameters. The algorithm is based on variational Bayesian inference, which approximates the target posterior distribution by optimizing a tractable family of distributions, making it more scalable than exact inference methods relying on sampling and that generally suffer from high variance and computational cost in high-dimensional spaces. We show how a regularizing prior can be used to select between competing models, each comprising a different number of parameters, identifying the simplest model that explains the experimental data. The regularization can further separate the degrees of freedom, e.g., quantum systems in the environment or processes, which contribute to major features in the observed dynamics, with respect to others featuring small coupling, which only contribute to a background. As an application of the introduced framework, we consider the problem of the identification of multiple individual nuclear spins with a single electron spin quantum sensor, relevant for nanoscale nuclear magnetic resonance and for the implementation of multiqubit quantum networking nodes. With the number of environmental spins unknown a priori, our Bayesian approach is able to correctly identify the model, i.e., the number of spins and their couplings. We benchmark the algorithm on both simulated and experimental data, using standard figures of merit, and demonstrate that we can estimate dozens of parameters within minutes. Our methodology is compatible with the implementation of real-time adaptive choice of experimental settings in multiparameter estimation, which has the potential to greatly decrease the estimation time in a variety of metrology, tomography, and characterization tasks by prioritizing the measurements with the highest content of information.

DOI: [10.1103/6yhz-9vr6](https://doi.org/10.1103/6yhz-9vr6)

I. INTRODUCTION

The increasingly fast-paced progress in the scale and complexity of quantum devices developed for quantum

computing [1–5], networking, [6–10] and sensing [11–15] requires the development of new frameworks to characterize, measure, and “learn” quantum systems [16–22].

Historically, the characterization of quantum systems has broadly followed two approaches in quantum information theory: The first originates in quantum tomography [23–25] and is based on learning the full representation of quantum states and processes, and more recently also their dynamical generators [26–28], such as the full set of Hamiltonian and/or Lindbladian operators describing the system evolution. In this scenario, the only assumption is that the state to be identified belongs to a large set of physical states. The second approach, rooted in quantum

*Contact author: federico.belliardo@gmail.com

†Contact author: y.altmann@hw.ac.uk

‡Contact author: c.bonato@hw.ac.uk

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

parameter estimation, aims at estimating a typically small set of parameters on which the state depends [29,30], with the goal of evaluating the ultimate precision limit on the knowledge of said parameters in terms of the (quantum) Cramér-Rao bound. In this scenario, we assume that we have a parametrized expression for the state, or the channel, that we want to characterize, which usually belongs to a small and strongly constrained set of states. The rapidly expanding fields of quantum sensing and metrology are primarily based on quantum parameter estimation, and typically deal with the practical problem of determining the external parameters encoded in a quantum system used as a probe, such as a magnetic field, temperature, or optical phase. In some cases, the characterization can be extended to model learning scenarios [31,32], where we want to identify a description of the quantum system among multiple different classes of models, each comprising a different set of free parameters. In this manuscript, we take an integrated approach that bases the characterization of quantum systems and the optimal selection of the model on parameter estimation. Although parameter estimation scenarios typically involve fewer parameters than corresponding tomography scenarios, the number of parameters could be sufficiently large that the task of model learning is in the regime of high-dimensional parameter estimation, even for relatively simple quantum systems.

The successful characterization of a quantum system should include a description that is sufficiently accurate for the purposes of the experimenter [17]. We require this description to include statistically significant confidence levels and a clear delimitation of the domain of validity of the performed model selection routine (i.e., clarity in the set of possible models that have been considered for the learning). The gold standard for satisfying both requirements is to produce a Bayesian distribution over the set of models and parameters considered [33].

A crucial benefit of Bayesian inference is that, from the probability density or mass function, one can compute uncertainties for every function of the model parameters and optimize experimental settings for subsequent measurements, for example, to minimize the overall sensing time. This is very important in quantum applications, which are often limited by slow data acquisition because the signals produced by individual quantum systems—such as streams of single photons or electrons—are generally weak.

Bayesian inference has been widely applied in the context of quantum technology: Bayesian quantum tomography [34–38] has been shown to be accurate [23] and to provide reliable and tight error bounds [39]. Although rarely analytically tractable, numerical approaches have been deployed to make tomography experimentally practical [40]. Its efficiency has been further improved with adaptive selection of experimental settings [41,42] and adaptive discovery of priors [33,43]. Bayesian techniques


can also be used to learn the generators of dynamics, such as the Hamiltonian [31,44–48] or the Lindblad master equation [32,49]. Applications of these concepts in the context of quantum sensing have demonstrated the adaptive choice of optimal experimental settings to minimize sensing time [50–61] and qubit control within a noisy environment [46,48,62–64]. In all applications involving experiment adaptation and feedback, it is crucial to minimize the processing time required to update the probability distribution according to Bayes rule, and generate the optimal control sequence for the next measurements. In quantum sensing, the processing time is ideally negligible compared to the data acquisition time, or at least the shortening of sensing time given by real-time adaptation needs to be larger than the computational time invested to achieve this.

Despite its success for small-scale quantum technologies, the extension of Bayesian approaches to learn highly multidimensional quantum systems is highly nontrivial, since exact Bayesian inference is often computationally intractable. As the parameter space grows, evaluating and sampling from the posterior becomes exponentially more expensive, in what is known as the “curse of dimensionality” [65,66]. This has restricted the application of Bayesian inference to experiments in quantum tomography [37,41,67,68] and quantum sensing [50,52,53,56,58,61,69,70] involving only a handful of parameters [71].


Here we introduce a powerful and practical solution to this challenge, deploying an approximate approach that extends the applicability of Bayesian inference to high-dimensional learning problems in quantum settings. To avoid the intractable scaling of exact Bayesian techniques for large numbers of parameters, we deploy variational Bayesian inference (VBI) [72], introducing a family of simpler distributions and selecting the one that best approximates the true posterior, see Fig. 1. In Sec. II, we show that, as the dimensionality of the estimation problem grows, we can choose distributions whose sampling time and memory complexity scale polynomially with the number of parameters, enabling us to efficiently estimate several tens of parameters.

In addition to enabling the estimation of large numbers of parameters, our approach is also applicable to cases where the system model is not uniquely defined, for example, for quantum systems coupling to an environment with a number of constituents unknown a priori (e.g., a spin bath). In Sec. III, we show that, for cases where the possible model classes are organized in a hierarchical structure with models higher in the hierarchy featuring higher complexity, our framework successfully identifies the model with minimal complexity that explains the data. This is, for example, the case of Hamiltonian or Lindbladian identification where the most complex models contain a large number of terms with their respective couplings and parameters, suggesting the need for a form

(a) Sketch of the VBI parameter learning algorithm



(b) Training of the ansatz



(c) Example of learning of a single parameter multi-modal probability density.




FIG. 1. Introduction to high-dimensional parameter estimation for quantum experiments using variational Bayesian inference. (a) Schematic representation of a quantum estimation experiment based on variational Bayesian inference. The physical system, here represented by an ensemble of spins, undergoes a series of measurements, controlled by the parameters $\mathbf{x}$, obtaining the outcomes $\mathbf{y}$. This information is used to train the posterior ansatz in variational Bayesian inference, i.e., a function that transforms $\boldsymbol{\theta} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ in samples extracted from the posterior distribution for $\boldsymbol{\theta}$, that contains the information on the batch of measurement just executed. The Bayesian posterior can be used in a real-time feedback algorithm to calculate the controls for the next batch of measurements, adaptively. (b) The space of distributions spanned by varying $\boldsymbol{\lambda}$ is indicated by the green set, and it corresponds to $\mathcal{P}_{\boldsymbol{\lambda}}(\Theta) := \{q_{\boldsymbol{\lambda}}(\boldsymbol{\theta}) \mid \boldsymbol{\lambda} \in \mathbb{R}^k\}$ which is a subset of the space of all possible probability densities over Θ , indicated with $\mathcal{P}(\Theta)$. The training of the ansatz is a procedure that starting from an initial distribution identified by $\boldsymbol{\lambda}_0$, through a sequence of update steps of the values of $\boldsymbol{\lambda}$ finds $\boldsymbol{\lambda}^*$ for which $q_{\boldsymbol{\lambda}^*}(\boldsymbol{\theta})$ minimizes the Kullback-Leibler divergence with the true posterior. The point $\boldsymbol{\lambda}^*$ is indicated by the red dot, while the posterior is the red star in the picture. (c) Example of application of variational Bayesian inference to a multimodal single-parameter distribution. The algorithm uses a mixture of Gaussians as ansatz. Maximization of the evidence lower-bound [Eq. (5)], over an increasing number of steps, minimizes the Kullback-Leibler divergence with the true distribution until a sufficiently good approximation is reached, using the means and standard deviations of the three Gaussian as training parameters. More complex distributions can be approximated by a sequence of invertible transformation, for example, implemented by neural networks [81] (e.g., normalized flow). In the plot, the ground truth distribution is shown in black, and the green lines correspond to the variational Bayesian inference approximations. The number on the top left of each subplot corresponds to the iterations index during training.

of “Occam’s razor” heuristic [31,32,49] privileging simple models that fit the data. A Bayesian approach provides explainable insights into the physics underpinning the observations, highlighting the constituent subsystems, the involved physical processes, and the corresponding coupling rates: this is generally less clear with deep-learning strategies to learn models for the generators of dynamics [73–75].

The general framework we introduce here can be applied to a variety of quantum systems. In Sec. IV, we consider the problem of identifying individual nuclear spins by a quantum sensor consisting of a single electronic spin. This example was chosen for its complexity, with the identification of several tens of individual spins demonstrated in experiments [76], and its practical relevance in both developing nanoscale magnetic resonance techniques [14] and quantum networks of nuclear spin qubits [77,78]. More

specifically, this application shows the power of our proposed framework by addressing a complex problem with a nonlinear signal, a number of spins not known in advance (as it depends on the random environment that each individual spin defect experiences), and a highly multidimensional parameter space. We demonstrate in simulation how our framework can perform model selection using a variational Bayesian inference approach, reliably identifying 15 ^{13}C spins (corresponding to 30 hyperfine parameters) with only a fraction of the total data collected in the laboratory. In processing the experimental data we use an ansatz for the Bayesian posterior distribution able to describe up to 40 individual spins, and we show how we are able to compete with state-of-the-art deep-learning approaches [79,80] with limited experimental data and training time. Our Bayesian approach generates an approximate posterior on time scales of a few minutes, and does not assume

prior knowledge of the settings used to acquire the input signal: this makes it perfectly compatible with the future integration of real-time sequential Bayesian design, with the goal to minimize the characterization/sensing time.

II. VARIATIONAL BAYESIAN INFERENCE FOR QUANTUM ESTIMATION

In this section, we first discuss how to construct the outcome probability distributions for parameter estimation in quantum experiments. We will then introduce variational Bayesian inference, and compare it to other approaches used in quantum metrology, such as particle filters and Markov chain Monte Carlo (MCMC), in terms of scalability and precision.

A. Bayesian formulation of the estimation problem

The goal of quantum parameter estimation is to determine the value of $\theta \in \Theta$ characterizing a completely positive, trace-preserving (CPTP) linear map [82,83], which we denote $\mathcal{E}_\theta(\cdot)$. This map corresponds to a physical quantum channel, which describes the (possibly nonunitary) evolution of an input quantum system. In order to complete the estimation task, a probe state $\rho(x)$ is evolved through the quantum channel, producing an encoded state $\rho(x, \theta) := \mathcal{E}_\theta(\rho(x))$. In this manuscript, consistently with the notation adopted in Refs. [84,85], we use $x \in \mathcal{X}$ to indicate the controls available to the experimenter. The goal is to estimate θ from M measurements performed on the set of states $\{\rho(x_t, \theta)\}_{t=1}^M$, and we assume that each encoded state can feature a different set of controls (x_t) for each experiment.

In the case of state estimation, we are given $\rho(\theta)$ directly, without the possibility of manipulating the probe state. After this encoding stage, a positive operator-valued measurement (POVM) is performed on the encoded state. Motivated by experimental feasibility, we assume that only a single copy of the encoded state is available for each measurement.

A generalized measurement on $\rho(x, \theta)$ is represented by the set of operators $\{E(x, y)\}_y$ fulfilling the normalization condition $\sum_y E(x, y) = \mathbb{1}$. The symbol y indicates the outcome of the measurement, while x represents collectively all the control parameters available to the experimenter at the measurement stage. In the following, we will use a single symbol x_t to indicate both the controls on the state preparation and the measurement in an experiment identified by the index $t \in \{1, \dots, M\}$. Throughout this manuscript, we use the term “experiment” to indicate a single step of encoding and measurement, while the set of all experiments carried out to estimate the same θ on the quantum channel, together with the data processing, is called the “estimation”. Given the whole set of M experiments, we indicate with $\mathbf{x}_M := (x_1, x_2, \dots, x_M)$ the

complete list of controls and with $\mathbf{y}_M := (y_1, y_2, \dots, y_M)$ the list of corresponding outcomes.

According to Born’s rule, the probability outcome for y_t is given by

$$p(y_t | x_t, \theta) := \text{Tr}[E(x_t, y_t) \rho(x_t, \theta)]. \quad (1)$$

We label $p(y_t | x_t, \theta)$ as *the model*. The experimental data are processed through Bayes’ formula to get the posterior probability of θ , i.e.,

$$P(\theta | \mathbf{x}_M, \mathbf{y}_M) := \frac{\pi(\theta) \prod_{t=1}^M p(y_t | x_t, \theta)}{\int d\theta \pi(\theta) \prod_{t=1}^M p(y_t | x_t, \theta)}. \quad (2)$$

The Bayesian approach, in addition to enabling uncertainty quantification from the posterior, is efficient in the use of data, and allows seamless incorporation of prior beliefs about the model and its parameters, which are indicated in the above formula with the distribution $\pi(\theta)$.

We will further use the following definitions and nomenclature. The joint likelihood is defined as

$$p(\mathbf{y}_M | \mathbf{x}_M, \theta) := \prod_{t=1}^M p(y_t | x_t, \theta) \quad (3)$$

and

$$P(\mathbf{y}_M | \mathbf{x}_M) := \int d\theta \pi(\theta) \prod_{t=1}^M p(y_t | x_t, \theta) \quad (4)$$

is called the “evidence” or marginal likelihood.

B. Computational approaches

Unless both prior and model fit within specific classes of distributions for which Bayes’ rule can be treated analytically, e.g., the exponential family and conjugate priors, the integral in the denominator of Eq. (2) can only be computed numerically. We focus on four options to apply Bayes’ rule numerically.

The first one, based on *sequential Monte Carlo* (SMC) and particle filtering [86–88], is very popular for quantum estimation experiments that are executed with few parameters [31,44,53,56,69,70,84,85,89]. These computational techniques, based on discretizing the parameter space, do not scale well to high-dimensional systems, as the memory and operations required for the Bayesian update scale exponentially in the number of parameters. Resampling techniques [86,90–95] are useful in mitigating this exponential dependence, but the computational costs remain prohibitive for dozens of parameters. SMC can readily incorporate real-time experiment design, as the posterior is updated online after each new measurement almost instantly.

A second popular computational approach is the use of MCMC [96–98] to approximate the posterior as the steady-state distribution of a Markov process [32,99]. MCMC techniques present a favorable cost scaling in the number of studied parameters, and allow the extension of quantum estimation to high-dimensional problems. Their main limitation is that the Markov chain only faithfully represents the posterior distribution when it has properly thermalized, which can take a significant time (burn-in). The burn-in time cannot be shortened by training multiple chains in parallel, so that even today’s GPUs have limited impact on improving the speed of this technique, which is sequential in nature.

A third popular option is *nested sampling* (NS) [100–102] which transforms the multidimensional posterior integration problem into a one-dimensional evidence calculation by iteratively sampling from prior-constrained shrinking likelihood contours. This method excels at efficiently exploring multimodal posteriors and computing Bayesian model evidence, and can offer a more favorable scaling for high-dimensional problems compared to plain MCMC. In a similar fashion to MCMC, nested sampling implementations might struggle with real-time updates due to sequential sampling needs.

The fourth computational approach is *approximate Bayesian inference*, in particular variational Bayesian inference (VBI) [72,103–107]. This algorithm prescribes the use of a trainable *ansatz*, often based on tailored neural

networks (NNs), to represent the posterior distribution. Although the update of the NN weights in the training process remains sequential, using large mini-batches of sampled parameters and training data at each step can be a very effective way of exploiting the parallelization offered by current GPU hardware. The fast estimation of an approximation to the actual posterior distribution makes real-time adaptive experiments possible.

These reasons motivate the choice to deploy VBI as a computational approach to target parameter estimation in highly multidimensional quantum experiments, with many data points to process. The next subsection briefly recalls the working principles of VBI. See also Fig. 2 for a comparison of the four techniques on a toy model involving multifrequency estimation.

C. Variational Bayesian inference

VBI assumes an ansatz probability density $q_\lambda(\boldsymbol{\theta})$, where λ are trainable parameters, and it aims to minimize the Kullback-Leibler divergence (measuring discrepancy between distributions) between the ansatz and the true posterior, i.e., $\text{KL}(q_\lambda(\boldsymbol{\theta}) \parallel P(\boldsymbol{\theta} \mid \mathbf{x}_M, \mathbf{y}_M))$. However, it is not possible to directly minimize this divergence, since it depends on the posterior whose normalization is typically unknown. It is instead possible to split the dependence of this divergence between λ , i.e., the training variables, and the unknown normalization of the posterior, i.e., the evidence, as in the following: [72,108],

$$\begin{aligned} \text{KL}(q_\lambda(\boldsymbol{\theta}) \parallel P(\boldsymbol{\theta} \mid \mathbf{x}_M, \mathbf{y}_M)) &:= -\mathbb{E}_{\boldsymbol{\theta} \sim q_\lambda} \left[\log \frac{P(\boldsymbol{\theta} \mid \mathbf{x}_M, \mathbf{y}_M)}{q_\lambda(\boldsymbol{\theta})} \right] \\ &= -\mathbb{E}_{\boldsymbol{\theta} \sim q_\lambda} \left[\log \pi(\boldsymbol{\theta}) + \sum_{t=1}^M \log p(y_t \mid x_t, \boldsymbol{\theta}) - \log P(\mathbf{y}_M \mid \mathbf{x}_M) - \log q_\lambda(\boldsymbol{\theta}) \right] \\ &= -\mathbb{E}_{\boldsymbol{\theta} \sim q_\lambda} \left[\log \pi(\boldsymbol{\theta}) + \sum_{t=1}^M \log p(y_t \mid x_t, \boldsymbol{\theta}) - \log q_\lambda(\boldsymbol{\theta}) \right] + \log P(\mathbf{y}_M \mid \mathbf{x}_M) \\ &:= -\text{ELBO}(\lambda) + \log P(\mathbf{y}_M \mid \mathbf{x}_M), \end{aligned} \quad (5)$$

where we have introduced the evidence lower-bound (ELBO). This quantity takes its name from the property $\text{ELBO}(\lambda) \leq \log P(\mathbf{y}_M \mid \mathbf{x}_M)$, which follows from $\text{KL}(q_\lambda(\boldsymbol{\theta}) \parallel P(\boldsymbol{\theta} \mid \mathbf{x}_M, \mathbf{y}_M)) \geq 0$. Instead of directly minimizing the KL divergence, VBI seeks to maximize the ELBO.

The natural candidate for the loss function to be minimized in the training is $-\text{ELBO}(\lambda)$, since it has the same

gradient of the KL-divergence,

$$\frac{\partial}{\partial \lambda} \text{KL}(q_\lambda(\boldsymbol{\theta}) \parallel P(\boldsymbol{\theta} \mid \mathbf{x}_t, \mathbf{y}_t)) = -\frac{\partial \text{ELBO}(\lambda)}{\partial \lambda}. \quad (6)$$

The next question to address is how to construct the ansatz. Assuming $\boldsymbol{\theta}$ consists of continuous parameters, the ansatz $q_\lambda(\boldsymbol{\theta})$ can be constructed starting from a set of normally




FIG. 2. Scalability of variational Bayesian inference for multidimensional parameter estimation. Schematic representation of a 3D Bayesian posterior distribution for three frequencies ($\omega_1, \omega_2, \omega_3$), estimated with four different techniques: (a) Variational Bayesian inference based on a Gaussian ansatz, which depends on 6 parameters only: the three components of the position of the Gaussian μ_1, μ_2, μ_3 and the standard deviations $\sigma_1, \sigma_2, \sigma_3$; (b) particle filtering, with an ensemble of particles used to approximate the posterior (in the figure, the color of the particles represents their weight); (c) Markov chain Monte Carlo, with the posterior distribution represented by an ensemble of samples obtained sequentially from a thermalized Markov chain, which has the Bayesian posterior as fixed point distribution; (d) nested sampling, which evaluates the volume of the region of parameters with a given log-likelihood, and reduces the computation of the multidimensional evidence in Bayesian inference to a 1D integral. (e) Comparison of the mean squared error ϵ in the estimation of n frequencies [Eq. (11)] with different Bayesian estimation algorithms (particle filter: red line; Markov chain Monte Carlo: green line; nested sampling: orange line; variational Bayesian inference: blue line), all operating on the same dataset. The error refers to a single instance of the estimation task, we do not report the statistical fluctuations on the plot. The shaded region corresponds to the error expectation value achievable without any measurements being performed, but guessing at random the frequencies according to the prior: this represent a limit in the performance of all estimation algorithms. (f) Comparison of run time for variational Bayesian inference, nested sampling and Markov chain Monte Carlo, on the same problem discussed in (e). The results show that variational Bayesian inference scales better with increasing problem dimensionality than competing approaches. All the simulations have been performed on the CPU. Particle filtering is not considered here as it is generally fast but, as evidenced in (e), it is not competitive in terms of estimation error in the high-dimensional regime. The results in (e) and (f) support the claim that variational Bayesian inference can implement Bayesian inference on a high-dimensional problem with the possibility to include real-time adaptivity.

distributed variables $\mathbf{z} \sim \mathcal{N}(0, \mathbb{I}_d)$, which are transformed through $\boldsymbol{\theta} := f_{\lambda}(\mathbf{z})$, called a normalizing flow [109] (d is the dimensionality of $\boldsymbol{\theta}$). If f_{λ} is invertible, the transformed probability distribution for $\boldsymbol{\theta}$ is given by

$$q_{\lambda}(\boldsymbol{\theta}) := \mathcal{G}(\mathbf{z}(\boldsymbol{\theta})) \left| \det \frac{\partial f_{\lambda}}{\partial \mathbf{z}}(\mathbf{z}(\boldsymbol{\theta})) \right|^{-1}, \quad (7)$$

where we have indicated with $\mathcal{G}(\mathbf{z})$ the probability density of a multivariate Gaussian and $\mathbf{z}(\boldsymbol{\theta}) := f_{\lambda}^{-1}(\boldsymbol{\theta})$. Crucially,

we need both operations, i.e., sampling from the ansatz and computing the density $q_{\lambda}(\boldsymbol{\theta})$, to be easy to perform in order to compute the ELBO. If we choose f_{λ} to be the affine transformation $\boldsymbol{\theta} = f_{L, \boldsymbol{\mu}}(\mathbf{z}) := L\mathbf{z} + \boldsymbol{\mu}$, with trainable L (lower triangular matrix and invertible) and $\boldsymbol{\mu}$, the ansatz will span the space of all multivariate Gaussian probability densities, when $\boldsymbol{\mu}$ and L are varied. To further extend the class of parametrized distributions, and account for non-Gaussianities, we include a series of nonlinear transformation before $f_{L, \boldsymbol{\mu}}$, so that the complete transformation

$f(\mathbf{z})$ is given by

$$f(\mathbf{z}) := f_1 \circ f_2 \circ \dots \circ f_K \circ f_{L,\mu}(\mathbf{z}). \quad (8)$$

These nonlinear transformations can be implemented as neural autoregressive layers [81], and are defined to be invertible and such that the corresponding log-determinant needed in the ELBO can be computed efficiently. Note that this choice is not exclusive and it is also possible to choose other neural architectures, or simply choose $q_\lambda(\theta)$ as a classical family of parametric (exponential) distributions. For instance, in the next subsection, $q_\lambda(\theta)$ is chosen to be a multivariate Gaussian density with diagonal covariance matrix.

D. Scalability of the Bayesian estimation procedure

In this subsection, we discuss an example of how variational Bayesian inference outperforms particle filtering as the number of parameters in the model increases. We consider a simple toy model where multiple frequencies are encoded on a probe quantum state when interacting with an environment, through the CPTP map,

$$\mathcal{E}_\omega(\cdot) = \frac{1}{n} \sum_{i=1}^n U_{\omega_i} \cdot U_{\omega_i}^\dagger. \quad (9)$$

The goal is to estimate the vector of unknown frequencies $\omega := (\omega_1, \omega_2, \dots, \omega_n)$, see Fig. 2. This corresponds to a Ramsey experiment with a single qubit, initially prepared in the state $|\psi\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, which evolves freely in the environment for a time τ , according to unitary operators $U_{\omega_i} := e^{-\frac{i}{2}\tau\omega_i\sigma_z}$. A qubit measurement in the superposition basis outputs the outcomes $|\pm 1\rangle$, with probabilities, respectively, given by

$$p(\pm 1 | \tau, \omega) := \frac{1}{2} \pm \frac{1}{2n} \sum_{i=1}^n \cos(\omega_i \tau). \quad (10)$$

The frequencies are randomly selected in the range $\omega_i \in [0, 1]$ MHz uniformly. We solve this multidimensional estimation problem for an increasing number of frequencies, comparing VBI with different Bayesian estimation techniques (MCMC, nested sampling, and particle filtering). The particle filter evaluates the probability distributions over a number $N_p = 16384$ of points (known as “particles”) in the multidimensional parameter space, implemented as detailed in Refs. [84,85]. We have chosen the dataset and the number of parameters n such that the size of the problem is similar to the estimation of hyperfine parameters discussed in Sec. IV. The dataset is composed of $M = 4096$ points with different τ , selected in such way that $\log(\tau/1 \mu\text{s})$ is sampled uniformly in $[-1, +5]$. This choice is motivated by the error analysis of phase

estimation in Ref. [110]. For each value of τ , $R = 1024$ experiments with a binary outcome are simulated.

Nested sampling is based on evaluating the volumes of the shells of parameter space with a given log-likelihood. For this comparison we have implemented adaptive shrinking of the prior volume and replacement of low-likelihood points through MCMC and slice sampling under likelihood constraints [111]. The algorithm features adaptive proposal distributions, likelihood-constrained slice sampling, and dynamic live point populations, which are all standard improvements over the baseline algorithm [112]. The initial and maximum number of live points are respectively 20 and 500, we have simulated 4096 iterations, with 500 MCMC steps used, inside the NS algorithm, to sample new live points in each iteration. The last error in Fig. 2 refers to inference using MCMC, which is based on sampling from a Markov chain whose equilibrium distribution is the Bayesian posterior. We have implemented adaptive Metropolis-Hastings MCMC, which generates posterior samples via random walk proposals from a multivariate normal [113]. We simulate 3×10^6 steps of a single Markov chain, with 10% burn-in. As described above, VBI does not evaluate the posterior over a set of particles or through sampling, but approximates it with a set of simpler functions chosen as ansatz. Here we use the “mean field” ansatz, i.e., we assume the parameters are uncorrelated and we approximate the posterior with a product of Gaussians. The trainable parameters of the posterior are therefore the mean values μ_i and the standard deviations of each Gaussian σ_i , i.e., for n frequencies the number of trainable parameters scales polynomially as $2n$. Labeling the Bayesian estimator for ω_i as $\hat{\omega}_i$, we compute the error in a single estimation as

$$\varepsilon := \frac{1}{n} \sum_{i=1}^n (\hat{\omega}_i - \omega_i)^2, \quad (11)$$

where we choose to order the list of frequencies and their estimators in ascending order, i.e., $\omega_1 < \omega_2 < \dots < \omega_n$ and $\hat{\omega}_1 < \hat{\omega}_2 < \dots < \hat{\omega}_n$.

Figure 2 compares the error ε as a function of the number of frequencies for all the techniques described here. The shaded gray region is the average error per frequency achievable when no measurements are carried out. This is not a constant; on the contrary, it slightly decreases as the number of estimated frequencies grows. This effect is due to the permutation invariance of the frequency labels and can be understood as follows. For an arbitrary frequency ω , if no measurements are performed, each of the $\hat{\omega}_i$ estimators is selected randomly and uniformly in the parameter domain. The larger the number of unknown frequencies n , the higher the probability that one of the n estimators $\hat{\omega}_i$ will randomly fall close to the selected ω . Therefore, we expect the average error to decrease with the number of frequencies. We observe that VBI with mean field ansatz

is as good as MCMC, but it has a more favorable time scaling.

III. LEARNING MODELS FOR QUANTUM SYSTEMS USING VARIATIONAL BAYESIAN INFERENCE

In this section we discuss the selection of a model within a set of possible models using VBI, as summarized in Fig. 3.

A. Model selection for a set of nested models

In the previous section, we have discussed the application of Bayesian inference to the estimation of a fixed number of parameters. We will now extend the discussion to model selection, to describe cases where there might be a number of different descriptions of a given quantum system, each with a different number of parameters, but all compatible with the observed data. This is useful, for example, to characterize the interaction of a quantum system with its environment, whose composition might not be precisely known. Examples of this situation are a central electron spin interacting with a number of nuclear spins, where both the number and couplings of the spins

are unknown a priori (discussed in detail in Sec. IV), or a superconducting qubit coupled to a bath of two-level fluctuators [114–117].

We assume the quantum system can be described by a discrete set of classes $\mathcal{C}$ (possibly infinite). Each class $c \in \mathcal{C}$ is described by a set of continuous parameters $\theta_c \in \Theta_c$, possibly with a different number of parameters for each class c .

The model selection problem involves finding the class c and the values of the associated parameters θ_c that explain the experimental data. To achieve this, instead of a single posterior distribution $P(\theta | x_t, y_t)$, we need to specify the list of marginal posteriors conditioned on the model's class, i.e., $\{P(\theta_c | c, x_t, y_t)\}_{c \in \mathcal{C}}$, and the list of class probabilities $\{P(c | x_t, y_t)\}_{c \in \mathcal{C}}$. This increased complexity affects the application of Bayes' rule, which in this case needs to account for both the posterior in the continuous parameters of each model and the probability over $\mathcal{C}$.

We assume that among different classes of models there is a natural partial ordering, defined by identifying a model as a submodel of a more general class [Fig. 3(a)]. For example, a system with $M > N$ qubits is more general than the model with N qubits (and includes it as a particular




FIG. 3. Model selection. (a) On the left, we show a schematic representation of a hierarchical set of model, where a is a submodel of b , which is a submodel of c^* . On the right, we represent the nested models of an environment constituted of different numbers N of qubits. (b) The regularization procedure described in Sec. III B enables selecting the simplest model that explains the data. L_1 regularization [Eq. (12)] enforces sparsity, setting the parameters that are not important exactly to zero, thereby pruning them from the model. L_2 regularization [Eq. (15)] results in values close to zero (but not exactly zero) for parameters which impact the outcome less: these can be then pruned in postprocessing by imposing a threshold. (c) Here we show an example of L_2 regularization. We consider a set of models with a central spin and, at maximum, $N = 27$ spins in the surrounding bath. We then perform a Ramsey experiment on the central spin, in which the central spin is prepared in an equal superposition and let freely evolve under the influence of the bath. The outcome signal consists of multiple decaying oscillations, corresponding to the distribution of magnetic field values given by the different spin configurations. Variational Bayesian inference with L_2 regularization learns the coupling parameters to the 27 environmental spins: in this case, 24 of them are close to zero. A threshold separates the 3 spins that contribute to the oscillating signal from the 24 background spins that contribute to the decaying envelope (corresponding to the T_2^* dephasing time). If we are both estimating T_2^* and all the spins in the environment, we could choose to use the L_1 regularization, set the coupling of the background spins exactly to zero in the ansatz and explain the envelope with a nonzero $1/T_2^*$. In Sec. IV, we consider a similar setting with a different experiment, using both (T_2^*) and the background spins.

case), when the couplings of the remaining qubits are set to zero.

Within this formalism, we assume we can construct a pruning function which, given a class $\mathbf{b} \in \mathcal{C}$ and the continuous parameters $\theta_{\mathbf{b}}$, returns the class $\mathbf{a} \in \mathcal{C}$ and its associated parameters such that $\mathbf{a}$ is a submodel of $\mathbf{b}$ that has been pruned of all the zero parameters $(\theta_{\mathbf{b}})_i = 0$. For example, given a model with M qubits, if $M - N$ qubits have zero couplings, then the pruning of these noninteracting qubits gives an equivalent model with N qubits.

In other words, we assume the existence of a class $\mathbf{c}^* \in \mathcal{C}$ such that we can view all models as its subclasses, after setting some of the parameters to zero, i.e., we can reduce the problem of model learning to parameter estimation for $\theta_{\mathbf{c}^*}$, where the goal is now both to minimize the number of nonzero elements of the vector of continuous parameters and to find the best values for such nonzero parameters.

Under these assumptions, the direct application of Bayes' rule over the set of classes and their parameters cannot identify the minimum viable explanation for the experimental data, as, if $\mathbf{a}$ is a submodel of $\mathbf{b}$, then $\mathbf{b}$ is at least as likely as $\mathbf{a}$ to be the underlying model, once the parameters $\theta_{\mathbf{a}}$ and $\theta_{\mathbf{b}}$ have been accordingly selected for both classes. In general, there could be a trade-off between the error in explaining the observed data and the complexity of the model selected, which can be tuned through a regularization or a thresholding parameter.

B. Selecting a minimal model

In this section we discuss the use of regularizing priors and thresholding to select a minimum viable model for explaining the data. We will outline the algorithm capable of extracting the salient features of a model explaining the data, in accordance with a hierarchical view of the class of models. According to the problem this algorithm is being applied to, regularization or thresholding could have different roles and contribute differently to the final outcome. A closer look at the structure of the particular problem at hand is therefore fundamental.

1. Regularizing prior

With a nested set of models, we approach the model selection problem with an "Occam's razor" heuristic, favoring the models that can fit the experimental data with the least number of parameters.

In order to penalize models with a higher number of nonzero parameters, we substitute the prior $\pi(\theta)$ in Eq. (2) with $\pi(\theta) \cdot r(\theta)$, where $r(\theta)$ is the L_1 -regularizer. This corresponds to a Laplace distribution

$$r(\theta) = r_1(\theta) := \frac{1}{(2\sigma)^d} e^{-\frac{\|\theta\|_1}{\sigma}}, \quad (12)$$

where the parameter σ acts as the regularization scale, d is the size of $\theta \in \mathbb{R}^d$, and $\|\theta\|_1 := \sum_{i=1}^d |\theta_i|$. With this

addition, the ELBO becomes

$$\begin{aligned} \text{ELBO}'(\lambda) := \mathbb{E}_{\theta \sim q_\lambda} & \left[\log \pi(\theta) + \log r(\theta) \right. \\ & \left. + \sum_{t=1}^M \log p(y_t | x_t, \theta) - \log q_\lambda(\theta) \right]. \end{aligned} \quad (13)$$

By inserting the explicit expression for $r(\theta)$, the ELBO becomes,

$$\text{ELBO}'(\lambda) := \text{ELBO}(\lambda) - \mathbb{E}_{\theta \sim q_\lambda} \left[\frac{\|\theta\|_1}{\sigma} \right], \quad (14)$$

where we have neglected a constant term. In this context, Laplace regularization acts as the LASSO in machine learning [118,119], and has been previously used, for example, in compressed sensing for quantum tomography [120] and to reduce the number of CNOT gates in the compilation of quantum circuits [121]. The L_1 -regularizer is the most commonly used to find a sparse solution, where those parameters not needed to explain the data will be set to zero and pruned. This regularization is in the spirit of the generalized Bayesian information criterion. The differences are that the norm does not count directly the number of nonzero parameters, it requires the introduction of an appropriate scale σ , and it can be used online during the training of the Bayesian posterior, unlike GIC.

A different choice for the regularization prescribes the use of a Gaussian distribution:

$$r(\theta) = r_2(\theta) := \frac{1}{(2\pi\sigma^2)^{\frac{d}{2}}} \exp\left(-\frac{\|\theta\|_2^2}{2\sigma^2}\right), \quad (15)$$

where $\|\theta\|_2^2 := \sum_{i=1}^d |\theta_i|^2$, which produces the regularized ELBO,

$$\text{ELBO}'(\lambda) := \text{ELBO}(\lambda) - \mathbb{E}_{\theta \sim q_\lambda} \left[\frac{\|\theta\|_2^2}{2\sigma^2} \right]. \quad (16)$$

The Gaussian regularizer corresponds to ridge regression [122] in machine learning. Instead of favoring a sparse solution, Gaussian regularization finds more balanced solutions, with the tendency to shrink θ toward zero.

The interpretation of the modified ELBO is that we pay a price each time one of the parameters θ_i is not exactly zero or close to zero, according to the choice of the norm. This will nudge the training toward models that explain the data correctly, i.e., maximize the ELBO, but also pushes unused parameters to zero.

The choice between L_1 or L_2 -regularization depends on the specific model under consideration. In practice, we

observed that the L_1 regularization is generally suited to sharply exclude features and select between competing models, favoring the least complex ones, while the L_2 is useful in scenarios where there is a background signal. In this latter case, we want to better account for this background in the training of the posterior, and postpone the proper model selection to the postprocessing, which is now a selection of the most salient features that stand out from the background.

In this work we have explored the use of both L_1 and L_2 regularization, and while we have observed L_2 to be better in simulation, we have concluded, empirically, that L_1 works better on the experimental data. The plots and numerical results reported, for simulated (L_2) and experimental data (L_1), reflect this choice. We stress that this result might be dependent on the particular dataset and choices of hyperparameters for the inference; a complete analysis of benefits of the two regularizers in model selection goes beyond the scope of this work.

2. Parameters thresholding and confidence level

Once the ansatz $q_\lambda(\theta)$ for the posterior distribution is trained, the most-likely least-complex model that describes the experimental observations needs to be extracted. According to the selected regularizer, a fraction of estimated parameters will be either exactly zero, or very close to zero. The (postprocessing) pruning steps using $q_\lambda(\theta)$ are described below.

- (1) $\theta \sim q_\lambda(\theta)$ is sampled and the mask $\theta'_j := \theta_j \cdot \mathbb{1}_{\theta_j > \theta_j^{\text{th}}}$ is applied. This ensures that when the j -th component of the sampled vector is less than a threshold θ_j^{th} , it is exactly set to zero through the multiplication with the indicator function. We define $\theta' = (\theta'_1, \theta'_2, \dots, \theta'_d)$.
- (2) The pruning function is then applied to θ' , obtaining the minimal model class c that only has the nonzero identified components, and its associated θ_c .
- (3) The sampling and the postprocessing delineated in (1) and (2) are repeated Z times, and the θ_c are classified according to the number of nonzero parameters, in order to build a set of samples $\mathcal{S}_c := \{\theta_1, \theta_2, \dots\}$ for each $c \in \mathcal{C}$ observed.

- (4) To each class c we associate a “pseudo-Bayesian” probability, defined as $p_c := |\mathcal{S}_c|/Z$, as the probability for the model c to be the source of the observed data, i.e., the fraction of the samples of the trained posterior, over all the samples, that have been reduced to the class c .

We identify $\{p_c\}_{c \in \mathcal{C}}$ as the probabilities $P(c | x_t, y_t)$ on $\mathcal{C}$, while $\mathcal{S}_c$ is a set of samples from the posterior distribution for θ conditioned on c . The class estimator will be the maximum a posteriori estimator (MAP), i.e., $\hat{c} := \text{argmax}_c p_c$.

C. Uncertainty estimation

We denote $\{p_c\}_{c \in \mathcal{C}}$ as pseudo-Bayesian probabilities, as they are not determined from the observed data only, but also depend on the regularization and the thresholding, in short from the model selection algorithm. The role of p_c is to offer a certainty level for the class c to be the simplest one explaining the observations. Ideally, we want p_c to be as concentrated as possible, minimizing uncertainty about the minimum complexity class that fits the data. From the set of samples $\mathcal{S}_c$, conditioned on the model’s class c , we can then evaluate any function of the parameters and their uncertainty after postprocessing the samples, see Sec. IV.

D. Nuisance parameters

Equation (1) defines the model used in VBI. In some cases, one might not want to treat all parameters in a Bayesian way, for example, because their uncertainties are irrelevant for the purpose of the experimenter or negligible, but these parameters still need to be included in the inference as they impact the observation model. Examples could be some intrinsic properties of a qubit (e.g., its coherence time), or the value of an externally applied magnetic or electric field. We call these “nuisance parameters” and we label them as φ . We then write the model as $p(y_t | \theta, \varphi, x_t)$, with the observation probability depending on the Bayesian parameters θ , as well as on the nuisance φ and the controls x_t . Nuisance parameters can be properties of the probe state, of the measurement, or of the evolution channel.

The ELBO depends now explicitly on φ ,

$$\text{ELBO}(\lambda, \varphi) := \mathbb{E}_{\theta \sim q_\lambda} \left[\log \pi(\theta) + \sum_{t=1}^M \log p(y_t | x_t, \theta, \varphi) - \log q_\lambda(\theta) \right]. \quad (17)$$

Maximizing the ELBO at the same time with respect to λ and φ will perform VBI on θ and maximum expected likelihood estimation on φ , which follows from the expression of the gradient of the ELBO with respect to φ , i.e.,

$$\frac{\partial \text{ELBO}(\theta, \varphi)}{\partial \varphi} = \sum_{t=1}^M \frac{\partial}{\partial \varphi} \mathbb{E}_{\theta \sim q_\lambda} [\log p(y_t | x_t, \theta, \varphi)]. \quad (18)$$

In the practical example discussed in Sec. IV, various decoherence parameters are treated as nuisance parameters and optimized with maximum likelihood.

IV. EXAMPLE: QUANTUM SENSING OF INDIVIDUAL NUCLEAR SPINS

As an example of the application of our multidimensional quantum estimation and model selection protocol, we consider the case of the identification of individual nuclear spins with a single electron spin quantum sensor. This is a problem of practical importance for different quantum technology applications [123].

The detection of single nuclear spins is motivated by the desire to push magnetic resonance to the nanoscale [14,124–131] in order to image single molecules with near-atomic spatial resolution. Quantum sensors based on single electronic spins are promising for achieving this target, as close proximity to the sample enables much higher sensitivity for magnetic dipole signals that decay fast as a function of distance, over a nanoscale sensing volume.

Additionally, individual nuclear spins are promising for the implementation of long-lived qubits [77,132–138], as their minimal interaction with the environment leads to long coherence times. These properties make individual nuclear spins attractive building blocks for networked quantum information processing, particularly in architectures where a single electronic spin in a solid-state defect serves as the interface between the local nuclear register and photonic communication channels.

The experiments to identify nuclear spins consist of long pulse sequences to achieve sufficient spectral resolution to discriminate small differences in hyperfine coupling strength. The complexity of the problem and the required fine spectral details lead to long data acquisition times and very large datasets, which are laborious to both acquire and analyze.

This complexity has stimulated work to find algorithms for automated data processing. Previous approaches [79,

80] have employed deep-learning techniques, framed the problem as pattern recognition or signal-to-image conversion, and demonstrated high performance and speed in identifying nuclear spins from experimental data traces. The drawback of deep-learning approaches is, however, that they are tailored to a specific set of measurements, with fixed hyperparameters (e.g., external magnetic field strength) and to a specific platform, requiring an extensive training time before being able to operate. This is incompatible with adaptivity, as adaptive protocols decide in real time which data is best to acquire and use as input for the evaluation.

Parallel work by Poteshman *et al.* [99,139] develops a model selection and parameter estimation procedure based on reversible-jump Markov chain Monte Carlo, which they deploy for the same case study investigated here. The main advantage of the variational approach presented in our manuscript is that it is more suitable for exploiting parallelization, with a much shorter training time for the posterior, which can enable the implementation of real-time sequential Bayesian experiment design. We will present a brief comparison between the two approaches in Sec. IV G.

A. Problem statement

We consider the experimental configuration in Fig. 4(a). As a quantum sensor, we consider the $S = 1$ electron spin associated with a nitrogen-vacancy (NV) center in diamond. The NV electronic spin is widely used in quantum sensing applications, as it can be optically initialized and readout and features excellent quantum coherence in a wide temperature range, up to room temperature and beyond [141–143].

In a suitable rotating frame and under the secular approximation, the NV single electron spin (indicated by the operator $\hat{S}$) interacts with surrounding ^{13}C nuclear spins (operators $\hat{\mathbf{I}}_k, k = 1 \dots K$) through the following Hamiltonian [144]:

$$\hat{\mathcal{H}} := D\hat{S}_z^2 + \gamma_e \mathbf{B} \cdot \hat{\mathbf{S}} + 2\pi \sum_{k=1}^K \left[A_{z,k} \hat{S}_z \hat{I}_{z,k} + A_{\perp,k} \hat{S}_z \hat{I}_{x,k} \right] + \sum_{k=1}^K \gamma_n B_z \hat{I}_{z,k}. \quad (19)$$

Here D is the NV zero-field splitting ($D = 2\pi \times 2.87$ GHz), B_z is the external magnetic field applied along the NV axis, γ_n is the ^{13}C gyromagnetic ratio, and $\{(A_{z,k}; A_{\perp,k}), k = 1, \dots, K\}$ are the hyperfine parameters, coupling the NV electron spin with the ^{13}C nuclear spins.

This model neglects the dipolar coupling between the nuclear spins, considering them as independent. While

this approximation is not fully valid, most nuclear–nuclear interactions are small (< 20 Hz) and can be approximately neglected on the time scales of the experiments considered. While we use this approximation here for simplicity, the extension of the algorithm to include nuclear spin interactions is straight-forward, only computationally more expensive due to the more complex Hamiltonian.




FIG. 4. Example of application to nanoscale magnetic resonance experiments. (a) We consider a single electron spin associated with an nitrogen-vacancy center in diamond, which can be optically initialized and read out. The electron spin is controlled by microwave pulses, and couples to a bath of independent ^{13}C nuclear spins in the diamond through the hyperfine interaction. (b) Measurements consist of dynamical decoupling sequences, comprising N_π repetitions of the $(\pi - \tau - \pi)$ unit on the electron spin, initially prepared into an equal superposition state. (c) Each nuclear spin, with hyperfine described by the values $\{A_{z,k}; A_{\perp,k}\}$ results in multiple resonances with linewidth on the order $\mathcal{O}(A_{\perp}/f_L^2)$, with f_L the Larmor frequency of the spin. (d) The signal generated by four spins includes multiple harmonics of index m ; the shift from the inverse of the Larmor frequency is controlled by A_z , i.e., the parallel hyperfine coupling of each spin. (e) The outcome signal, upon repeating R experiments consisting of spin initialization, DD sequence, and electron spin readout, is the probability for the electron spin to be in $|0\rangle$ as a function of the interpulse delay τ in the DD sequence. The measurement signal presents a series of dips occurring at values of τ related to the parallel hyperfine of a given nuclear spin by $\tau_m \approx (m - \frac{1}{2})/(2f_L + A_z)$. The data in the figure are a subset of the experimental data from Abobeih *et al.* [140].

Our goal is to estimate the hyperfine parameters $\{(A_{z,k}; A_{\perp,k}), k = 1, \dots, K\}$, with the number K of spins a priori unknown. We denote $\mathbf{A}_z := (A_{z,1}, A_{z,2}, A_{z,3}, \dots, A_{z,K})$ and $\mathbf{A}_{\perp} := (A_{\perp,1}, A_{\perp,2}, A_{\perp,3}, \dots, A_{\perp,K})$ the parallel and orthogonal hyperfine couplings, respectively, and with

$$\mathbf{A} := (A_{z,1}, A_{\perp,1}, A_{z,2}, A_{\perp,2}, \dots, A_{z,K}, A_{\perp,K}) \quad (20)$$

the full vector of hyperfine couplings.

We consider here dynamical decoupling measurements performed by preparing the electron spin in an equal superposition $(|m_s = 0\rangle + |m_s = 1\rangle)/\sqrt{2}$, and applying a sequence $(\tau - \pi - \tau)$ repeated N_π times, where “ π ”

represents a π -pulse on the electron spin, and τ the interpulse delay [Fig. 4(b)]. In the frequency domain, this corresponds to a filter [145,146] with central frequency determined by the interpulse delay τ , and width inversely proportional to N_π . This is just one of the possible pulse sequences investigated in the literature for the identification of nuclear spins: while more sophisticated sequences have been investigated [77,147–149], we take dynamical decoupling as a benchmark example.

Assuming single-shot electron spin readout [150,151], the outcome of a measurement on the electron spin results in a random Bernoulli variable with outcome probability given by Refs. [140,144]

$$p(y = 0, 1 \mid x = (\tau, N_\pi), \boldsymbol{\theta} = \mathbf{A}) := \frac{1}{2} \left(1 + (-1)^b e^{-\left(\frac{2N_\pi\tau}{T_2}\right)^\xi} \prod_{k=1}^K \mathcal{M}(\mathbf{A}_k, \tau, N_\pi) \right), \quad (21)$$

with

$$\mathcal{M}(\mathbf{A}_k, \tau, N_\pi) := 1 - m_{k,x}^2 \cdot \frac{(1 - \cos \alpha_k)(1 - \cos \beta)}{1 + \cos \alpha_k \cos \beta - m_{k,x} \sin \alpha_k \sin \beta} \cdot \sin^2 \left(\frac{N_\pi \varphi_k}{2} \right), \quad (22)$$

where $\tilde{f}_k := \sqrt{(A_{z,k} + f_L)^2 + A_{\perp,k}^2}$, $\tilde{\omega}_k := 2\pi\tilde{f}_k$, $\alpha_k := \tilde{\omega}_k\tau$, $\beta := \omega_L\tau$, $m_{k,z} := (A_{z,k} + f_L)/\tilde{f}_k$, $m_{k,x} := A_{\perp,k}/\tilde{f}_k$, $\cos\varphi_k := \cos\alpha_k \cos\beta - m_{k,z} \sin\alpha_k \sin\beta$, and finally the Larmor frequency of the ^{13}C is defined as $f_L := \gamma_n B_z$ and $\omega_L := 2\pi f_L$. We have written Eq. (21) in a way that makes it clear that b is the binomial outcome of the measurement, $x := (\tau, N_\pi)$ is the control, i.e., the parameters that characterize the design of the DD pulse sequence, and $\theta = \mathcal{A}$, i.e., the hyperfine couplings are to be treated in a Bayesian way. Note that with K spins in the estimation the dimension of the space of parameters is $d = 2K$. The time T_2 is the characteristic time of the decoherence process, mainly determined by the interactions between the ^{13}C spins in the bath, which are not explicitly included in the model in Eq. (19). We treat T_2 as a nuisance parameter (Sec. III D).

As depicted in Figs. 4(b) and 4(c), in the limit of magnetic field strength larger than the hyperfine, each nuclear spin results in a series of resonances in the signal, occurring at the interpulse delay

$$\tau_{m,k} \approx \frac{m - \frac{1}{2}}{2f_L + A_{z,k}}, \quad (23)$$

where m is an integer. Each resonance is Lorentzian in shape, with width $\mathcal{O}(A_{\perp,k}/f_L^2)$ and depth $\mathcal{O}(N_\pi A_{\perp,k}/f_L)$. An example of the DD signal as a function of the interpulse delay τ is shown in Fig. 4(d). An intuitive discussion of the Hamiltonian in Eq. (19) and the corresponding DD signal is presented in Appendix A.

The model [Eq. (21)] manifests two symmetries. Given a set of hyperfine couplings $\mathcal{A}$ we can permute the labels of the spins and obtain the same signal. Furthermore, the model is symmetric for $A_{\perp,k} \rightarrow -A_{\perp,k}$, which gives yet another degeneracy. In the estimation, we will fix $A_{\perp,k} \geq 0$. The development of symmetry-preserving ansatzes is an open problem in VBI, which we briefly discuss in Sec. V. To conclude this section, we observe that the spatial location of a nuclear spin is not unequivocally determined by knowledge of the $(A_z, A_\perp)$ couplings. The hyperfine couplings are however relevant parameters needed to address the spins individually and implement quantum gates [144].

B. Outcome model in VBI

We define a single experiment to comprise $R = 10^3$ repetitions of the process of initialization of the NV center, application of the N_π -pulses DD sequence, and measurement. We define the outcome of an experiment as

$$y := \frac{1}{R} \sum_{r=1}^R b_r + \delta, \quad (24)$$

where $\delta \sim \mathcal{N}(0, \eta_0^2)$ is the measurement noise in addition to the shot noise, and we have:

$$\lim_{R \rightarrow \infty} \mathbb{E}[y] = p := p(1 \mid (\tau, N_\pi), \mathcal{A}). \quad (25)$$

The variance (noise) on the measurement y is the sum of the shot noise plus the Gaussian error, which models for us all other sources of errors in the experiment ($\eta_0 := 10^{-2}$ for the data generated in simulation), i.e.,

$$\text{Var}(y) := \frac{p(1-p)}{R} + \eta_0^2. \quad (26)$$

In the estimation we choose to model the outcomes as Gaussian variables, sampled from the distribution

$$y \sim \mathcal{N}\left(p, \frac{p(1-p)}{R} + \eta^2\right), \quad (27)$$

where η (initialized at $\eta_0 = 10^{-2}$), and T_2^{-1} (initialized at zero), are the nuisance parameters in the estimation, i.e., $\varphi := (T_2^{-1}, \eta)$, and are therefore estimated with expected maximum likelihood, as explained in Sec. III D. We stress that, while the generation of the simulation dataset involves sampling from a binomial distribution, the model we use in VBI treats y as a Gaussian variable. We have used the same model for processing both simulated and experimental data. Notice also that the probability distribution for y has a complicated dependence on the hyperfine couplings $\mathcal{A}$ through the probability p . In the following sections we indicate with $\mathcal{G}(y \mid x, \mathcal{A}, \varphi)$ the density function of the Gaussian in Eq. (27), which plays the role of the model for the outcome probability in nano-NMR, i.e., $p(y \mid x, \theta)$ in Eq. (1).

C. Algorithm deployment

Here we illustrate the application of the multiparameter quantum estimation and model selection procedure described in Secs. II C and III to the nanoscale magnetic resonance problem.

The expression for the ELBO in Eq. (5) is evaluated numerically by approximating the expectation value over the ansatz with a Monte Carlo approach, i.e., we compute $\mathcal{B}$ extractions of hyperfine couplings from the current ansatz $q_\lambda(\mathcal{A})$ to approximate the ELBO as the following average:

$$\text{ELBO} := \frac{1}{\mathcal{B}} \sum_{j=1}^{\mathcal{B}} \left[\log \pi(\mathcal{A}) + \log r(\mathcal{A}) + \sum_{t=1}^{M_B} \log \mathcal{G}(y_t \mid x_t, \mathcal{A}, \varphi) - \log q_\lambda(\mathcal{A}) \right]. \quad (28)$$

Throughout the paper we use an uninformative, unnormalized prior $\pi(A) \propto 1$. We always use $M_B = M$, i.e., we process the whole dataset in each computation of the ELBO, but it is possible to select smaller mini-batches to ease the computational requirements of the inference process. For the processing of nano-NMR we use a normalizing flow as the ansatz, constituted by five layers of neural autoregressive layers [81] followed by one last affine layer. See Sec. II C for a definition of the normalizing flow.

D. Posterior postprocessing

In Sec. III B 2 we have outlined the working principles of the postprocessing of the trained posterior. This is used to prune unnecessary parameters and obtain the minimal model that explains the data. For the particular case of nano-NMR, a lower threshold $A_\perp^{\text{th}}$ is imposed, while there is no lower threshold on A_z , although we choose to process and visualize only samples within $A_z < A_z^{\text{th}} \simeq 80$ kHz. When the estimation is carried out with an ansatz for the

posterior containing more spins than required to accurately describe the experimental data, the spins in excess are pushed below the threshold by the maximization of the ELBO, and are finally cut by the thresholding. Also spins that are in the signal, but are too weak to be above threshold, are cut. The threshold therefore represents the cut-off of a blind zone, beyond which the coupling of the spins with the NV center is too weak for the spins to be individually detected from the dataset considered.

The thresholding is performed only on $A_\perp$ and not on A_z because the orthogonal component of the hyperfine coupling is the parameter controlling the amplitude and width of the dips, i.e., the main features determining whether a dip is detectable within the noise level.

With an ansatz containing K spins, excluding the class with zero spins, the thresholding procedure described in Sec. III B 2 leaves us with K sets $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_K$, which contain all the samples classified according to the number of spins, i.e.,

$$\mathcal{S}_1 := \left\{ \begin{array}{c} (A_z^1, A_\perp^1)_1, \\ (A_z^1, A_\perp^1)_2, \\ (A_z^1, A_\perp^1)_3, \\ \vdots \\ (A_z^1, A_\perp^1)_{|\mathcal{S}_1|} \end{array} \right\}, \quad \mathcal{S}_2 := \left\{ \begin{array}{c} (A_z^1, A_\perp^1, A_z^2, A_\perp^2)_1, \\ (A_z^1, A_\perp^1, A_z^2, A_\perp^2)_2, \\ (A_z^1, A_\perp^1, A_z^2, A_\perp^2)_3, \\ \vdots \\ (A_z^1, A_\perp^1, A_z^2, A_\perp^2)_{|\mathcal{S}_2|} \end{array} \right\}, \dots \quad \mathcal{S}_n := \left\{ \begin{array}{c} (A_z^1, A_\perp^1, A_z^2, A_\perp^2, \dots, A_z^n, A_\perp^n)_1, \\ (A_z^1, A_\perp^1, A_z^2, A_\perp^2, \dots, A_z^n, A_\perp^n)_2, \\ (A_z^1, A_\perp^1, A_z^2, A_\perp^2, \dots, A_z^n, A_\perp^n)_3, \\ \vdots \\ (A_z^1, A_\perp^1, A_z^2, A_\perp^2, \dots, A_z^n, A_\perp^n)_{|\mathcal{S}_n|} \end{array} \right\}, \dots \quad (29)$$

So far we have been following the general framework laid down in Sec. III. The rest of this section is specific to the determination of the number of spins generating the signal (model selection) and the determination of the hyperfine couplings (parameter estimation). For each dataset $\mathcal{S}_n$ containing $|\mathcal{S}_n|$ samples, we break the correlations between the spins in each single sample from the posterior and we build an ensemble of single spins

$$\tilde{\mathcal{S}}_n := \{(A_{z,j}, A_{\perp,j})\}_{j=1}^{n|\mathcal{S}_n|}, \quad (30)$$

for $n \leq K$. These are now samples from a 2D probability distribution, which is a marginalized version of the full posterior distribution. Although this process discards much information, it is a necessary step to obtain a condensed version of the posterior for two-dimensional plotting and clustering. We then apply the K-means clustering algorithm to this dataset, with an optimal number of clusters obtained by testing the inertia for each number of clusters between 1 and n . This procedure leaves us with a partition of $\tilde{\mathcal{S}}_n$ in sets of points belonging to different

clusters:

$$\tilde{\mathcal{S}}_n = \bigcup_{j=1}^{L_n} K_{n,j}, \quad (31)$$

with L_n being the total number of clusters. For each cluster of points $K_{n,j}$ we then compute three quantities:

- (1) the average value of the 2D vector of couplings, i.e.,

$$\mu_{n,j} := \frac{1}{|K_{n,j}|} \sum_{(A_z, A_\perp) \in K_{n,j}} (A_z, A_\perp). \quad (32)$$

- (2) the covariance matrix of the cluster, i.e.,

$$\Sigma_{n,j} := \frac{1}{|K_{n,j}|} \sum_{(A_z, A_\perp) \in K_{n,j}} \begin{pmatrix} A_z \\ A_\perp \end{pmatrix} \begin{pmatrix} A_z & A_\perp \end{pmatrix} - \mu_{n,j}^T \mu_{n,j}. \quad (33)$$

(3) the “weight” of the cluster, i.e.,

$$\#S_{n,j} := \frac{|K_{n,j}|}{|\mathcal{S}_n|}. \quad (34)$$

The weight of the cluster gives us information on the number of spins in each cluster, and it can be a fractional number. The weights of all clusters are normalized to the number of spins n , i.e., $\sum_{j=1}^{L_n} \#S_{n,j} = n$.

To understand its physical meaning, consider an example case where all samples in $\mathcal{S}_n$ contain a spin at a certain position in the coupling space $(A_z, A_\perp)$. Since the couplings are permutationally invariant, the same spin can have a different label in different samples from the posterior. When breaking the correlations to build $\tilde{\mathcal{S}}_n$, exactly $|\mathcal{S}_n|$ pairs $(A_z, A_\perp)$ refer to the same physical spin. If these are clustered together in $K_{n,j}$, because they are well-separated from all the other spins, a weight of $\#S_{n,j} = 1$ is assigned to the cluster according to the definition. This indicates a single spin at that position in the hyperfine couplings space. If, for example, a posterior distribution for a single spin produces samples that cluster with equal probability around two well-separated positions A and B, then each of these clusters will carry a weight $\#S = 0.5$, meaning that the posterior, given the observed data, predicts as equally likely configurations with a spin at position A or a spin at position B.

All intermediate configurations are possible, with an arbitrary number of spins, not exceeding K , distributed among any number of clusters.

E. Defining performance metrics

We benchmark the performance of the model selection algorithm by computing the standard machine learning (ML) metrics of precision, recall, F_1 score, and the mean absolute error. These are widely used in the ML literature and are based on counting the true positives (TPs), false positives (FPs), and false negatives (FNs). Broadly speaking, we look at clusters obtained in the postprocessing of the posterior, and if an appropriate number of spins in the ground truth is close to the cluster, we consider these spins to be counted among the true positives. If the algorithm signals a cluster, but no spins are to be found nearby, we classify this cluster as a false positive. Finally, if no cluster is in the vicinity of a spin in the ground truth, this spin is counted as one of the false negatives. The ML figures of

merit are defined as follows:

$$\begin{aligned} \text{Precision} &:= \frac{\text{TP}}{\text{TP} + \text{FP}}, \\ \text{Recall} &:= \frac{\text{TP}}{\text{TP} + \text{FN}}, \\ F_1 &:= \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}. \end{aligned} \quad (35)$$

Precision is the fraction of correctly identified spins over the total number of spins in the ground truth, while recall is the fraction of correctly identified spins over the total number of identified spins. We arbitrarily take the F_1 score as the harmonic mean of precision and recall: in general, they should be combined in a figure of merit that acknowledges the costs associated with a false positive and a false negative detection based on the specific application under consideration.

The figures of merit are computed as follows. We fix the number n of spins in the conditioned posterior, and for each cluster, indexed by j , we compute the Mahalanobis distance [152] from each of the spins, indexed by k , i.e.,

$$d_{j,k} := \sqrt{(\boldsymbol{\mu}_{n,j} - \mathbf{A}_k)^T \Sigma_{n,j}^{-1} (\boldsymbol{\mu}_{n,j} - \mathbf{A}_k)}. \quad (36)$$

From this distance, we compute the ML metrics as follows:

- (i) We identify all spins for which $d_{j,k} \leq t$, with $t \sim \mathcal{O}(1)$ being an arbitrarily chosen threshold value. These are flagged as true positives, i.e., they have been correctly identified in the cluster. We indicate the number of true positives in a cluster with TP_j . In this manuscript, we have chosen $t = 4$. Assuming that the cluster is Gaussian, this means a probability of 13.5% for a point belonging to the cluster to be flagged as a false negative.
- (ii) The weight of the cluster $\#S_{n,j}$ is rounded, and if the difference between $\lfloor \#S_{n,j} \rfloor$ and the number of true spins identified through the Mahalanobis distance is positive, these spins are considered false positives. The number of false positives in the cluster is

$$\text{FP}_j := \lfloor \#S_{n,j} \rfloor - \text{TP}_j. \quad (37)$$

- (iii) The total number of false positives is $\sum_j \text{FP}_j$, assuming $\text{FP}_j > 0$. The total number of true positives is the number of spins that have been flagged as such, and the number of false negatives is the number of true spins that have not been flagged.

A cluster can contain multiple spins, and we do not trigger any event in the case where more spins are close to the cluster than the weight of the cluster itself. Underestimating the number of spins in a cluster is not penalized, as

TABLE I. Range of τ values for different M used in simulation.

M	256	512	1024	2048	4096
$\tau_{\min}$ (μs)	6.0	6.0	6.0	6.0	6.0
$\tau_{\max}$ (μs)	7.25	8.5	11.0	16.0	26.0

long as a cluster is identified that is close to the spins. This is in line with the idea of using VBI for a fast exploration of the environment with the goal of obtaining a rough representation of the environment to be used in the adaptive step.

Mathematically, this difference can be seen in the fact that $\text{FP}_j > 0$ negatively impacts the precision, while $\text{FP}_j < 0$ does not contribute to the false negatives, and therefore to the recall, as long as there is a cluster in the vicinity of all the spins in the ground truth.

F. Simulation results

In this section, we report the results of simulations carried out for a magnetic field $B = 403$ G, a number $N_\pi = 32$ of π pulses in the DD sequence, $R = 10^3$, batch size $\mathcal{B} = 128$, and $I = 8192$ training steps. We impose L_2 regularization, with $\sigma = 10^{-2}$ and a relatively high value $A_\perp^{\text{th}} = 8$ kHz for the threshold on the orthogonal hyperfine components.

We simulate the signal generated by DD experiments on a single electron spin by environments comprising between 6 and 16 spins, while the ansatz for the posterior is always trained on $K = 20$ spins. Model selection consists of the identification of the exact number of spins generating the signal. The extra spins in the posterior, with couplings close to zero, are pruned by the thresholding procedure.

We have carried out 50 simulated estimations in total, i.e., one for each number of spins $6 \leq n \leq 16$ and for each number of measured points in the signal $M = 256, 512, 1024, 2048, 4096$. In Table I we report the windows used for the interpulse time delay in DD, corresponding to each number of measurement points M in the dataset. A data point is taken every 10 ns, which is the temporal resolution of the instruments.

In Figs. 5 and 6, we report the outcome of our model selection algorithms for, respectively, $M = 512$ and $M = 4096$ measurement steps. The measurement points are uniformly distributed in the interval $(\tau_{\min}, \tau_{\max})$. Subplot (a) compares the simulated signal (blue line) with the algorithm fit (red line), showing excellent match in all cases.

Subplot (b) shows the 2D marginalized posterior generated by the VBI procedure (blue), with the ground truth hyperfine values marked as red circles. We have selected the conditional probability corresponding to the most likely number of spins, according to p_n , as defined in Sec. III B 2. The probability density represented by the blue

regions indicates the uncertainty around each spin, as estimated by the algorithm. The red stripe marks the threshold on the orthogonal hyperfine coupling. Note that, in plotting only the marginal of the posterior, we are discarding the information about the correlations between the spins (Sec. IV D). This information can however be useful in the computation of the error for quantities depending on multiple spins, or in providing more accurate confidence levels to inform adaptive measurements.

Subplots (c) and (d) list, respectively, the ground truth hyperfine values for the $n = 15$ spins and the estimated values. As explained above, the number of spins in each cluster is not necessarily an integer, and it can be interpreted in a probabilistic sense as the average number of spins at that position according to the posterior.

The total measurement duration is computed as $T_{\text{tot}} = \sum_{j=1}^M R(2N_\pi \tau_j + T_{\text{oh}})$, with T_{oh} being an overhead time that takes into account 10 μs for the initialization and read-out of the NV center for each DD sequence, and 37 ns for each π -pulse [140]. We ignore the nondeterministic overhead introduced by NV charge state instabilities, spectral diffusion, and (optional) measurement-based initialization of the NV's own ^{14}N nuclear spin.

In Fig. 7 we report the ML metrics (recall, precision, and F_1 -score) obtained in simulation for multiple environments and multiple numbers of measurements.

Finally, we report the results for the average error (in kHz) for the simulated estimation of A_z and of $A_\perp$, in Tables II and III, respectively.


In both tables, the configuration of spins from which the simulated signal is created is representative of a distribution of spins with relatively large hyperfine couplings only. Real configurations typically include a large number of spins further away from the central electron spin, with small hyperfine couplings that produce a background signal. Such spins are not simulated here as they produce a weak background signal and cannot be accurately estimated. While we have neglected the spins further from the NV center here for simplicity, we will see in the next section how VBI-based model selection performs well also under realistic circumstances.

The training times for the algorithm, reported in Table IV, barely depends on the true number of spins in the environment, while exhibiting a stronger dependence on the number of spins in the ansatz and number of measurements. In Table IV we report also T_{tot} , i.e., the total time required for the data acquisition, which, by construction, does not depend on the number of spins in the environment.


G. Experimental results

We deploy our algorithm on experimental data from Abobeih *et al.* [140], consisting of dynamical decoupling measurements on a single NV in diamond with natural

(a) Simulated signal vs. predicted signal



(b) Marginal posterior



(c) Ground truth

A_z (kHz)	$A_\perp$ (kHz)
-50.07	57.93
-46.52	5.173
-40.06	31.67
-20.40	17.33
-8.387	16.00
-0.6844	30.53
-0.2706	28.25
-0.1114	35.89
2.467	22.63
10.35	24.80
10.77	2.085
15.18	14.47
24.75	6.334
34.79	23.76
37.35	23.44

(d) Estimation

A_z (kHz)	$A_\perp$ (kHz)	#S
-50.13	57.68	1.0000
-40.19	31.75	1.0003
-20.40	17.20	1.0001
-8.531	15.44	1.0050
-0.3661	29.41	1.6106
-0.1114	35.06	1.2322
0.2228	8.992	0.8939
1.958	24.13	1.1576
10.39	24.89	1.0000
15.07	14.21	1.0000
35.89	9.326	0.0491
36.10	23.62	2.0000
64.30	9.804	0.0512

FIG. 5. Example of application to a simulated bath of 15 spins, using $M = 512$ values of the interpulse delay τ in the DD measurement. (a) Dynamical decoupling signal [Eq. (21)] as a function of the interpulse delay τ with $\tau \in [6, 8.5] \mu\text{s}$. We compare the simulated data (blue curve) to the signal calculated from the VBI analysis output parameters (red curve). The uncertainty in the prediction is represented by the light red lines around the measurement signal, being approximately one standard deviation from the average prediction. (b) Marginal posterior distribution trained on the observed signal, conditioned on the presence of 13 spins (which has been found by the algorithm to be the most likely number of clusters). Three of the spins fall within the blind zone, represented by the red stripe, and cannot be individually detected based on the simulated data. The simulated data includes shot noise and additional measurement noise, as reported in Eqs. (26) and (27). The red dots represent the true values of the couplings, the green crosses are the centers of the clusters, and the blue shaded regions indicate the regions where the posterior probability concentrates. (c) List of the hyperfine values for the ground truth spins [used to generate the simulated DD signal in (a)]. (d) List of the hyperfine values for the identified spins, given as the centers of the clusters in the trained posterior, with the corresponding weights. The spins not used in the explanation of the signal (i.e., with couplings close to zero) are pushed into the blind zone by L_2 regularization and pruned by the thresholding. The confidence that the environment comprises 13 spins is $p_{13} = 0.4545$, the total signal acquisition time (T_{tot}) is 2 min.


1.1% abundance of ^{13}C isotopes, preselected only for the absence of ^{13}C spins with hyperfine coupling larger than 500 kHz. Data were collected at cryogenic temperature, exploiting single-shot electron spin readout through spin-selective resonant excitation, with average fidelity 95% [140].

The experimental signal has been processed with an ansatz for the posterior containing $K = 40$ spins, i.e., 80 spin parameters. The batch size was set to $\mathcal{B} = 256$, and


we used the L_1 regularization with $\sigma = 10^{-2}$, and $R = 500$. The threshold on the hyperfine coupling has been decreased to $A_\perp^{\text{th}} = 3 \text{ kHz}$ compared to the simulations in the previous section. The other parameters are identical to the simulated estimation discussed in Sec. IV F.

As illustrated in Fig. 8, we are able to successfully identify at least 8 spins from the experimental data. All the other identified spins above threshold, and below threshold, contribute to the fit of the signal. As ground truth

(a) Simulated signal vs. predicted signal



(b) Marginal posterior



(c) Ground truth

A_z (kHz)	$A_\perp$ (kHz)
-50.07	57.93
-46.52	5.173
-40.06	31.67
-20.40	17.33
-8.387	16.00
-0.6844	30.53
-0.2706	28.25
-0.1114	35.89
2.467	22.63
10.35	24.80
10.77	2.085
15.18	14.47
24.75	6.334
34.79	23.76
37.35	23.44

(d) Estimation

A_z (kHz)	$A_\perp$ (kHz)	#S
-50.07	57.77	1.000
-40.03	31.77	1.000
-20.42	17.22	1.000
-8.356	15.96	1.000
-0.5093	31.04	0.5072
-0.3661	28.19	1.374
-0.2865	34.76	1.129
2.435	22.60	0.9898
10.35	24.70	1.000
15.14	14.12	1.000
34.79	23.83	1.000
37.32	23.13	1.000

FIG. 6. Example of application to a simulated bath of 15 spins, using $M = 4096$ values of the interpulse delay τ in the DD measurement. This figure reports the results of the algorithm for the same situation as Fig. 5, only with more data, corresponding to $M = 4096$ values of τ . The legend for all the elements of the plot is the same as in Fig. 5. By utilizing more data, the algorithm identifies with great confidence the presence of 13 spins in the environment outside of the blind zone ($p_{13} = 90\%$). Here $T_{\text{tot}} = 36$ min.

values are not available in the experiment, we report with red dots the hyperfine couplings estimated in previous work by Jung *et al.* [79]. In Ref. [79] the authors identify 14 spins using only the same $N_\pi = 32$ DD dataset that we have processed, but extended to $\tau \in [6, 50]$ μs . The training time of the posterior for the experimental data is 8 min, to be compared with 8 h for the model training and 50 s of processing for each spin of the deep-learning approach.

The spins not used in the explanation of the signal (i.e., with couplings close to zero) are pushed into the blind zone by regularization and pruned by the thresholding. The confidence that the environment comprises 22 spins is $p_{22} = 0.2042$, the total signal acquisition time (T_{tot}) is 16 min.

An interesting question is how the performance of our VBI-based approach compares to MCMC, utilized for

example in Ref. [139] in fitting the experimental data. To assess this, we have implemented a simple Metropolis-Hastings MCMC for a fixed number of parameters and without parallel tempering (different than the reversible-jump implementation, with parallel tempering described in Ref [139]).

In Fig. 8(d), we report histograms of the log-likelihoods resulting from the applications of, respectively, VBI (left) and MCMC (right). VBI reaches higher log-likelihood values than our MCMC baseline, indicating a better fit within our model class. In addition, the estimation time reported in Table IV scales more favorably with the number of data points for VBI than for Metropolis-Hastings MCMC in our implementation, making VBI more suitable for rapidly updating a posterior in the context of online Bayesian experimental design.




FIG. 7. Machine learning metrics for the deployment of the algorithm on simulated data. We report the machine learning figures of merit for the model selection and estimation of hyperfine couplings in nano-NMR. Samples from the trained posterior are postprocessed to extract clusters from the identified spins, as described in Secs. III B 2 and IV D. The centers and covariance matrices of these clusters are used to identify the false positive and false negative spins, from which the recall, precision, and F₁ score are computed, as defined in Eq. (35). The values of all these figures of merit are in [0, 1], with one corresponding to maximum performance. Subplots (a), (b), and (c) are tables reporting recall, precision, and F₁ score respectively, as a function of the number of spins and number of measurements. Recall represents the impact of the false negatives, precision the impact of the false positives, and the F₁ score is the harmonic mean of precision and recall. We observe that the performance using $M = 4096$ measurements degrades compared to $M = 2048$; this happens because the posterior has not been trained to convergence. (d) Probability p_s of the selected number of spins, as defined in Sec. III B 2. This parameter is a proxy for the confidence level in the identification of the number of spins.

TABLE II. Average error per spin for the parallel component of the hyperfine coupling A_z , shown in kHz (high-precision values), as evaluated in simulation.

Num. spins	Number of measurements				
	256	512	1024	2048	4096
6	0.732	0.095	0.064	0.035	0.116
7	1.035	0.605	0.076	0.100	0.089
8	0.255	0.191	0.115	0.088	0.140
9	0.140	0.119	0.089	0.123	0.088
10	0.812	0.102	0.107	0.060	0.056
11	0.318	0.159	0.430	0.089	0.076
12	0.573	0.653	0.477	0.134	0.078
13	0.175	0.684	0.605	0.107	0.121
14	0.116	0.286	0.091	0.038	0.255
15	0.589	0.334	0.032	0.462	0.045
16	1.210	0.557	0.140	0.557	0.525

TABLE III. Average error per spin for the orthogonal component of the hyperfine coupling $A_{\perp}$, shown in kHz (high-precision values), as evaluated in simulation.

Num. spins	Number of measurements				
	256	512	1024	2048	4096
6	1.45	0.350	0.716	0.239	0.382
7	1.19	2.55	0.477	0.589	0.255
8	2.39	1.19	0.732	0.621	0.350
9	1.59	0.509	0.477	0.589	0.557
10	3.18	0.493	0.350	0.366	0.509
11	1.53	0.557	0.334	0.366	0.350
12	1.59	1.34	1.13	0.302	0.286
13	1.75	1.37	0.541	0.334	0.430
14	1.24	0.955	0.398	0.255	0.828
15	0.732	0.525	0.207	1.54	0.175
16	3.34	0.923	0.684	1.75	1.75

TABLE IV. Training time of the variational Bayesian inference ansatz (in min) for various numbers of measurements, together with the time T_{tot} required for acquiring the corresponding signal in the experiment (in min) and the training time for the Markov chain Monte Carlo. The training of the variational Bayesian inference is performed on a GPU, while the Markov chain Monte Carlo is done on a CPU, see Sec. IV H for more information on the hardware.

	Number of measurements				
	256	512	1024	2048	4096
T_{tot}	1	2	5	16	36
Training time VBI	5	5	6	8	12
Training time MCMC	15	22	35	108	144

While MCMC is able to fit a number of spins, it appears to not be necessarily converging to the right mode of the posterior. On the other hand, VBI is probably underestimating uncertainties, and therefore it remains closer to the maximum likelihood estimator than MCMC. VBI also appears to lead to heavy tails in the log-likelihood histograms, which are absent in the MCMC histograms. A more accurate and complete comparison between VBI and MCMC, also involving reversible jumps, is out of the scope of this manuscript and is left to future work.

A natural question is how our approach compares to methods based on deep learning [79], in terms of the number of spins successfully detected and the associated estimation errors in the limit of large datasets. We report a detailed performance analysis in Appendix B, showing that the architecture utilized in this work is unable to match the results of Ref. [79], even when increasing the number of layers in the normalizing flow. In particular, we observe that the training process slows down significantly as the ansatz complexity increases; numerical instabilities appear in the gradient descent already at 32 layers, accompanied by prohibitive memory requirements. We conjecture that these challenges originate from the highly peaked structure of the likelihood when large datasets are used.

A comprehensive investigation of alternative architectures to optimize variational Bayesian inference for individual nuclear spin identification in the large-data regime is beyond the scope of this manuscript. However, our observation that increasing the number of autoregressive layers improves the fit quality indicates that capturing the non-Gaussian structure of the posterior peaks is essential. This, in turn, suggests that greater expressive power is required from the ansatz, for example, through the use of a mixture of flows or a diffusion-based posterior representation.

H. Methods

The algorithm has been coded using Pyro [153,154], a universal probabilistic programming language (PPL) [155]

written in Python and supported by PyTorch on the backend. All simulations have been run on an NVIDIA A40 GPU with 48 GB of dedicated VRAM, while the CPU simulation have been executed on an Apple M1. The training of the posterior ansatz has been executed with the ADAM optimizer and an exponentially decaying learning rate schedule, interpolating between an initial and final learning rates that were respectively set to $\ell(I = 1) = 10^{-3}$ and $(I = 8196) = 10^{-5}$.


V. DISCUSSION AND OUTLOOK

This work presents a framework for parameter estimation and model selection for quantum experiments, based on variational Bayesian inference. We show that VBI enables the characterization of quantum systems with a number of parameters significantly beyond what is currently feasible with techniques such as particle filtering [89]. We expect this framework to be broadly applicable across many different physical platforms, to high-dimensional estimation problems related to quantum tomography [41,45], Hamiltonian learning [26,31,89,156,157], the characterization of NISQ quantum circuits [158], multiparameter quantum sensing [159], and others.


In addition to parameter estimation for high-dimensional systems, our framework includes the capability to select models within a class of nested models featuring different numbers of unknown parameters. This is a topical challenge of great practical importance in quantum science and technology, as it captures the physics of a known quantum system interacting with an environment where not all relevant degrees of freedom might be known. Examples of this are a spin qubit coupled to an environment comprising an unknown number of other electronic and nuclear spins [14,127,140,144], a superconducting qubit coupling to an unknown number of two-level fluctuators [114,115], the explainable identification of the processes governing the dynamics of an unknown quantum system [32,49], and generally multiqubit systems featuring an unknown number of coherent coupling and decoherence processes (but limited within a known library of processes).

We have benchmarked our approach on a specific problem of high technological interest, i.e., the identification of individual nuclear spins from measurements on a single electron spin quantum sensor. This problem is sufficiently complex, in terms of number of free parameters and nonlinearity of the signal, to showcase the power of our framework in tackling the characterization of multi-dimensional quantum systems with unknown numbers of components. We have demonstrated, on simulated data and on real experimental data, how we can infer the hyperfine values for the spins that contribute to the observed signal, even when their number is not known in advance. Our protocol, parallelized on a GPU, can identify up to 16 spins in a few minutes, a timescale much shorter than that required

(a) Measured signal vs. predicted signal



(b) Marginal posterior



(c) Estimation

A_z (kHz)	$A_{\perp}$ (kHz)	$\#S$	A_z (kHz)	$A_{\perp}$ (kHz)	$\#S$
-76.87	5.045	1.0504	0.9708	17.73	1.0242
-48.40	10.27	1.1260	2.897	15.20	1.9814
-36.03	26.40	1.0000	8.324	21.09	1.0000
-28.85	4.775	0.2882	14.50	8.945	1.0821
-21.53	10.92	2.1320	18.30	8.053	1.0754
-13.72	14.59	1.0019	20.91	40.95	1.0000
-11.76	4.631	0.9572	24.72	22.41	1.0000
-11.43	56.58	1.0000	38.56	4.600	0.3245
-3.756	10.89	1.0590	53.09	4.520	0.3665
-2.976	4.536	0.7719	68.82	3.883	0.4880
-1.241	14.67	1.9698	76.95	4.122	0.3014

(d) Log-likelihood comparison




FIG. 8. Benchmarking the algorithm on experimental data. (a) Dynamical decoupling signal [Eq. (21)] as a function of the interpulse delay τ for a subset of the data from Abobeih *et al.* [140], with τ roughly between 10 and 18 μs ($M = 2048$ measurement points with 4 ns between each data point). The blue curve represents the experimental data, and the red curve shows the fit obtained using the VBI algorithm. The prediction uncertainty is indicated by the light red lines around the prediction, corresponding approximately to one standard deviation. (b) Marginal posterior distribution trained on the observed signal, conditioned on the most likely number of spins, identified by the algorithm as $n = 22$. A total of $K = 40$ spins are fit to model the weakly interacting spin bath and better reproduce the signal. The red dots represent the spins estimated by the deep-learning approach of Ref. [79] using the $N_{\pi} = 32$ and $N_{\pi} = 256$ complete datasets, which required approximately 4 h of data acquisition. (c) List of the hyperfine values and corresponding weights for the identified spins, divided into two tables according to the sign of A_z . Spins highlighted in red indicate false positives. Around the origin, more spins are not identified precisely. See the caption of Fig. 5 for interpretation of the marginal posterior. (d) Comparison of the log-likelihood distributions obtained with VBI and MCMC. We consider Metropolis-Hastings MCMC with an optimized Gaussian proposal, similar to Ref. [99]. We use a single chain, extracting 10^4 samples and fixing the number of spins to 50 (no reversible jump is implemented here).

for other Bayesian techniques such as Markov chain Monte Carlo.

While we have used nuclear spin detection as an example of the application of the technique, a direction of future work is to push the identification of nuclear spins to the limit of the largest number of spins and the smallest uncertainty on their hyperfine couplings and actual physical positions. This requires deploying the algorithm on larger datasets, with more values of the interpulse delay τ , as here we have only used a portion of the data in Aboeih *et al.* [140]. For sufficiently large values of τ , the dipolar coupling between nearby nuclear spins becomes important, requiring updating the Hamiltonian in Eq. (19) with the corresponding terms. The resulting Hamiltonian has analytical solutions for pairs of nuclear spins [125,130,140,160], while numerical solutions can be found with the cluster expansion method for larger spin clusters [161–165]. Furthermore, the learning will have to be integrated with different types of measurements, such as the DD-RF sequence [77] interleaving dynamical decoupling with phase-controlled driving of nuclear spins, and correlated double-resonance sequences to map connected chains through the network of spins [76]. To give an idea of the power of these techniques, and the complexity of the sensing and data processing procedures, these have recently enabled mapping a 50-nuclear-spin network comprising 1225 spin-spin interactions in the vicinity of an NV center [76]. This is an important application as the measurement time is much longer than for the dynamical decoupling datasets presented here, leading to potentially big advantages. Additionally, not all couplings are required to resolve the structure, and adaptive protocols can therefore be extremely advantageous in deciding which ones to characterize based on information from the previous measurements.

Our framework is purely based on classical algorithms and, as with any classical algorithm processing quantum systems [17–19], complexity scales exponentially with size. It is therefore most useful for a number of parameters which is too large for exact Bayesian inference but still sufficiently small for an efficient classical simulation of the probability outcomes that are required for applying VBI. Our techniques are thus best suited to problems where the quantum system of interest is simple enough to permit efficient evaluation of its dynamics, but where the number of parameters characterizing it is sufficiently large to constitute a high-dimensional parameter estimation problem. The example described in Sec. IV of this manuscript fits in this sweet spot.

It is also important to note that, while we train the full multidimensional posterior, for the clustering analysis that gives us an estimator for the couplings we do not leverage any of the correlations between the spins, discarding some information captured by the posterior for simplicity. In an adaptive scenario, the full posterior would be used to

decide the next optimal measurements to perform, including correlations between the parameters and confidence levels.

Our parameter learning and model selection algorithm produces an approximate posterior for all the models and parameters for the quantum system, given a certain set of measured data. The algorithm exploits parallelization, leading to short computation times on the order of a few minutes with modern GPUs (see Table IV). This fast-computed approximate posterior can be used to optimally choose subsequent measurements and probe state design in order to maximize the final precision of the estimation, in optimal Bayesian experimental design [166]. There are different possibilities to develop heuristics for the optimization of subsequent measurements, ranging from information-theoretic quantities such as Fisher information or information gain, to deploying reinforcement learning [84,85,167]. While information-theoretic quantities are typically expensive to compute, our framework can be extended to approximate them. For example, one could use the surrogate information gain:

$$\text{SIG}(\mathbf{x}) := \mathbb{E}_{\mathbf{y}} \left\{ \text{Var}_{\boldsymbol{\theta} \sim q_{\lambda}(\boldsymbol{\theta})} [p(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta})] \right\}, \quad (38)$$

which is easy to compute, since the variance is computed on the distribution $q_{\lambda}(\boldsymbol{\theta})$, and $p(\mathbf{y} | \mathbf{x}, \boldsymbol{\theta})$ is the probability for the observation of the outcomes $\mathbf{y}$ given the controls $\mathbf{x}$. We could then sample the next measurements among those that feature a large SIG, realizing thereby optimal experimental design based on VBI [168].

The posterior training times (Table IV) highlighted for the nanoscale magnetic resonance example in Sec. IV are suited for integration of adaptivity, with the goal of greatly reducing data acquisition time. As a reminder of the importance of smart data acquisition for this application, the total sensing time for a full dynamical decoupling trace with interpulse delay τ ranging from 10 to 300 μs and 4 ns discretization [140] is more than 8 days ($N_{\pi} = 32$, $R = 10^3$ repetitions). This is the bare sensing time, neglecting non-deterministic overheads due to the NV charge and spectral instabilities, the initialization of the NV's $^{14/15}\text{N}$, nuclear spin, and any setup-related issues such as periodic optical realignment, calibration, or drifts.

We envision adaptivity to be deployed by grouping data acquisition in batches. Training the posterior for a batch of 1024 measurements takes about 6 min (Table IV). Data acquisition for 1024 values of the interpulse delay τ takes about 9 min in the range [6, 10] μs and more than one hour for $\tau \in [60, 64]$ μs . The goal of adaptivity will be to utilize “cheap” information gained at shorter τ or smaller number N_{π} of pulses to prioritize measurements with the largest information content for larger values of τ and N_{π} , where data acquisition is more time-consuming. Posterior training durations comparable to the shortest data acquisition batches offer a good investment for the expected large

gains in avoiding long measurements with small information content. Our preliminary analysis [169] suggests that the amount of DD data collected can be greatly reduced without sacrificing estimation performance.

The data acquisition timescales described here refer to experiments performed at low temperature on NV centers deeper inside the diamond [140], where fast high-fidelity single-shot electron spin readout is available [150]. Room-temperature experiments and/or NV centers close to the diamond surface often use averaged readout; while Bayesian inference is still feasible and advantageous [151], the lack of single-shot readout leads to a significant slowdown of data acquisition, which makes the posterior training time completely negligible.

The operation of any technology requires dealing with imperfections, such as fabrication flaws, instabilities, and the impact of uncontrolled environmental factors. All these factors are often complicated to model and predict yet need to be included in the Bayesian inference process; otherwise, they will lead to incorrect estimation and suboptimal (if not detrimental) feedback control loops. A promising avenue for future exploration is to integrate VBI with “gray-box” techniques [170–172] which couple a physics-based model for the system with a data-driven model for the imperfections, which can be dependent on external parameters such as temperature, electric/magnetic fields, etc.

Another interesting open question is how to account for symmetries in the problem [173–180], such as the permutational symmetries of spin labeling for the example described in Sec. IV. Assuming that the prior model presents the same symmetries, the exact posterior $P(\theta | \mathbf{x}, \mathbf{y})$, calculated from Bayes’ rule, inherits the symmetries of the model $p(\theta | \mathbf{x}, \mathbf{y})$. However, the variational ansatz for the posterior distribution does not necessarily preserve these symmetries. For discrete symmetries, the posterior distribution exhibits multiple equivalent modes in the parameter space. In this case, the reverse KL-divergence used in variational inference [Eq. (5)] is known to manifest a mode-seeking behavior, which theoretically leads to the posterior concentrating on a single mode. In practice, during training, the posterior probability density tends to concentrate in only a few of the multiple equivalent regions of the parameter space, neglecting the others. An interesting future direction is to explore the use of symmetry-preserving ansatzes, with the idea that by stopping the ansatz from “leaking” into other modes, its expressivity will be better exploited in the approximation of the posterior.

ACKNOWLEDGMENTS

We thank D. Yudilevich, B. Varona-Uriarte and L. Sarra for useful discussions. We gratefully acknowledge the computational resources of the “High Performance

Computing” facility at Heriot-Watt University. This work is funded by the Engineering and Physical Sciences Research Council (EP/S000550/1, EP/V053779/1, EP/Z533208/1, EP/Z533191/1) and the European Innovation Council (QuSPARC, Grant agreement 101186889). This work is also supported by the project 23NRM04 NoQTeS, which has received funding from the European Partnership on Metrology, cofinanced from the European Union’s Horizon Europe Research and Innovation Programme and by the Participating States. Finally, this project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 852410).

C.B., F.B., and Y.A. conceived the project. F.B. developed the algorithm (with input by Y.A. and C.B.), carried out the simulations and analyzed the results (in collaboration with C.B., Y.A., and E.M.G.). M.H.A. and T.H.T. provided the experimental data. F.B. and C.B. led the writing of the manuscript, with contributions from all co-authors.

CODE AVAILABILITY

The open source library BrightCarbon can be found in the <https://gitlab.com/federico.belliardo/brightcarbonsoftware> repository and can be installed with pip following the instructions in the README file.

DATA AVAILABILITY

The data that support the findings of this article are openly available [181].

APPENDIX A: INTUITIVE INTERPRETATION OF THE EXPERIMENT

In this Appendix, we discuss intuitively the properties of the Hamiltonian in Eq. (19). We perform an estimation with N_π fixed and $\tau \in (\tau_{\min}, \tau_{\max})$ uniformly distributed. In the limit of large magnetic field (i.e., $f_L \gg \{A_{z,k}\}$), the probability distribution in Eq. (21) is characterized by multiple resonances for each spin, at values of the interpulse delay τ [144] given by Eq. (23). Each resonance is Lorentzian, with width $\mathcal{O}(A_{\perp,k}/f_L^2)$ and depth $\mathcal{O}(N_\pi A_{\perp,k}/f_L)$.

From the requirement that the distance between the location of two dips is greater than the width of the resonances we can define a *Rayleigh criterion* for being able to individually identify multiple spins in the environment. This tells us that we need to look at least at the resonances of order

$$m_{\min} := \frac{\max A_{\perp}}{4 \min(\Delta A_z)}, \quad (\text{A1})$$

to distinguish every spin. Here, $\min(\Delta A_z)$ indicates the minimum distance between the hyperfine couplings of

each pair of spins in the detection volume. For realistic numbers (e.g., $\max A_\perp = 80$ kHz), resolving hyperfines with $\min \Delta A_z = 0.1$ kHz requires taking data for resonances up to $m_{\min} = \mathcal{O}(10^2)$. This is a fairly time-consuming data acquisition, particularly as larger values of m result in a larger interpulse delay τ and, consequently, longer data acquisition. A possible strategy to address this is to first take data on resonances with lower m , e.g., $m < \mathcal{O}(10)$, and then use adaptive Bayesian experimental design to select the specific interpulse delays τ corresponding to resonances in the $m_{\min} = \mathcal{O}(10^2)$ range that give the most information in distinguishing the required spectral shifts, based on the information accumulated. The development of optimal adaptive schemes for nuclear spin identification is however beyond the scope of this manuscript, and will be the subject of future dedicated work.

In order to detect all the resonances, the discretization in the interpulse delay ($\delta\tau$) should be of the order of the width of the resonance, i.e., $A_\perp/f_L$. The m -th resonances fit within a window of size $1/4f_L$, giving a number of discrete interpulse delay values of

$$M_{\min} := \frac{1}{4f_L} \left(\frac{\min A_\perp}{f_L^2} \right)^{-1} = \mathcal{O} \left(\frac{f_L}{4 \min A_\perp} \right). \quad (\text{A2})$$

For $\min A_\perp = 8$ kHz and $f_L = 0.43$ MHz we get $M_{\min} \simeq 85$, which is in the ball park of the number of different τ needed to start distinguishing the spins individually in our simulations. Additionally, for a dip to be detectable, the shot-noise must be less than its depth. This sets a condition on the number of repetitions R for each DD sequence, which reads

$$\sqrt{\frac{p(1-p)}{R}} \leq \frac{N_\pi \min A_\perp}{f_L}, \quad (\text{A3})$$

where

$$p := 1 - \frac{N_\pi \min A_\perp}{f_L} \quad (\text{A4})$$

and this equation is effectively just noise $<$ signal. For small dips the condition on $R_{\min}$ is

$$R \geq R_{\min} := \frac{f_L}{N_\pi \min A_\perp}, \quad (\text{A5})$$

Combining all the previous considerations, the condition on the total measurement time for being able to distinguish all the spins is

$$T_{\min} = \mathcal{O} \left(\frac{f_L \max A_\perp}{\min \Delta A_z (\min A_\perp)^2} \right). \quad (\text{A6})$$

For the numerical values reported above we get $T_{\min} \simeq 3$ s, which is at least two orders of magnitude smaller than

the total measurement time in our simulations. This could be an indication that we could actually use far less time in nano-NMR experiments than currently done [169], or maybe the prefactor of $T_{\min}$ is not of order one.

The discrepancy in the estimated time could be solved by assuming that there are additional sources of noise on the system, and therefore we need to increase $R_{\min}$, which for our range of parameters is only ~ 4 , far from the used value of $R \simeq \mathcal{O}(10^3)$. Since we have assumed additional sources of noise in the analysis of this manuscript, the above analysis might indicate that the algorithm we propose works close to the theoretical limits in the individual identification of the spins.

To summarize, we have considered the large-field limit of the task of determining the hyperfine couplings for independent nuclear spins in the environment of a single electron spin, and we have established lower bounds on the time and measurements needed in the learning. In the limit $\omega_L \rightarrow \infty$, the task is equivalent to the problem of searching and identifying the resonances by scanning over an interval of times τ s. This is similar to the setting of Allen *et al.* [182], suggesting that the estimation could be sped up by quantum algorithms, to achieve the Heisenberg-Grover limit.

APPENDIX B: EXTENDING THE DATASET AND THE EXPRESSIVITY OF THE ANSATZ

In this Appendix, we explore the performance of our architecture in the limit of a large dataset. We fit an extended dataset corresponding to $\tau \in [6, 50] \mu\text{s}$ for $N_\pi = 32$ and $\tau \in [18, 28] \mu\text{s}$ for $N_\pi = 256$. We first examine the fit for the $N_\pi = 32$ dataset alone (see Fig. 9), which contains approximately 10^4 points. We then jointly fit the data for $N_\pi = 32$ together with the $N_\pi = 256$ data (for a total of $\sim 1.25 \times 10^4$ points). Our results show that, as the number of data points increases, the complexity of the variational ansatz becomes increasingly critical in determining the quality of the fit.

As in the main text, we employ a normalizing-flow ansatz constructed from either 5 or 16 autoregressive layers. In its base configuration, this architecture utilizes deep neural networks with 8 hidden layers, each consisting of 256 units.

In Fig. 9, we compare the performance obtained on the full $N_\pi = 32$ dataset using two architectures: one with 5 autoregressive layers and another with 16. We also report the training loss [the negative ELBO from Eq. (28)]. The reduced posterior indicates that the larger ansatz yields a substantially better fit to the data, a finding further supported by the lower asymptotic loss achieved during training. To ensure a fair comparison, all remaining hyperparameters were held fixed across training runs.




FIG. 9. Result of the fit for the dataset with $N_\pi = 32$ for $\tau \in [6, 50]$ μs . (a) Restricted posterior obtained using a normalizing flow ansatz with 5 autoregressive layers. (b) Restricted posterior obtained with 16 autoregressive layers. In both posteriors, the results of the clustering algorithm are superimposed; green crosses indicate the cluster centers, with the number of spins in each cluster labeled accordingly. (c) Loss function, specifically the negative evidence lower-bound from Eq. (28), as a function of the training step. The 16-layer ansatz converges more rapidly to the asymptote and achieves a lower loss, despite requiring more computational time per training step.

In Fig. 10, we present the results of the combined fit using both the $N_\pi = 32$ and $N_\pi = 256$ datasets. For this more challenging setting, we employed 16 autoregressive layers and increased the network depth from 8 to 12 hidden layers. Although the simpler ansatz fails completely, the more expressive architecture also produces unsatisfactory results, despite a substantial increase in training time from 47 min to 67 min.

We conjecture that increasing the ansatz complexity further may be necessary to obtain a satisfactory fit for datasets containing multiple pulse sequences; however, pursuing this direction lies beyond the scope of the present work. In particular, a more computationally efficient architecture is likely required.

In all reduced posteriors presented in Figs. 9 and 8, we include 47 spins as a reference, following the identifications reported in Ref. [79]. We exclude a single spin with an abnormally large A_z coupling that lies outside our

analyzed window. The authors of Ref. [79] note that some redundancies may be present in their analysis, where a single spin is occasionally resolved as a pair of closely spaced spins, potentially reducing the effective number to 31. For the $N_\pi = 32$ dataset, only 14 of these spins have been confirmed through independent analyses in subsequent literature. By inspecting Figs. 8 and 9, and by relaxing the identification criteria to include clusters that are not perfectly centered on known spin locations, it is reasonable to argue that as many as 12 spins can be resolved directly from the posterior, with additional spins originating from the surrounding bath.

When incorporating the $N_\pi = 256$ dataset, a significantly larger number of weakly coupled spins is identified, although only 7 of them have been independently confirmed. In total, 21 spins identified by the deep-learning method of Ref. [79] have received external confirmation. Finally, we observe in Fig. 10 that the 16-layer ansatz fails




FIG. 10. Reduced posteriors and training loss for the joint fit of datasets with $N_{\pi} = 32$ and $N_{\pi} = 256$ pulses. The description provided for Fig. 9 also applies here. Notably, while the final training loss values for the simple and complex architectures are comparable, the fit using fewer layers fails entirely to capture the characteristic features of the environment.

to resolve the spins with the strongest couplings, instead fitting a large number of spins with comparatively smaller coupling strengths.

[1] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, High-threshold and low-overhead fault-tolerant quantum memory, *Nature* **627**, 778 (2024).
[2] Google Quantum AI and Collaborators *et al.*, Quantum error correction below the surface code threshold, *Nature* **638**, 920 (2025).
[3] N. Maskara, S. Ostermann, J. Shee, M. Kalinowski, A. McClain Gomez, R. Araiza Bravo, D. S. Wang, A. I. Krylov, N. Y. Yao, M. Head-Gordon, M. D. Lukin, and S. F. Yelin, Programmable simulations of molecules and materials with reconfigurable quantum processors, *Nat. Phys.* **21**, 289 (2025).
[4] D. Bluvstein *et al.*, A fault-tolerant neutral-atom architecture for universal quantum computation, *Nature* **649**, 39 (2026).

[5] N.-C. Chiu, E. C. Trapp, J. Guo, M. H. Abobeih, L. M. Stewart, S. Hollerith, P. L. Stroganov, M. Kalinowski, A. A. Geim, S. J. Evered, S. H. Li, X. Lyu, L. M. Peters, D. Bluvstein, T. T. Wang, M. Greiner, V. Vuletić, and M. D. Lukin, Continuous operation of a coherent 3,000-qubit system, *Nature* **646**, 1075 (2025).
[6] P.-J. Stas, Y. Q. Huan, B. Machielse, E. N. Knall, A. Suleymanzade, B. Pingault, M. Sutula, S. W. Ding, C. M. Knaut, D. R. Assumpcao, Y.-C. Wei, M. K. Bhaskar, R. Riedinger, D. D. Sukachev, H. Park, M. Lončar, D. S. Levonian, and M. D. Lukin, Robust multi-qubit quantum network node with integrated error detection, *Science* **378**, 557 (2022).
[7] S. L. N. Hermans, M. Pompili, H. K. C. Beukers, S. Baier, J. Borregaard, and R. Hanson, Qubit teleportation between non-neighbouring nodes in a quantum network, *Nature* **605**, 663 (2022).
[8] C. Delle Donne *et al.*, An operating system for executing applications on quantum network nodes, *Nature* **639**, 321 (2025).
[9] A. J. Stolk *et al.*, Metropolitan-scale heralded entanglement of solid-state qubits, *Sci. Adv.* **10**, eadp6442 (2024).

- [10] C. M. Knaut, A. Suleymanzade, Y.-C. Wei, D. R. Assumpcao, P.-J. Stas, Y. Q. Huan, B. Machielse, E. N. Knall, M. Sutula, G. Baranes, N. Sinclair, C. De-Eknamkul, D. S. Levonian, M. K. Bhaskar, H. Park, M. Lončar, and M. D. Lukin, Entanglement of nanophotonic quantum memory nodes in a telecom network, *Nature* **629**, 573 (2024).
- [11] E. Boto, N. Holmes, J. Leggett, G. Roberts, V. Shah, S. S. Meyer, L. D. Muñoz, K. J. Mullinger, T. M. Tierney, S. Bestmann, G. R. Barnes, R. Bowtell, and M. J. Brookes, Moving magnetoencephalography towards real-world applications with a wearable system, *Nature* **555**, 657 (2018).
- [12] B. Stray *et al.*, Quantum sensing for gravity cartography, *Nature* **602**, 590 (2022).
- [13] N. Aslam, H. Zhou, E. K. Urbach, M. J. Turner, R. L. Walsworth, M. D. Lukin, and H. Park, Quantum sensors for biomedical applications, *Nat. Rev. Phys.* **5**, 157 (2023).
- [14] R. Budakian *et al.*, Roadmap on nanoscale magnetic resonance imaging, *Nanotechnology* **35**, 412001 (2024).
- [15] K. D. Briegel, N. R. Von Grafenstein, J. C. Draeger, P. Blümler, R. D. Allert, and D. B. Bucher, Optical widefield nuclear magnetic resonance microscopy, *Nat. Commun.* **16**, 1281 (2025).
- [16] N. Mohseni, T. Fösel, L. Guo, C. Navarrete-Benlloch, and F. Marquardt, Deep learning of quantum many-body dynamics via random driving, *Quantum* **6**, 714 (2022).
- [17] V. Gebhart, R. Santagati, A. A. Gentile, E. M. Gauger, D. Craig, N. Ares, L. Bianchi, F. Marquardt, L. Pezzè, and C. Bonato, Learning quantum systems, *Nat. Rev. Phys.* **5**, 141 (2023).
- [18] M. Krenn, J. Landgraf, T. Foesel, and F. Marquardt, Artificial intelligence and machine learning for quantum technologies, *Phys. Rev. A* **107**, 010101 (2023).
- [19] A. Dawid, *Machine Learning in Quantum Sciences* (Cambridge University, Cambridge, 2025), 1st ed.
- [20] G. Acampora *et al.*, Quantum computing and artificial intelligence: Status and perspectives, [arXiv:2505.23860](https://arxiv.org/abs/2505.23860) [quant-ph].
- [21] Y. Alexeev *et al.*, Artificial intelligence for quantum computing, *Nat. Commun.* **16**, 10829 (2025).
- [22] Y. Du, Y. Zhu, Y.-H. Zhang, M.-H. Hsieh, P. Rebentrost, W. Gao, Y.-D. Wu, J. Eisert, G. Chiribella, D. Tao, and B. C. Sanders, Artificial intelligence for representing and characterizing quantum systems, [arXiv:2509.04923](https://arxiv.org/abs/2509.04923).
- [23] R. Blume-Kohout, Optimal, reliable estimation of quantum states, *New J. Phys.* **12**, 043034 (2010).
- [24] M. Paris and J. Řeháček, (Eds.), *Quantum State Estimation*, Lecture Notes in Physics (Springer, Berlin, Heidelberg, 2004), Vol. 649.
- [25] N. R. von Grafenstein, K. D. Briegel, J. Casanova, and D. B. Bucher, Coherent signal detection in the statistical polarization regime enables high-resolution nanoscale NMR spectroscopy, [arXiv:2501.02093](https://arxiv.org/abs/2501.02093) [physics].
- [26] N. Wiebe, C. Granade, C. Ferrie, and D. Cory, Quantum Hamiltonian learning using imperfect quantum resources, *Phys. Rev. A* **89**, 042314 (2014).
- [27] J. Wang, S. Paesani, R. Santagati, S. Knauer, A. A. Gentile, N. Wiebe, M. Petruzzella, J. L. O'Brien, J. G. Rarity, A. Laing, and M. G. Thompson, Experimental quantum Hamiltonian learning, *Nat. Phys.* **13**, 551 (2017).
- [28] A. Anshu, S. Arunachalam, T. Kuwahara, and M. Soleimanifar, Sample-efficient learning of interacting quantum systems, *Nat. Phys.* **17**, 931 (2021).
- [29] M. Hayashi, *Asymptotic Theory of Quantum Statistical Inference* (World Scientific, Singapore, 2005).
- [30] A. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory* (Edizioni Della Normale, Pisa, 2011).
- [31] A. A. Gentile, B. Flynn, S. Knauer, N. Wiebe, S. Paesani, C. E. Granade, J. G. Rarity, R. Santagati, and A. Laing, Learning models of quantum systems from experiments, *Nat. Phys.* **17**, 837 (2021).
- [32] S. Wallace, Y. Altmann, B. D. Gerardot, E. M. Gauger, and C. Bonato, Learning the dynamics of Markovian open quantum systems from experimental data, *Phys. Rev. Res.* **7**, 033217 (2025).
- [33] J. M. Lukens, K. J. H. Law, A. Jasra, and P. Lougovski, A practical and efficient approach for Bayesian quantum state estimation, *New J. Phys.* **22**, 063038 (2020).
- [34] K. Jones, Principles of quantum inference, *Ann. Phys.* **207**, 140 (1991).
- [35] V. Bužek, R. Derka, G. Adam, and P. Knight, Reconstruction of quantum states of spin systems: From quantum Bayesian inference to quantum tomography, *Ann. Phys.* **266**, 454 (1998).
- [36] R. Schack, T. A. Brun, and C. M. Caves, Quantum Bayes rule, *Phys. Rev. A* **64**, 014305 (2001).
- [37] K. S. Kravtsov, S. S. Straupe, I. V. Radchenko, N. M. T. Houlsby, F. Huszár, and S. P. Kulik, Experimental adaptive Bayesian tomography, *Phys. Rev. A* **87**, 062122 (2013).
- [38] O. Di Matteo, J. Gamble, C. Granade, K. Rudinger, and N. Wiebe, Operational, Gauge-free quantum tomography, *Quantum* **4**, 364 (2020).
- [39] M. Christandl and R. Renner, Reliable quantum state tomography, *Phys. Rev. Lett.* **109**, 120403 (2012).
- [40] C. Granade, J. Combes, and D. G. Cory, Practical Bayesian tomography, *New J. Phys.* **18**, 033024 (2016).
- [41] F. Huszár and N. M. T. Houlsby, Adaptive Bayesian quantum tomography, *Phys. Rev. A* **85**, 052120 (2012).
- [42] C. Granade, C. Ferrie, and S. T. Flammia, Practical adaptive quantum tomography, *New J. Phys.* **19**, 113017 (2017).
- [43] S. Mondal and A. K. Dutta, A Bayesian quantum state tomography along with adaptive frameworks based on linear minimum mean square error criterion, *New J. Phys.* **25**, 123001 (2023).
- [44] N. Wiebe, C. Granade, C. Ferrie, and D. G. Cory, Hamiltonian learning and certification using quantum resources, *Phys. Rev. Lett.* **112**, 190501 (2014).
- [45] N. Wiebe and C. Granade, Efficient Bayesian phase estimation, *Phys. Rev. Lett.* **117**, 010503 (2016).
- [46] M. D. Shulman, S. P. Harvey, J. M. Nichol, S. D. Bartlett, A. C. Doherty, V. Umansky, and A. Yacoby, Suppressing qubit dephasing using real-time Hamiltonian estimation, *Nat. Commun.* **5**, 5156 (2014).
- [47] V. Gebhart, A. Smerzi, and L. Pezzè, Bayesian quantum multiphase estimation algorithm, *Phys. Rev. Appl.* **16**, 014035 (2021).
- [48] F. Berritta, J. Benestad, L. Pahl, M. Mathews, J. A. Krzywdą, R. Assouly, Y. Sung, D. K. Kim, B. M. Niedzielski, K. Serniak, M. E. Schwartz, J. L. Yoder, A. Chatterjee,

- J. A. Grover, J. Danon, W. D. Oliver, and F. Kuemmeth, Efficient qubit calibration by binary-search Hamiltonian tracking, *PRX Quantum* **6**, 030335 (2025).
- [49] L. Fioroni, I. Rojko, and F. Reiter, A learning agent-based approach to the characterization of open quantum systems, *Phys. Rev. Appl.* **24**, 034011 (2025).
- [50] C. Bonato, M. S. Blok, H. T. Dinani, D. W. Berry, M. L. Markham, D. J. Twitchen, and R. Hanson, Optimized quantum sensing with a single electron spin using real-time adaptive measurements, *Nat. Nanotechnol.* **11**, 247 (2016).
- [51] D. T. Lennon, H. Moon, L. C. Camenzind, L. Yu, D. M. Zumbühl, G. A. D. Briggs, M. A. Osborne, E. A. Laird, and N. Ares, Efficiently measuring a quantum device using machine learning, *npj Quantum Inf.* **5**, 79 (2019).
- [52] R. Santagati, A. A. Gentile, S. Knauer, S. Schmitt, S. Paesani, C. Granade, N. Wiebe, C. Osterkamp, L. P. McGuinness, J. Wang, M. G. Thompson, J. G. Rarity, F. Jelezko, and A. Laing, Magnetic-field learning using a single electronic spin in diamond with one-photon readout at room temperature, *Phys. Rev. X* **9**, 021019 (2019).
- [53] M. Valeri, E. Polino, D. Poderini, I. Gianani, G. Corrielli, A. Crespi, R. Osellame, N. Spagnolo, and F. Sciarrino, Experimental adaptive Bayesian estimation of multiple phases with limited data, *npj Quantum Inf.* **6**, 92 (2020).
- [54] S. Dushenko, K. Ambal, and R. D. McMichael, Sequential Bayesian experiment design for optically detected magnetic resonance of nitrogen-vacancy centers, *Phys. Rev. Appl.* **14**, 054036 (2020).
- [55] R. Kaubruegger, D. V. Vasilyev, M. Schulte, K. Hammerer, and P. Zoller, Quantum variational optimization of Ramsey interferometry and atomic clocks, *Phys. Rev. X* **11**, 041045 (2021).
- [56] T. Joas, S. Schmitt, R. Santagati, A. A. Gentile, C. Bonato, A. Laing, L. P. McGuinness, and F. Jelezko, Online adaptive quantum characterization of a nuclear spin, *npj Quantum Inf.* **7**, 56 (2021).
- [57] K. Craigie, E. M. Gauger, Y. Altmann, and C. Bonato, Resource-efficient adaptive Bayesian tracking of magnetic fields with a quantum sensor, *J. Phys.: Condens. Matter* **33**, 195801 (2021).
- [58] S. E. D'Aurelio, M. Valeri, E. Polino, V. Cimini, I. Gianani, M. Barbieri, G. Corrielli, A. Crespi, R. Osellame, F. Sciarrino, and N. Spagnolo, Experimental investigation of Bayesian bounds in multiparameter estimation, *Quantum Sci. Technol.* **7**, 025011 (2022).
- [59] M. Caouette-Mansour, A. Solyom, B. Ruffolo, R. D. McMichael, J. Sankey, and L. Childress, Robust spin relaxometry with fast adaptive Bayesian estimation, *Phys. Rev. Appl.* **17**, 064031 (2022).
- [60] I. Zohar, B. Haylock, Y. Romach, M. J. Arshad, N. Halay, N. Drucker, R. Stöhr, A. Denisenko, Y. Cohen, C. Bonato, and A. Finkler, Real-time frequency estimation of a qubit without single-shot-readout, *Quantum Sci. Technol.* **8**, 035017 (2023).
- [61] M. J. Arshad, C. Bekker, B. Haylock, K. Skrzypczak, D. White, B. Griffiths, J. Gore, G. W. Morley, P. Salter, J. Smith, I. Zohar, A. Finkler, Y. Altmann, E. M. Gauger, and C. Bonato, Real-time adaptive estimation of decoherence timescales for a single qubit, *Phys. Rev. Appl.* **21**, 024026 (2024).
- [62] P. Cappellaro, Spin-bath narrowing with adaptive parameter estimation, *Phys. Rev. A* **85**, 030301 (2012).
- [63] E. Scerri, E. M. Gauger, and C. Bonato, Extending qubit coherence by adaptive quantum environment learning, *New J. Phys.* **22**, 035002 (2020).
- [64] F. Berritta, T. Rasmussen, J. A. Krzywda, J. Van Der Heijden, F. Fedele, S. Fallahi, G. C. Gardner, M. J. Manfra, E. Van Nieuwenburg, J. Danon, A. Chatterjee, and F. Kuemmeth, Real-time two-axis control of a spin qubit, *Nat. Commun.* **15**, 1676 (2024).
- [65] R. Bellman, *Dynamic Programming* (Princeton University, Princeton, NJ, 1984).
- [66] T. Bengtsson, P. Bickel, and B. Li, in *Institute of Mathematical Statistics Collections* (Institute of Mathematical Statistics, Beachwood, Ohio, USA, 2008), pp. 316–334.
- [67] H.-H. Lu, K. V. Mylswamy, R. S. Bennink, S. Seshadri, M. S. Alshaykh, J. Liu, T. J. Kippenberg, D. E. Leaird, A. M. Weiner, and J. M. Lukens, Bayesian tomography of high-dimensional on-chip biphoton frequency combs with randomized measurements, *Nat. Commun.* **13**, 4338 (2022).
- [68] J. C. Chapman, J. M. Lukens, B. Qi, R. C. Pooser, and N. A. Peters, Bayesian homodyne and heterodyne tomography, *Opt. Express* **30**, 15184 (2022).
- [69] M. Valeri, V. Cimini, S. Piacentini, F. Ceccarelli, E. Polino, F. Hoch, G. Bizzarri, G. Corrielli, N. Spagnolo, R. Osellame, and F. Sciarrino, Experimental multiparameter quantum metrology in adaptive regime, *Phys. Rev. Res.* **5**, 013138 (2023).
- [70] F. Belliardo, V. Cimini, E. Polino, F. Hoch, B. Piccirillo, N. Spagnolo, V. Giovannetti, and F. Sciarrino, Optimizing quantum-enhanced Bayesian multiparameter estimation of phase and noise in practical sensors, *Phys. Rev. Res.* **6**, 023201 (2024).
- [71] We stress that this exponential growth is not related to the exponential complexity in time and memory of simulating and certifying quantum systems, which we are not trying to address in this paper.
- [72] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, Variational inference: A review for statisticians, *J. Am. Stat. Assoc.* **112**, 859 (2017).
- [73] P. P. Mazza, D. Zietlow, F. Carollo, S. Andergassen, G. Martius, and I. Lesanovsky, Machine learning time-local generators of open quantum dynamics, *Phys. Rev. Res.* **3**, 023084 (2021).
- [74] G. Cemin, F. Carnazza, S. Andergassen, G. Martius, F. Carollo, and I. Lesanovsky, Inferring interpretable dynamical generators of local quantum observables from projective measurements through machine learning, *Phys. Rev. Appl.* **21**, L041001 (2024).
- [75] L. Schorling, P. Vaidhyathan, J. Schuff, M. J. Carballido, D. Zumbühl, G. Milburn, F. Marquardt, J. Foerster, M. A. Osborne, and N. Ares, Meta-learning characteristics and dynamics of quantum systems, *arXiv:2503.10492* [quant-ph].
- [76] G. L. Van De Stolpe, D. P. Kwiatkowski, C. E. Bradley, J. Randall, M. H. Abobeih, A. Breitweiser, L. C. Bassett, M. Markham, D. J. Twitchen, and T. H. Taminiau, Mapping a 50-spin-qubit network through correlated sensing, *Nat. Commun.* **15**, 2006 (2024).

- [77] C. E. Bradley, J. Randall, M. H. Abobeih, R. C. Berrevoets, M. J. Degen, M. A. Bakker, M. Markham, D. J. Twitchen, and T. H. Taminiau, A ten-qubit solid-state spin register with quantum memory up to one minute, *Phys. Rev. X* **9**, 031045 (2019).
- [78] J. Randall, C. E. Bradley, F. V. Van Der Gronden, A. Galicia, M. H. Abobeih, M. Markham, D. J. Twitchen, F. Machado, N. Y. Yao, and T. H. Taminiau, Many-body-localized discrete time crystal with a programmable spin-based quantum simulator, *Science* **374**, 1474 (2021).
- [79] K. Jung, M. H. Abobeih, J. Yun, G. Kim, H. Oh, A. Henry, T. H. Taminiau, and D. Kim, Deep learning enhanced individual nuclear-spin detection, *npj Quantum Inf.* **7**, 41 (2021).
- [80] B. Varona-Uriarte, C. Munuera-Javaloy, E. Terradillos, Y. Ban, A. Alvarez-Gila, E. Garrote, and J. Casanova, Automatic detection of nuclear spins at arbitrary magnetic fields via signal-to-image AI model, *Phys. Rev. Lett.* **132**, 150801 (2024).
- [81] C.-W. Huang, D. Krueger, A. Lacoste, and A. Courville, Neural autoregressive flows, in *Proceedings of the 35th International Conference on Machine Learning*, edited by J. Dy and A. Krause (PMLR, 2018), Vol. 80, pp. 2078–2087.
- [82] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University, 2012).
- [83] A. S. Holevo, *Quantum Systems, Channels, Information: A Mathematical Introduction* (De Gruyter, 2019).
- [84] F. Belliardo, F. Zoratti, F. Marquardt, and V. Giovannetti, Model-aware reinforcement learning for high-performance Bayesian experimental design in quantum metrology, *Quantum* **8**, 1555 (2024).
- [85] F. Belliardo, F. Zoratti, and V. Giovannetti, Applications of model-aware reinforcement learning in Bayesian quantum metrology, *Phys. Rev. A* **109**, 062609 (2024).
- [86] N. Gordon, D. Salmond, and A. Smith, Novel approach to nonlinear/non-gaussian Bayesian state estimation, *IEE Proc. F (Radar Signal Process.)* **140**, 107 (1993).
- [87] A. Doucet, N. Freitas, and N. Gordon, (Eds.), *Sequential Monte Carlo Methods in Practice* (Springer New York, New York, NY, 2001).
- [88] P. Del Moral, A. Doucet, and A. Jasra, Sequential Monte Carlo samplers, *J. R. Stat. Soc. Ser. B: Stat. Methodol.* **68**, 411 (2006).
- [89] C. E. Granade, C. Ferrie, N. Wiebe, and D. G. Cory, Robust online Hamiltonian learning, *New J. Phys.* **14**, 103013 (2012).
- [90] G. Kitagawa, Monte Carlo filter and smoother for Non-Gaussian nonlinear state space models, *J. Comput. Graph. Stat.* **5**, 1 (1996).
- [91] J. S. Liu and R. Chen, Sequential Monte Carlo methods for dynamic systems, *J. Am. Stat. Assoc.* **93**, 1032 (1998).
- [92] J. Liu and M. West, in *Sequential Monte Carlo Methods in Practice* (Springer New York, New York, NY, 2001), pp. 197–223.
- [93] O. Cappe, S. J. Godsill, and E. Moulines, An overview of existing methods and recent advances in sequential Monte Carlo, *Proc. IEEE* **95**, 899 (2007).
- [94] T. Li, M. Bolic, and P. M. Djuric, Resampling methods for particle filtering: Classification, implementation, and strategies, *IEEE Signal Process. Mag.* **32**, 70 (2015).
- [95] L. M. Murray, A. Lee, and P. E. Jacob, Parallel resampling in the particle filter, *J. Comput. Graph. Stat.* **25**, 789 (2016).
- [96] A. E. Gelfand and A. F. M. Smith, Sampling-based approaches to calculating marginal densities, *J. Am. Stat. Assoc.* **85**, 398 (1990).
- [97] S. Brooks, A. Gelman, G. Jones, and X.-L. Meng, *Handbook of Markov Chain Monte Carlo* (Chapman and Hall/CRC, New York, 2011), 1st ed.
- [98] G. L. Jones and Q. Qin, Markov chain Monte Carlo in practice, *Annu. Rev. Stat. Appl.* **9**, 557 (2022).
- [99] A. N. Potesman, M. Onizhuk, C. Egerstrom, D. P. Mark, D. D. Awschalom, F. J. Heremans, and G. Galli, High-throughput spin-bath characterization of spin-defects in semiconductors, *Phys. Rev. Appl.* **24**, 054048 (2025).
- [100] J. Skilling, in *AIP Conference Proceedings* (AIP, Garching (Germany), 2004), Vol. 735, pp. 395–405.
- [101] J. Skilling, Nested sampling for general Bayesian computation, *Bayesian Anal.* **1**, 833 (2006).
- [102] G. Ashton *et al.*, Nested sampling for physical scientists, *Nat. Rev. Methods Primers* **2**, 39 (2022).
- [103] C. Peterson and J. Anderson, A mean field theory learning algorithm for neural networks, *Complex Syst.* **1**, 995 (1987).
- [104] G. E. Hinton and D. Van Camp, in *Proceedings of the Sixth Annual Conference on Computational Learning Theory—COLT '93* (ACM, Santa Cruz, California, USA, 1993), pp. 5–13.
- [105] R. M. Neal and G. E. Hinton, in *Learning in Graphical Models* (Springer Netherlands, Dordrecht, 1998), pp. 355–368.
- [106] L. K. Saul, T. Jaakkola, and M. I. Jordan, Mean field theory for sigmoid belief networks, *J. Artif. Intell. Res.* **4**, 61 (1996).
- [107] T. S. Jaakkola and M. I. Jordan, A variational approach to Bayesian logistic regression models and their extensions, in *Proceedings of the Sixth International Workshop on Artificial Intelligence and Statistics*, edited by D. Madigan and P. Smyth (PMLR, 1997), Vol. R1, pp. 283–294..
- [108] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings* (2014).
- [109] D. Rezende and S. Mohamed, Variational inference with normalizing flows, in *Proceedings of the 32nd International Conference on Machine Learning*, edited by F. Bach and D. Blei (PMLR, Lille, France, 2015).
- [110] F. Belliardo and V. Giovannetti, Achieving Heisenberg scaling with maximally entangled states: An analytic upper bound for the attainable root-mean-square error, *Phys. Rev. A* **102**, 042613 (2020).
- [111] R. M. Neal, Slice sampling, *Ann. Stat.* **31**, 705 (2003).
- [112] E. Higson, W. Handley, M. Hobson, and A. Lasenby, Dynamic nested sampling: An improved algorithm for parameter estimation and evidence calculation, *Stat. Comput.* **29**, 891 (2019).
- [113] H. Haario, E. Saksman, and J. Tamminen, An adaptive metropolis algorithm, *Bernoulli* **7**, 223 (2001).
- [114] J. Burnett, L. Faoro, I. Wisby, V. L. Gurtovoi, A. V. Chernykh, G. M. Mikhailov, V. A. Tulin, R.

- Shaikhaidarov, V. Antonov, P. J. Meeson, A. Ya. Tzalenchuk, and T. Lindström, Evidence for interacting two-level systems from the $1/f$ noise of a superconducting resonator, *Nat. Commun.* **5**, 4119 (2014).
- [115] S. E. De Graaf, L. Faoro, L. B. Ioffe, S. Mahashabde, J. J. Burnett, T. Lindström, S. E. Kubatkin, A. V. Danilov, and A. Ya. Tzalenchuk, Two-level systems in superconducting quantum devices due to trapped quasiparticles, *Sci. Adv.* **6**, eabc5055 (2020).
- [116] M. Lucas, A. V. Danilov, L. V. Levitin, A. Jayaraman, A. J. Casey, L. Faoro, A. Ya. Tzalenchuk, S. E. Kubatkin, J. Saunders, and S. E. De Graaf, Quantum bath suppression in a superconducting circuit by immersion cooling, *Nat. Commun.* **14**, 3522 (2023).
- [117] B.-J. Liu, Y.-Y. Wang, T. Sheffer, and C. Wang, Observation of discrete charge states of a coherent two-level system in a superconducting qubit, *Phys. Rev. Lett.* **133**, 160602 (2024).
- [118] A. Gelman, A. Jakulin, M. G. Pittau, and Y.-S. Su, A weakly informative default prior distribution for logistic and other regression models, *Ann. Appl. Stat.* **2**, 1360 (2008).
- [119] R. Tibshirani, Regression shrinkage and selection via the Lasso, *J. R. Stat. Soc.: Ser. B (Methodol.)* **58**, 267 (1996).
- [120] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Quantum state tomography via compressed sensing, *Phys. Rev. Lett.* **105**, 150401 (2010).
- [121] N. A. Nemkov, E. O. Kiktenko, I. A. Luchnikov, and A. K. Fedorov, Efficient variational synthesis of quantum circuits with coherent multi-start optimization, *Quantum* **7**, 993 (2023).
- [122] R. Kress, *Numerical Analysis* (Springer Science & Business Media, 2012).
- [123] J. C. Marcks, B. Pingault, J. Zhang, C. Zeledon, F. J. Hermans, and D. D. Awschalom, Nuclear spin engineering for quantum information science, *J. Mater. Res.* **40**, 1433 (2025).
- [124] T. Staudacher, F. Shi, S. Pezzagna, J. Meijer, J. Du, C. A. Meriles, F. Reinhard, and J. Wrachtrup, Nuclear magnetic resonance spectroscopy on a (5-nanometer)³ sample volume, *Science* **339**, 561 (2013).
- [125] F. Shi, X. Kong, P. Wang, F. Kong, N. Zhao, R.-B. Liu, and J. Du, Sensing and atomic-scale structure analysis of single nuclear-spin clusters in diamond, *Nat. Phys.* **10**, 21 (2014).
- [126] J. Zopes, K. S. Cujia, K. Sasaki, J. M. Boss, K. M. Itoh, and C. L. Degen, Three-dimensional localization spectroscopy of individual nuclear spins with sub-Angstrom resolution, *Nat. Commun.* **9**, 4678 (2018).
- [127] J. Zopes, K. Herb, K. S. Cujia, and C. L. Degen, Three-dimensional nuclear spin positioning using coherent radio-frequency control, *Phys. Rev. Lett.* **121**, 170801 (2018).
- [128] K. S. Cujia, J. M. Boss, K. Herb, J. Zopes, and C. L. Degen, Tracking the precession of single nuclear spins by weak measurements, *Nature* **571**, 230 (2019).
- [129] K. S. Cujia, K. Herb, J. Zopes, J. M. Abendroth, and C. L. Degen, Parallel detection and spatial mapping of large nuclear spin clusters, *Nat. Commun.* **13**, 1260 (2022).
- [130] K. Herb, T. F. Segawa, L. A. Völker, J. M. Abendroth, E. Janitz, T. Zhu, and C. L. Degen, Multidimensional spectroscopy of nuclear spin clusters in diamond, *Phys. Rev. Lett.* **132**, 133002 (2024).
- [131] J. Du, F. Shi, X. Kong, F. Jelezko, and J. Wrachtrup, Single-molecule scale magnetic resonance spectroscopy using quantum diamond sensors, *Rev. Mod. Phys.* **96**, 025001 (2024).
- [132] N. Zhao, J. Honert, B. Schmid, M. Klas, J. Isoya, M. Markham, D. Twitchen, F. Jelezko, R.-B. Liu, H. Fedder, and J. Wrachtrup, Sensing single remote nuclear spins, *Nat. Nanotechnol.* **7**, 657 (2012).
- [133] G. Wolfowicz, P.-A. Mortemousque, R. Guichard, S. Simmons, M. L. W. Thewalt, K. M. Itoh, and J. J. L. Morton, ²⁹Si nuclear spins as a resource for donor spin qubits in silicon, *New J. Phys.* **18**, 023021 (2016).
- [134] N. Kalb, A. A. Reiserer, P. C. Humphreys, J. J. W. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, Entanglement distillation between solid-state quantum network nodes, *Science* **356**, 928 (2017).
- [135] S. Asaad, V. Mourik, B. Joecker, M. A. I. Johnson, A. D. Baczewski, H. R. Firgau, M. T. Mądzik, V. Schmitt, J. J. Pla, F. E. Hudson, K. M. Itoh, J. C. McCallum, A. S. Dzurak, A. Laucht, and A. Morello, Coherent electrical control of a single high-spin nucleus in silicon, *Nature* **579**, 205 (2020).
- [136] M. H. Abobeih, Y. Wang, J. Randall, S. J. H. Loenen, C. E. Bradley, M. Markham, D. J. Twitchen, B. M. Terhal, and T. H. Taminiau, Fault-tolerant operation of a logical qubit in a diamond quantum processor, *Nature* **606**, 884 (2022).
- [137] C. E. Bradley, S. W. De Bone, P. F. W. Möller, S. Baier, M. J. Degen, S. J. H. Loenen, H. P. Bartling, M. Markham, D. J. Twitchen, R. Hanson, D. Elkouss, and T. H. Taminiau, Robust quantum-network memory based on spin qubits in isotopically engineered diamond, *npj Quantum Inf.* **8**, 122 (2022).
- [138] J. Reiner, Y. Chung, S. H. Misha, C. Lehner, C. Moehle, D. Poulos, S. Monir, K. J. Charde, P. Macha, L. Kranz, I. Thorvaldson, B. Thorgriemsson, D. Keith, Y. L. Hsueh, R. Rahman, S. K. Gorman, J. G. Keizer, and M. Y. Simmons, High-fidelity initialization and control of electron and nuclear spins in a four-qubit register, *Nat. Nanotechnol.* **19**, 605 (2024).
- [139] A. N. Poteshman, J. Yun, T. H. Taminiau, and G. Galli, Trans-dimensional Hamiltonian model selection and parameter estimation from sparse, noisy data, *Quantum* **10**, 2055 (2026).
- [140] M. H. Abobeih, J. Cramer, M. A. Bakker, N. Kalb, M. Markham, D. J. Twitchen, and T. H. Taminiau, One-second coherence for a single electron spin coupled to a multi-qubit nuclear-spin environment, *Nat. Commun.* **9**, 2552 (2018).
- [141] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. Hollenberg, The nitrogen-vacancy colour centre in diamond, *Phys. Rep.* **528**, 1 (2013).
- [142] R. Schirhagl, K. Chang, M. Loretz, and C. L. Degen, Nitrogen-vacancy centers in diamond: Nanoscale sensors for physics and biology, *Annu. Rev. Phys. Chem.* **65**, 83 (2014).

- [143] J. F. Barry, J. M. Schloss, E. Bauch, M. J. Turner, C. A. Hart, L. M. Pham, and R. L. Walsworth, Sensitivity optimization for NV-diamond magnetometry, *Rev. Mod. Phys.* **92**, 015004 (2020).
- [144] T. H. Taminiau, J. J. T. Wagenaar, T. Van Der Sar, F. Jelezko, V. V. Dobrovitski, and R. Hanson, Detection and control of individual nuclear spins using a weakly coupled electron spin, *Phys. Rev. Lett.* **109**, 137602 (2012).
- [145] M. J. Biercuk, A. C. Doherty, and H. Uys, Dynamical decoupling sequence construction as a filter-design problem, *J. Phys. B: At. Mol. Opt. Phys.* **44**, 154002 (2011).
- [146] C. L. Degen, F. Reinhard, and P. Cappellaro, Quantum sensing, *Rev. Mod. Phys.* **89**, 035002 (2017).
- [147] A. Laraoui, F. Dolde, C. Burk, F. Reinhard, J. Wrachtrup, and C. A. Meriles, High-resolution correlation spectroscopy of ^{13}C spins near a nitrogen-vacancy centre in diamond, *Nat. Commun.* **4**, 1651 (2013).
- [148] J. Casanova, Z.-Y. Wang, J. F. Haase, and M. B. Plenio, Robust dynamical decoupling sequences for individual nuclear-spin addressing, *Phys. Rev. A* **92**, 042304 (2015).
- [149] M. Zahedian, V. Vorobyov, and J. Wrachtrup, Blueprint for efficient nuclear spin characterization with color centers, *Phys. Rev. B* **109**, 214111 (2024).
- [150] L. Robledo, L. Childress, H. Bernien, B. Hensen, P. F. A. Alkemade, and R. Hanson, High-fidelity projective readout of a solid-state spin quantum register, *Nature* **477**, 574 (2011).
- [151] H. T. Dinani, D. W. Berry, R. Gonzalez, J. R. Maze, and C. Bonato, Bayesian estimation for quantum sensing in the absence of single-shot detection, *Phys. Rev. B* **99**, 125413 (2019).
- [152] P. C. Mahalanobis, On the generalized distance in statistics, *Proc. Natl. Inst. Sci. India* **2**, 49 (1936).
- [153] E. Bingham, J. P. Chen, M. Jankowiak, F. Obermeyer, N. Pradhan, T. Karaletsos, R. Singh, P. Szerlip, P. Horsfall, and N. D. Goodman, Pyro: Deep universal probabilistic programming, *JMLR* **20**, (2019).
- [154] D. Phan, N. Pradhan, and M. Jankowiak, Composable effects for flexible and accelerated probabilistic programming in NumPyro, *arXiv:1912.11554* [stat].
- [155] G. Baudart, M. Hirzel, and L. Mandel, Deep probabilistic programming languages: A qualitative study, *CoRR*, *arXiv:1804.06458* [cs].
- [156] A. Valenti, E. Van Nieuwenburg, S. Huber, and E. Greplova, Hamiltonian learning for quantum error correction, *Phys. Rev. Res.* **1**, 033092 (2019).
- [157] A. Valenti, G. Jin, J. Léonard, S. D. Huber, and E. Greplova, Scalable Hamiltonian learning for large-scale out-of-equilibrium quantum dynamics, *Phys. Rev. A* **105**, 023302 (2022).
- [158] S. Duffield, M. Benedetti, and M. Rosenkranz, Bayesian learning of parameterised quantum circuits, *Mach. Learn.: Sci. Technol.* **4**, 025007 (2023).
- [159] L. Pezzè and A. Smerzi, Advances in multiparameter quantum sensing and metrology, *arXiv:2502.17396*.
- [160] N. Zhao, J.-L. Hu, S.-W. Ho, J. T. K. Wan, and R. B. Liu, Atomic-scale magnetometry of distant nuclear spin clusters via nitrogen-vacancy spin in diamond, *Nat. Nanotechnol.* **6**, 242 (2011).
- [161] J. R. Maze, J. M. Taylor, and M. D. Lukin, Electron spin decoherence of single nitrogen-vacancy defects in diamond, *Phys. Rev. B* **78**, 094303 (2008).
- [162] W. Yang and R.-B. Liu, Quantum many-body theory of qubit decoherence in a finite-size spin bath, *Phys. Rev. B* **78**, 085315 (2008).
- [163] W. Yang and R.-B. Liu, Quantum many-body theory of qubit decoherence in a finite-size spin bath. II. Ensemble dynamics, *Phys. Rev. B* **79**, 115320 (2009).
- [164] W. Yang, W.-L. Ma, and R.-B. Liu, Quantum many-body theory for electron spin decoherence in nanoscale nuclear spin baths, *Rep. Prog. Phys.* **80**, 016001 (2017).
- [165] M. Onizhuk and G. Galli, PyCCE: A python package for cluster correlation expansion simulations of spin qubit dynamics, *Adv. Theor. Simul.* **4**, 2100254 (2021).
- [166] T. Rainforth, A. Foster, D. R. Ivanova, and F. B. Smith, Modern Bayesian experimental design, *Stat. Sci.* **39**, 100 (2024).
- [167] K. Reuer, J. Landgraf, T. Fösel, J. O'Sullivan, L. Beltrán, A. Akin, G. J. Norris, A. Remm, M. Kerschbaum, J.-C. Besse, F. Marquardt, A. Wallraff, and C. Eichler, Realizing a deep reinforcement learning agent for real-time quantum feedback, *Nat. Commun.* **14**, 7138 (2023).
- [168] L. Sarra and F. Marquardt, Deep Bayesian experimental design for quantum many-body systems, *Mach. Learn.: Sci. Technol.* **4**, 045022 (2023).
- [169] B. Varona-Uriarte, F. Belliardo, T. H. Taminiau, C. Bonato, E. Garrote, and J. Casanova, Computationally tractable offline quantum experimental design for nuclear spin detection, *arXiv:2508.21450* [quant-ph].
- [170] T. Bohlin, *Practical Grey-Box Process Identification*, *Advances in Industrial Control* (Springer, London, 2006).
- [171] É. Genois, J. A. Gross, A. Di Paolo, N. J. Stevenson, G. Koolstra, A. Hashim, I. Siddiqi, and A. Blais, Quantum-tailored machine-learning characterization of a superconducting qubit, *PRX Quantum* **2**, 040355 (2021).
- [172] A. Youssry, Y. Yang, R. J. Chapman, B. Haylock, F. Lenzini, M. Lobino, and A. Peruzzo, Experimental gray-box quantum system identification and control, *npj Quantum Inf.* **10**, 9 (2024).
- [173] M. Biloš and S. Günnemann, Scalable normalizing flows for permutation invariant densities, in *Proceedings of the 38th International Conference on Machine Learning*, edited by M. Meila and T. Zhang (PMLR, 2021), pp. 957–967.
- [174] K. Rasul, I. Schuster, R. Vollgraf, and U. Bergmann, Set flow: A permutation invariant normalizing flow, *arXiv:1909.02775* [cs].
- [175] B. Zwartsenberg, A. Scibior, M. Niedoba, V. Lioutas, Y. Liu, J. Sefas, S. Dabiri, J. W. Lavington, T. Campbell, and F. Wood, Conditional permutation invariant flows, *Trans. Machine Learn. Res.* (2023).
- [176] H. Maron, O. Litany, G. Chechik, and E. Fetaya, On learning sets of symmetric elements, edited by Hal Daumé III and A. Singh (PMLR, 2020), Vol. 119, pp. 6734–6744.
- [177] E. H. Thiede, T. S. Hy, and R. Kondor, The general theory of permutation equivariant neural networks and higher order graph variational encoders, *arXiv:2004.03990* [cs].

- [178] A. Zhou, K. Yang, K. Burns, A. Cardace, Y. Jiang, S. Sokota, J. Z. Kolter, and C. Finn, Permutation equivariant neural functionals, *Advances in Neural Information Processing Systems*, edited by A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (Curran Associates, Inc., 2023), pp. 24966–24992.
- [179] B. Tolooshams, A. Chandrashekar, R. Zirvi, A. Mammadov, J. Yao, C. Wang, and A. Anandkumar, EquiReg: Equivariance regularized diffusion for inverse problems, [arXiv:2505.22973](https://arxiv.org/abs/2505.22973) [cs].
- [180] J. Köhler, L. Klein, and F. Noe, Equivariant Flows: Exact Likelihood Generative Learning for Symmetric Densities, in *Proceedings of the 37th International Conference on Machine Learning* (PMLR, 2020), pp. 5361–5370.
- [181] <https://doi.org/10.5281/zenodo.20189740>.
- [182] R. R. Allen, F. Machado, I. L. Chuang, H.-Y. Huang, and S. Choi, Quantum computing enhanced sensing, [arXiv:2501.07625](https://arxiv.org/abs/2501.07625) [quant-ph].