Optimal Quantum Control with Poor Statistics

Frédéric Sauvage^{✉*} and Florian Mintert[✉]

Physics Department, Blackett Laboratory, Imperial College London, Prince Consort Road, London SW7 2BW, United Kingdom

(Received 24 September 2019; revised 1 July 2020; accepted 13 November 2020; published 17 December 2020)

Control of quantum systems is a central element of high-precision experiments and the development of quantum technological applications. Control pulses that are typically temporally or spatially modulated are often designed based on theoretical simulations. As we gain control over larger and more complex quantum systems, however, we reach the limitations of our capabilities of theoretical modeling and simulations, and learning how to control a quantum system based exclusively on experimental data can help us to exceed those limitations. Because of the intrinsic probabilistic nature of quantum mechanics, it is fundamentally necessary to repeat measurements on individual quantum systems many times in order to estimate the expectation value of an observable with good accuracy. Control algorithms requiring accurate data can thus imply an experimental effort that negates the benefits of avoiding theoretical modeling. We present a control algorithm that finds optimal control solutions in the presence of large measurement shot noise and even in the limit of single-shot measurements. The algorithm builds up on Bayesian optimization that is well suited to handle noisy data; but since the commonly used assumption of Gaussian noise is not appropriate for projective measurements with a low number of repetitions, we develop Bayesian inference that respects the binomial nature of shot noise. With several numerical and experimental examples we demonstrate that this method is capable of finding excellent control solutions with minimal experimental effort.

DOI: [10.1103/PRXQuantum.1.020322](https://doi.org/10.1103/PRXQuantum.1.020322)

I. INTRODUCTION

Quantum control allows us to perform highly accurate experiments and to turn synthetic quantum systems into devices ranging from the most precise available clocks [1] to qubit registers [2–4] on which several elementary quantum gates can be executed. Optimization of the typically time-dependent control fields is traditionally performed based on a theoretical model. Given the increasingly complex quantum systems that we are gaining control over, it often becomes infeasible to find a sufficiently accurate model, let alone to simulate its dynamics. The design of control fields based on experimental observations, but without resort to theoretical modeling [5] has thus become more and more important.

While it is generally important that control algorithms converge quickly to a good solution, optimal control based on experimental data also requires that such solutions can be found using a small amount of data and data with

limited accuracy. This is particularly problematic with modern experiments on quantum systems since the vast majority are performed on individual systems, rather than ensembles as was the case in the early days of quantum mechanics. Since obtaining expectation values of an elementary observable with good accuracy requires a large number of repetitions of the same experiment, there is a trade-off in the acquisition of data. More repetitions yield data with higher accuracy, which is clearly helpful in the search for a good control pulse. The repetitions, however, also imply additional effort, which may even make a search prohibitively expensive.

Our present goal is to devise a control algorithm that does not require accurate estimation of expectation values, and that can find good solutions of control problems even in the extreme case of poor statistics due to intrinsic shot noise.

Among the most popular methods to perform optimizations directly based on experimental outcomes, exact or approximate gradient-based routines [6–10] and Nelder-Mead [11–15] have the advantage of simplicity and fast convergence when the underlying optimization landscape is well behaved. Their performance can however be limited in the presence of many local minima [16,17] or plateaus [18]. Evolutionary algorithms [16,19,20] and the recently introduced reinforcement learning techniques [17,

*frederic.sauvage15@imperial.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.

[21–24] may overcome this problem at the expense of a large number of iterations that require a large set of data.

Bayesian optimization [25–28], on the other hand, has been shown to offer a good compromise between the ability to find high-quality solutions quickly, based on limited, potentially noisy data, and reliability in distinguishing local from global maxima [29–34]. Given its data efficiency and its underlying probabilistic modeling approach, we deem Bayesian optimization an excellent starting point for our present purpose. Despite these appealing aspects, however, Bayesian optimization in its usually employed form is not yet optimally adapted to work with data that is affected by measurement noise. In this paper we thus develop a framework based on Bayesian optimization that takes into account this specific type of noise that is characteristic of quantum systems. With several examples, we show that this variant of Bayesian optimization converges much faster to optimal solutions than conventional Bayesian optimization, making it ideally suited for situations in which experimental data is a scarce resource, for example, due to low repetition rates.

In Sec. II we review pertinent elements of Bayesian optimization and develop the detailed modeling of measurement noise that will enable optimal control based on extremely noisy data. Specifically, Secs. II A, II C, and II E are devoted to an introduction to Bayesian optimization, while the explicit modeling of measurement noise is discussed in Sec. II B, followed by a discussion of the formulation of control targets in Sec. II D. The explicit evaluation of predictions of Bayesian inference in the present framework is detailed in Sec. II F. The discussion in Sec. III of the performance of this framework with a pedagogical example of a single qubit, the preparation of a Greenberger-Horne-Zeilinger (GHZ) state, the realization of ordered states of ultra-cold gases and the state preparation in a noisy, intermediate-scale quantum (NISQ) device might be appreciated also by readers who decided to skip details of the discussion in Sec. II.

II. OPTIMAL CONTROL ALGORITHM

Optimal control [35–37] is applicable in any situation in which a system can be manipulated with external control fields, as realized for example by a laser- or microwave pulse. The central task is the identification of the control pulse that achieves a desired goal as accurately as possible. A figure of merit that characterizes how well such a goal has been achieved is often called a target functional. Typical examples for such figures of merit include state fidelities [38], gate fidelities [13,39,40], the expectation values of an entanglement witness, and also more non-linear quantities like entanglement measures [21,38,41,42] and the Fisher information [14,43].

In order to turn an optimization task into a practical problem, it is typically required to characterize control

pulses in terms of a finite number, i.e., a vector θ of tuneable control parameters. These are often given in terms of piecewise-constant amplitudes of external electric or magnetic fields [6–10,44], but also temporal shapes of laser- or microwave pulses with parametrization in terms of Fourier series [11,12,38,45–47], or rotation angles of quantum gates [48,49] are objects of optimization.

In control based on theoretical modeling, the control parameters are tuneable parameters in the Hamiltonian or quantum gate, but in control based on experimental data the control parameters would typically be inputs to an experimental device. Crucially, it is not necessary to know how these inputs relate to the Hamiltonian describing the system, and even systems whose microscopic workings are not understood are accessible to this type of control.

The dependence of the figure of merit F on the control parameters θ defines the control landscape, and the task at hand is identifying the value of θ for which $F(\theta)$ adopts its maximum. In most practical problems, however, one is often satisfied with the identification of a value of θ that results in a sufficiently large value of $F(\theta)$.

It is hardly ever possible to find the dependence of the figure of merit F on the control parameters θ analytically. In addition to the necessity of an underlying theoretical model, this would require constructing the solution of the Schrödinger equation as an analytic function of the control parameters, which is only feasible in exceptionally elementary situations. Typically, one thus probes the value of $F(\theta)$ only for a discrete set of data points, and it is essential to use this limited information wisely in order to gauge where the true maximum or at least a particularly large value of the control landscape can be found.

Denoting the vectors of control parameters for which the control landscape has been probed by θ_j ($j = 1, \dots, M$) and denoting the corresponding observations by $\mathcal{D} = [F(\theta_1), \dots, F(\theta_M)]$, this task can be formalized as:

given the observations $\mathcal{D}$, what control pulse θ_{M+1} is expected to yield a particularly large value of F ?

In practice, this question can be interpreted in different ways: one may strive for a control pulse for which one is highly confident that it will yield a good result; or, one may strive for a more optimistic approach, testing a control pulse that holds the potential to yield an exceptional result even if one is not confident that this will actually be the case, as sketched in Fig. 1.

Bayesian optimization provides the solid statistical footing for taking this decision. It allows us to consider the full spectrum of conceivable control landscapes that are consistent with the given experimental observations. In the spirit of Bayesian analysis, one can assign a probability to every landscape that characterizes one's belief for this landscape to coincide with the actual control landscape. Since this provides a full probability distribution for conceivable control landscapes, it not only allows us

FIG. 1. Schematic example of the estimation of a control landscape based on a finite number of data points. In this example, the six given data points are consistent with a parabola shown in black, but they are also consistent with many other landscapes, such as that shown in red. Assuming the black landscape to be the correct landscape, one would predict the control parameter θ' to yield the maximum of the landscape, but one would predict the control parameter $\hat{\theta}$ when assuming the red landscape to be the correct landscape. Even though the red landscape might not seem likely to be the correct one given the six data points, it can be worth considering since it promises a larger value of F than the black landscape.

to construct an expected control landscape, but also to estimate when deviations from the expected behavior are sufficiently likely to justify exploring control parameters with high risk, but also high gain.

The ability to make predictions on the control landscape is then the basis for an iterative procedure. Based on the available information, one selects the most promising control pulse. Its performance then needs to be assessed experimentally. With the additional knowledge on the performance of this new pulse, one is able to make a more educated prediction for a subsequent pulse, and the repetition of these steps is likely to result in the identification of a pulse that achieves the desired goal.

A. Surrogate model

Since no finite number of data points will be enough to reconstruct the complete control landscape $F(\theta)$, Bayesian optimization aims at approximating this landscape in terms of a set of conceivable control landscapes that are called proxy or surrogate models. A simple example of such a surrogate model is given in Fig. 1, where the black parabola represents the actual – but unknown – control landscape, and the red curve represents one possible surrogate model. In practice, one typically considers a continuous set of surrogate models that are consistent with the data obtained from the experiment.

The central quantity required to make statistically sound predictions on the actual control landscape is the probability distribution for the actual control landscape to be given by a certain surrogate model f . Since this shall be done based on prior experimental observations, this is given by

the conditional probability distribution

$$P(f|\mathcal{D}) \quad (1)$$

for f , given the set of observations $\mathcal{D}$ made on the experiment. The construction of this predictive probability density follows Bayes' basic rule

$$P(f|\mathcal{D}) = \frac{P(f)P(\mathcal{D}|f)}{P(\mathcal{D})}, \quad (2)$$

in terms of the prior probability distribution $P(f)$ of surrogate models, the conditional probability distribution $P(\mathcal{D}|f)$ to obtain the acquired data for a given control landscape f , and the probability $P(\mathcal{D})$ to acquire the data $\mathcal{D}$ in a given sequence of experiments; the latter can be understood as a normalization constant and can be obtained from the condition $\int df P(f|\mathcal{D}) = 1$.

1. The prior distribution

The prior distribution $P(f)$ describes our limited knowledge on the control landscape before making any observations. In the absence of detailed knowledge about a given system, it is chosen based on general, physically reasonable assumptions; in practice, one would tend to find a landscape that varies smoothly with the control parameters more likely than a landscape with excessive oscillations or even discontinuities.

In the asymptotic limit, i.e., after infinitely many observations, Bayesian inference will identify the correct landscape essentially independently of the specific choice of prior distribution. There is thus nothing like a unique, correct prior distribution, but we are free in its choice. For Bayesian optimization to find a good solution quickly, however, it is important to use a reasonable prior distribution that is sufficiently broad to contain the actual control landscape, but that does not assign high probability to landscapes with unnatural features such as excessively wild oscillations or close-to-discontinuous dependencies. We defer the details of the practical definition of $P(f)$ to Sec. II F 1, and in the following take for granted that we can actually specify such a distribution.

2. Expected observations

The most relevant term in Eq. (2) for our present purposes is the conditional probability density $P(\mathcal{D}|f)$. Even though one hardly ever knows the true control landscape, one can characterize accurately what observations one would expect with a given control landscape.

In a perfectly idealized, noiseless situation, one would expect to observe the true value $f(\theta)$ of the underlying landscape f in an experiment performed with the control pulse θ . In this case, the conditional probability density for the figure of merit to adopt the value of x would thus be given by $P(x|f) = \delta[x - f(\theta)]$.

In a more realistic situation, any type of imperfection, such as finite resolution in a measurement or experimental noise can be incorporated into the term $P(\mathcal{D} | f)$, and instead of the above infinitely sharply peaked distribution, one obtains a probability distribution with a finite width. This can be done for essentially any type of imperfection, or combinations therefore, but we develop it in more detail for the case of measurement noise in the following.

B. Measurement noise

In the presence of large shot noise, it is helpful to express a figure of merit in terms of probabilities of measurement outcomes, rather than expectation values of observables. In most practical situations, the figure of merit will not just be the probability to obtain a given outcome for one specific measurement, but it will be a function of several physical observables, or, equivalently, a function of probabilities of outcomes of several measurements. Since all these cases require the proper description of individual measurements, we elaborate on the modeling of a single observable first, and defer the discussion of control targets comprised of several observables to Sec. II D.

Since the probability to obtain a given measurement outcome can only adopt values in the interval between 0 and 1, one should impose the same condition on the surrogate model f satisfying this condition, the conditional probability to obtain a detector click given the control landscape f is given by

$$P(1 | f(\theta)) = f(\theta), \quad (3)$$

whereas the probability to obtain no detector click is given by

$$P(0 | f(\theta)) = 1 - f(\theta). \quad (4)$$

This is readily generalized to the case of N repetitions of the same measurement, in which the probability to obtain n detector clicks is given by the binomial distribution

$$P(n | f(\theta)) = \binom{N}{n} f(\theta)^n [1 - f(\theta)]^{N-n}. \quad (5)$$

The estimation of a control landscape is not based on observations with one single control pulse θ , but rather based on observations with several different control pulses θ_j utilized in different repetitions of the experiment. Since, for any given surrogate model f , the measurement outcomes in experiments performed with different control pulses are statistically independent, the probability $P(\{n_j\} | f)$ to observe a sequence of n_1 to n_M detector clicks in a sequence of experiments performed with the control pulses

θ_1 to θ_M is given by the product

$$P(\{n_j\} | f) = \prod_{j=1}^M P(n_j | f(\theta_j)) \quad (6)$$

of the probabilities $P(n_j | f(\theta_j))$ to obtain n_j detector clicks in the experiments performed with the control pulse θ_j . This long product of binomial distributions is the explicit form of the general object $P(\mathcal{D} | f)$ entering the Bayesian inference above in Eq. (2). With this modeling in hand, one can explicitly construct the predictive probability density $P(f(\theta) | \mathcal{D})$, based on any level of measurement noise, including the extreme case of no repetitions, i.e., $N = 1$.

C. Predictions

With everything laid out so far, we are now ready to address the central step of predicting properties of the control landscape. As argued above, this will always be based on a discrete set of data points obtained from a discrete set of control pulses θ_j (with $j = 1, \dots, M$). It is essential that properties of the control landscape can be predicted for any conceivable control pulse θ . Since the goal is to do so for points corresponding to pulses that have not been experimentally explored yet, it is helpful to first consider the conditional probability density

$$P(f(\theta) | \vec{f}) \quad (7)$$

for the control landscape to contain the point $f(\theta)$, given that it contains the points $\vec{f} = [f(\theta_1), \dots, f(\theta_M)]$. This is a quantity that can be characterized without any experimental observation, as it is given by

$$P(f(\theta) | \vec{f}) = \frac{P(f(\theta_1), \dots, f(\theta_M), f(\theta))}{P(f(\theta_1), \dots, f(\theta_M))} \quad (8)$$

in terms of the prior distribution $P(f)$ described above in Sec. II A 1. Once the prior is defined explicitly, the conditional density $P(f(\theta) | \vec{f})$ is thus too.

Experimental observations are necessary in order to assess the probability for the actual control landscape to contain the points $\vec{f}$. Because of the noise in the data acquisition, this can indeed only be done in terms of a probability distribution, and it is given exactly by $P(\vec{f} | \mathcal{D})$. This posterior distribution can be obtained according to Bayes' rule as

$$P(\vec{f} | \mathcal{D}) = \frac{P(\vec{f})P(\mathcal{D} | \vec{f})}{P(\mathcal{D})}. \quad (9)$$

With the probabilities $P[f(\theta) | \vec{f}]$ and $P(\vec{f} | \mathcal{D})$, one can finally recast the desired predictive probability as

$$P(f(\theta) | \mathcal{D}) = \int d^M f P(f(\theta) | \vec{f}) P(\vec{f} | \mathcal{D}), \quad (10)$$

where $d^M f$ denotes the M -dimensional integral over the components of $\vec{f}$. This is exactly the quantity that enables predictions for any point on the landscape before any experimental data resulting from the pulse corresponding to this point is available. As anticipated above, it is not simply the expected value of the control landscape, but it is the full probability distribution that allows one to estimate whether there is a sufficiently high chance to find a particularly high value of the control landscape for any new control pulse.

D. Target functionals

So far, we have restricted the discussion to the case in which the figure of merit is exactly the probability for a detector click. In practice, however, many control targets F depend on the outcomes of several measurements. They can thus be expressed as the function

$$F = \mathcal{Q}(p_1, \dots, p_K) \quad (11)$$

of the probabilities p_k for the outcomes of those measurements. In this case, one can employ an independent surrogate model f_k for each of the K required probabilities, such that each f_k can be used to make predictions for the probability p_k . The surrogate model f of the control target F is then expressed as

$$f = \mathcal{Q}(f_1, \dots, f_K), \quad (12)$$

with exactly the same functional dependence as in Eq. (11). With this, one can construct the overall predictive distribution $P(f(\theta) | \mathcal{D})$ for the actual control target based on each individual predictive probability distribution $P(f_k(\theta) | \mathcal{D})$ for the individual probabilities.

E. Decision rules

With the probabilistic modeling for the control landscape defined, the remaining question is how to use the predictive probability distribution $P(f(\theta) | \mathcal{D})$ [Eq. (10)] in order to decide which control pulse to use in the next step of the iterative optimization.

At any level of given information obtained from the experiment can the expected control landscape be expressed as

$$\langle f(\theta) \rangle = \int df f(\theta) P(f(\theta) | \mathcal{D}). \quad (13)$$

Note that the notation $P(f(\theta) | \mathcal{D})$ is shorthand for $P(f(\theta) = p | \mathcal{D})$. Since the probability p lies in the interval $[0, 1]$, Eq. (13) can also be expressed as

$$\langle f(\theta) \rangle = \int_0^1 dp p P(f(\theta) = p | \mathcal{D}), \quad (14)$$

which is a regular expectation value with respect to the probability distribution $P(f(\theta) | \mathcal{D})$. Analogously, the variance

$$\sigma^2(\theta) = \langle [f(\theta) - \langle f(\theta) \rangle]^2 \rangle \quad (15)$$

around the expected landscape, or any other function A of $f(\theta)$, can readily be obtained in terms of a regular integral

$$\langle A[f(\theta)] \rangle = \int_0^1 dp A(p) P(f(\theta) = p | \mathcal{D}). \quad (16)$$

Particularly in cases of few experimental observations, the expected landscape $\langle f(\theta) \rangle$ does not necessarily approximate the exact landscape $F(\theta)$ well for any value of θ , but there can be substantial uncertainty. Since a strategy ignoring this will likely miss the global maximum, it is essential to take it into account, e.g., in terms of the variance or higher-order statistical moments.

A strategy based on both expectation and uncertainty allows one to identify values of θ for which large values of F can be expected. In practice, this can be realized in terms of the acquisition function

$$a(\theta) = \langle F(\theta) \rangle + \alpha \sigma(\theta), \quad (17)$$

with a real, non-negative scalar α that balances the weight of the expected landscape and the confidence in this estimate. Using the value of θ that maximizes $a(\theta)$ for the next query to the experiment allows the algorithm to broadly explore the control landscape, thus minimizing the risk of getting trapped in a local extremum.

The choice of the numerical value of α can be made depending on practical requirements. A small value will result in rapid convergence to some reasonably good solution that might not be the best solution available, whereas a large value reduces the risk of not finding the best solution, but will typically result in slower convergence. In practice, one may also change the value as the optimization progresses, starting with a large value favoring exploration of the entire landscape, followed by a decrease to focus on particularly promising domains.

F. Implementation

In Secs. II A to II E we conceptually described the basic structure of Bayesian optimization and its specificities in the presence of substantial measurement noise. In this final subsection, we discuss the most relevant aspects for the explicit implementation of these concepts, but we invite any reader who is more interested in the performance of the present techniques than the underlying details to jump ahead to Sec. III.

1. The prior distribution

In order to implement the algorithm laid out so far, one needs to define the prior distribution $P(f)$ of surrogate models explicitly. Crucially, this does not require an explicit parametrization of possible control landscapes, but this can be done elegantly in terms of Gaussian processes as detailed in the following.

A random process extends the concept of a vector of random variables to an infinite collection of random variables. Any function g can thus be understood as the continuum limit of a set of discrete random numbers $g(\theta_i)$. Since optimization landscapes are functions of continuous control parameters, random processes are the appropriate mathematical structure for the present purposes. In particular, the description of potential landscapes in terms of random processes avoids the necessity of an explicit parametrization, but ensures that properties like continuity and differentiability of the control landscapes are guaranteed.

A random process such that any finite collection of variables $[g(\theta_1), \dots, g(\theta_N)]$ follows a Gaussian distribution is referred to as a Gaussian process. While a Gaussian distribution is parameterized by a mean vector and a positive semidefinite covariance matrix, a Gaussian process is specified by a mean function $\mu(\theta) = \langle g(\theta) \rangle$ and a positive semidefinite kernel function $k(\theta, \theta')$ that defines the covariance

$$\langle g(\theta)g(\theta') \rangle - \langle g(\theta) \rangle \langle g(\theta') \rangle = k(\theta, \theta') \quad (18)$$

between $g(\theta)$ and $g(\theta')$ for arbitrary θ and θ' , where $\langle \cdot \rangle$ denotes the average, under this distribution. This covariance is typically taken to be exponentially decaying in the distance between the two arguments, which implies that knowledge of a function at some point θ allows one to estimate the function around this point with high confidence, whereas the ability for prediction decreases with the distance from this point.

A typical function $g(\theta)$ consistent with a Gaussian process will generally not qualify as a proper surrogate model, since it will not be bound to the interval $[0, 1]$ as a probability has to. A Gaussian process on its own is therefore not suitable for the present purposes, but one can arrive at a suitable distribution with a *squashing function* π , such that the functions $f(\theta) = \pi(g(\theta))$ are bounded by the interval $[0, 1]$. With $g(\theta)$ resulting from a Gaussian process, this construction yields a probability distribution for function $f(\theta)$ with the desired properties. In the following we use the cumulative distribution function of a standard normal distribution $\pi(x) = \int_{-\infty}^x dy \exp(-y^2/2)/\sqrt{2\pi}$, but essentially any monotonic function mapping the real axis to the interval $[0, 1]$ could be used.

With a specific choice of mean function $\mu(\theta)$ and kernel $k(\theta, \theta')$, one thus has the desired explicit realization of the prior distribution $P(f)$. Without any prior knowledge about the optimization landscape, one would typically

choose a vanishing mean, $\mu(\theta) = 0$, resulting in a flat expected landscape $\langle f \rangle = \frac{1}{2}$.

There are many possible choices for the kernel function $k(\theta, \theta')$ [50], and we choose them from the class of Matérn functions

$$C_j(\theta - \theta') = VR_j(x) \exp(-x) \quad (19)$$

with $x = |\theta - \theta'|/l$, defined in terms of a correlation length l and a variance V . Depending on the optimization problem at hand, we use the three different polynomials

$$R_0(x) = 1, \quad (20)$$

$$R_1(x) = 1 + \sqrt{3}x, \quad (21)$$

$$\text{and } R_2(x) = 1 + \sqrt{5}x + 5/3x^2. \quad (22)$$

The choice of polynomials depends mostly on the expected roughness of the control landscape. The polynomial R_2 effectively ensures that any f that occurs with nonvanishing probability is at least twice differentiable, R_1 results in control landscapes that are at least once differentiable, and R_0 enforces only that the landscapes are continuous [50].

Choosing suitable values for the variance V and correlation length scale l would require some knowledge of the scales over which the value of F changes. Since even a very limited amount of data can be used to identify the values of V and l , one can adapt them during the course of the optimization. We do so following standard practice [27, 50] by minimizing the log marginal likelihood term $\log P(\mathcal{D})$, with $P(\mathcal{D})$ defined in the context of Eq. (2).

2. Integration

The explicit construction of $P(f(\theta) | \mathcal{D})$ in Eq. (10) requires the solution of an M -dimensional integral. Already after a few steps in the iterative control algorithm a numerical evaluation of such an M -dimensional integral is prohibitively expensive. An efficient, but still accurate estimate is thus essential for the practical value of the control scheme. In the following we discuss how this is routinely done in many applications, and how the specificities of poor statistics require a different approach in order to maintain accuracy.

a. Integration with the binomial noise model. Many problems of Bayesian optimization are based on entirely Gaussian models whose integrals have analytic solutions. In the present case, however, only the first factor $P(f(\theta) | \vec{f})$ in Eq. (10) is Gaussian because of the underlying Gaussian process, but the second factor $P(\vec{f} | \mathcal{D})$ is non-Gaussian because $P(\mathcal{D} | \vec{f})$ is given in terms of the binomial distribution in Eq. (6).

In order to perform the required integration efficiently, we employ the Laplace approximation [50] for $P(\mathcal{D} | \vec{f})$ due to its conceptual simplicity. It entails approximating

$P(\mathcal{D} | \vec{f})$ by a Gaussian distribution, such that $P(f(\theta) | \mathcal{D})$ in Eq. (10) becomes an M -dimensional, analytically solvable integral over the product of two Gaussian distributions. It is thus possible to efficiently perform all integrations despite the detailed and non-Gaussian underlying noise given in Eq. (6), and we discuss the implications of the approximate integration in detail in Sec. III B below.

b. Gaussian noise. In a standard approach that does not aim at detailed modeling of measurement noise, one would estimate the probabilities p_k of measurement results in terms of the frequencies $p_k^{(e)} = n_k/N$ extracted from the observations. These estimated probabilities admit an estimation of the figure of merit F as

$$F \simeq F_e = \mathcal{Q}(p_1^{(e)}, \dots, p_K^{(e)}) \quad (23)$$

with the function $\mathcal{Q}$ defined in Eq. (11). The fact that this estimate does not necessarily coincide with the exact value of F is then taken into account with a phenomenological Gaussian probability distribution

$$P(F_e(\theta) | f(\theta)) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{[F_e(\theta) - f(\theta)]^2}{2\sigma^2}\right), \quad (24)$$

which, in this case, is defined in terms of a single surrogate model f . In contrast to the binomial modeling above, there is the phenomenological parameter σ ; its value is not determined by basic principles, but needs to be chosen in accordance with observations similarly to the choice of values for the Gaussian process discussed above in Sec. II F 1.

III. POOR STATISTICS AT WORK

In order to demonstrate the strength and limitations of the control algorithm developed in Sec. II, we apply it to four exemplary problems. The problem of state preparation of a single qubit discussed in Sec. III A is intended to be of pedagogical nature, giving insight into the workings of the methodology. This is followed by the analysis of the state preparation of a three-qubit GHZ state in terms of single- and two-qubit gates (Sec. III B), and the preparation of a Mott-insulating quantum phase in a Bose-Hubbard system (Sec. III C). In addition to these demonstrations based on numerical simulations, the problem of state preparation on a publicly available NISQ device [51] discussed in Sec. III D gives experimental evidence of the benefits of the present approach.

In these examples the efficiency of the present framework will be judged in terms of its ability to converge towards good solutions with minimal experimental effort. This effort is quantified as the total number N_r of experimental runs, i.e., individual cycles of preparation, evolution, and measurements. Assessing the expectation value

of a single observable with N repetitions, for example, requires $N_r = N$ runs, whereas assessing several commuting observables in terms of a single-shot measurement requires only a single run. The explicit examples in this section demonstrate that the present framework permits a reduction in the number of experimental runs by working with few repetitions N , i.e., poorly resolved expectation values of observables, allowing one to explore more control pulses at a given number of experimental runs N_r . If a given experiment poses a limitation on the number of control pulses to be tested, one may select the lowest number of repetitions N that is sufficient to converge within this number of pulses.

With the exception of the illustrative example of Sec. III A, the present framework with proper binomial modeling of measurement noise (as developed in Sec. II) will be systematically benchmarked against Bayesian optimization with Gaussian modeling of measurement noise (as sketched in Sec. II F 2 b), and the simultaneous perturbation stochastic approximation (SPSA) framework for which we follow the implementation described in Ref. [52]. These two techniques are especially competitive when data are noisy and limited [6,29–33,53–55] and as such provide strong benchmarks to compete with.

As the success of any optimization depends on the random choice of initial guesses but also on measurement data that are intrinsically probabilistic, most of the results presented in Sec. III correspond to optimizations that are repeated several times, with the exception of Fig. 2 in Sec. III A that depicts one specific instance of an optimization for illustrative purposes. For the results obtained based on numerical simulations, these repetitions are taken over different random seeds, and measurement data are sampled following Born's rule.

All the following numerical simulations and optimizations are based on QuTip [56], QuSpin [57], GPy [58], and GPyOpt [59]; and interfacing with the NISQ device in Sec. III D is based on Qiskit [60].

A. A single qubit to warm up

A simple toy model of pedagogical value is given by the state $|\psi(\theta)\rangle$ of a single qubit parameterized with a scalar control parameter θ . The goal lies in maximizing the fidelity with respect to a given state $|\phi\rangle$, and we assume for simplicity that a projective measurement in a basis including the state $|\phi\rangle$ can be taken. In the present case, the target functional is thus directly the probability to project onto $|\phi\rangle$, i.e., $F(\theta) = |\langle\phi | \psi(\theta)\rangle|^2$.

Rather than defining a control landscape in terms of the dynamics induced by a specific Hamiltonian, we consider the exemplary function

$$F(\theta) = \sin^2\left(\frac{1}{2} \sin\left(3\theta + \frac{9}{10}\right) + \frac{2\theta}{3} + \frac{1}{5}\right). \quad (25)$$

FIG. 2. Three stages of a typical run of Bayesian optimization with extremely bad measurement statistics after (a) 30, (b) 31, and (c) 100 steps. The top panel depicts the actual control landscape (dashed red line), the expected control landscape (solid blue line), and a 95% confidence interval (solid gray line) around the expected landscape. The actual probability density is represented by shades of blue. Simulated, random outcomes of a projective measurement are depicted by red circles and green squares, where green corresponds to the most recent data point. Assuming a sufficiently smooth landscape, it is possible to predict control landscapes ranging continuously between the values 0 and 1 despite the digital data. The lower panel depicts the acquisition function [Eq. (17) with $\alpha = 4$ for (a) and (b) and $\alpha = 0$ for (c)], the maximum of which determines the value of θ for the next upcoming step in the iterative algorithm. During the course of the optimization (from left to right), the surrogate model becomes a better approximation of the actual landscape in the vicinity of the maxima, but the algorithm avoids the effort that would be required to approximate the landscape well in other domains.

This landscape, depicted by a dashed red line in Fig. 2, reaches the maximum value of 1 and has two additional local maxima in the interval $[0, 4]$; it is thus well suited to illustrate the ability of the algorithm to distinguish local from global extrema.

In Fig. 2 we depict the progress in data acquisition and the estimates of the control landscape during the course of the optimization using the kernel defined in Eq. (19) with P_2 given in Eq. (22). The optimization starts after taking 30 initial single-shot projective measurements for randomly chosen θ_j ($j = 1, \dots, 30$), as shown in Fig. 2(a). The measurement results depicted by red circles and a green square can only adopt the values of 0 and 1, and each value is randomly drawn following the actual control landscape. Based on those 30 observations, one obtains a rough estimate of the actual control landscape; the expected landscape $\langle f(\theta) \rangle$ is depicted in blue, a 95% confidence interval is depicted by the gray contour, and the actual probability density is represented by shades of blue. As one can see, the expected control landscape captures the most salient features like the approximate locations of the three maxima, but, given the very limited data, it fails to reproduce accurately the actual landscape. In particular, it fails to distinguish the local from the global maximum.

Based on this model, the decision rule described in Sec. II E determines for which value of θ to take the next measurement. The acquisition function [Eq. (17)], with a value of $\alpha = 4$, is depicted in red in the bottom panel of Fig. 2(a) and takes its maximal value for $\theta_{31} \approx 0.1$ (red vertical line). This decision reflects the high uncertainty in the model in this region where no experimental observations are available yet.

Figure 2(b) contains the additional data point (green square), resulting from the projective measurement with the value θ_{31} of the control parameter. The updated surrogate model in (b) shows that the surrogate model in (a) has overestimated the value of $F(\theta_{31})$. In addition to the better estimate of this value, the uncertainty around θ_{31} has also slightly decreased from (a) to (b). As a result of these two effects, the next probe is taken for the value $\theta_{32} \approx 2.3$, i.e., in the vicinity of the maximum where the value of the expected control landscape is moderately high and, in addition, the uncertainty is large.

In Fig. 2(c) we depict the surrogate model after 100 queries to the experiment. As one can see, the framework suggested several probes in the vicinity of the three maxima, and managed to avoid queries that would not have resulted in substantial added value. Consequently, the surrogate model approximates the actual control landscape more accurately close to these maxima.

After these 100 iterations the algorithm has identified the value $\theta_f = 1.86$ [as indicated in green in Fig. 2(c)] as optimal, which is very close to the true optimal value $\theta_o = 1.876$. With the value $F(\theta_f) = 0.996$ of the actual control landscape, the algorithm has thus found a solution with an infidelity of 0.004, after a number of measurements that would have only been enough to determine a single point of the control landscape with a resolution of 0.01.

B. GHZ states

With the flavor of the workings of optimal quantum control based on poor statistics laid out in Sec. III A, we can now proceed to a more challenging control problem. We consider the gate sequence depicted in Fig. 3 for the

FIG. 3. Gate sequence for the preparation of a three-qubit GHZ state in terms of two controlled-NOT gates and six single-qubit gates with adjustable parameter θ_j . The addition of noisy unitaries N_ϵ and readout error helps to demonstrate optimal control in the presence of experimental noise in addition to the measurement shot noise.

preparation of a three-qubit GHZ state $|\Psi\rangle = (|000\rangle + |111\rangle)/\sqrt{2}$ that could be realized for example in trapped ion [61–63] or superconducting qubit [64–66] hardware. It consists of five single-qubit gates $R_x(\theta_j) = \exp(i\theta_j\sigma_x)$ ($j = 1, \dots, 5$), one single-qubit gate $R_y(\theta_6) = \exp(i\theta_6\sigma_y)$, and two controlled-NOT gates, as shown in Fig. 3.

The goal is to find suitable angles $\theta_j \in [0, 2\pi]$ such that the circuit maps the initial state $|000\rangle$ into the GHZ state $|\Psi\rangle$. Below we also consider additional noisy unitaries N_ϵ , but for the moment, they are treated as identities.

1. Control targets

The fidelity F for the state ρ resulting from the gate sequence is defined as $F = \langle \Psi | \rho | \Psi \rangle$. Unless one is able to perform a projective measurement in a basis including the state $|\Psi\rangle$, however, one is bound to perform measurements on each individual qubit, as indicated by the three detectors in Fig. 3. In practice, one would therefore construct the fidelity in terms of local measurements.

Fidelities are often expressed in terms of expectation values of observables. In the case of a GHZ state, the fidelity can be decomposed as

$$F = \frac{1}{8} \left[1 + \sum_{k=1}^4 \text{tr}(\rho S_k) - \sum_{k=5}^7 \text{tr}(\rho S_k) \right], \quad (26)$$

where $S_1 = \sigma_x \otimes \sigma_x \otimes \sigma_x$, S_k with $k = 2, 3$, and 4 are the permutations of $\mathbb{1} \otimes \sigma_z \otimes \sigma_z$, and S_k with $k = 5, 6$, and 7 are the permutations of $\sigma_x \otimes \sigma_y \otimes \sigma_y$.

Since the framework developed so far allows us to estimate probabilities rather than expectation values, it is more appropriate to express the fidelity in Eq. (26) in terms of probabilities for measurement outcomes. Denoting $P_k = (S_k + \mathbb{1})/2$ as the projector onto the subspace spanned by the eigenstates with eigenvalue $+1$ of the observable S_k , the state fidelity for the GHZ state can be written as

$$F = \frac{1}{4} \left[\sum_{k=1}^4 \text{tr}(\rho P_k) - \sum_{k=5}^7 \text{tr}(\rho P_k) \right]. \quad (27)$$

In contrast to the simpler case discussed above in Sec. III A, the fidelity is thus not given by the probability of one single measurement outcome, but it is given in terms of the seven probabilities $p_k = \text{tr}(\rho P_k)$. It is thus necessary to employ seven surrogate models f_k as detailed in Sec. II E, and to estimate the state fidelity as

$$f = \frac{1}{4} \left[\sum_{k=1}^4 f_k - \sum_{k=5}^7 f_k \right], \quad (28)$$

consistent with Eq. (27).

2. Optimization

The following optimizations are obtained with R_0 [Eq. (20)] for binomial modeling and with R_2 [Eq. (22)] for Gaussian modeling. In the examples, we consider convergence in terms of the total number of runs N_r required to reach a given fidelity. Since the three observables S_5 , S_6 , and S_7 commute, an optimization based on the assessment of the fidelity [Eq. (27)] for M different sets of control parameters, with each measurement repeated N times, requires $N_r = 5MN$ runs.

With few repetitions for a measurement, in particular, in cases of single-shot measurements ($N = 1$), any individual run of an optimization has a substantial probabilistic component. The following comparisons will thus be based on statistics of the infidelity $\mathcal{I} = 1 - F$, such as the median and quartiles, over 30 random instances of the same control task.

In Fig. 4 we depict the decrease in infidelity as a function of the total number of runs N_r . Different colors correspond to different numbers of repetitions ranging from the single-shot case $N = 1$ to $N = 1000$; results from binomial and Gaussian modeling are depicted with filled and open symbols. Fluctuations in convergence towards high fidelities due to statistical fluctuations in measurement results are characterized in terms of the interquartile intervals depicted by shaded regions.

In the case of single-shot measurements (depicted in black), control based on Gaussian modeling essentially fails, whereas binomial modeling yields infidelities in the percent regime. With a larger number of repetitions N , Gaussian modeling also results in low infidelities, but for any given choice of N (with the exception of $N = 1000$), control based on binomial modeling converges faster; after a given number of runs N_r , it yields infidelities about an order of magnitude lower than control based on Gaussian modeling. The gap between the two methods decreases slowly as N increases, and for $N = 1000$, one might have expected Gaussian modeling to become comparable to binomial modeling. There are, however, instances in which the Gaussian version slightly outperforms the binomial version. This can be attributed to the Laplace approximation (Sec. II C). Close to the extreme values of 0 and 1 of

FIG. 4. Convergence of the control algorithms towards low infidelities $\mathcal{I}$ as a function of the total number of runs N_r of the circuit required. The median values of $\mathcal{I}$ are indicated by filled symbols for the binomial modeling of measurement noise and by open symbols for Gaussian modeling. The different colors correspond to different numbers of repetitions of each measurement, ranging from $N = 1$ in black to $N = 1000$ in orange. Interquartile intervals are depicted by shaded regions. Optimizations with the proper binomial modeling converge substantially faster than optimizations with Gaussian modeling of measurement noise, except in the case of a large number of repetitions $N = 1000$. The inset depicts results for fewer repetitions ranging from $N = 1$ in black to $N = 10$ in orange, and binomial modeling only. Interquartile intervals are not displayed for visual clarity.

the underlying probabilities this approximation becomes less accurate [50], resulting respectively in an overestimation and underestimation of the actual value of the probabilities. Signatures of this effect can also be seen in Fig. 2(c), where the expected control landscape is systematically below the true landscape in the vicinity of the global maximum. Even though this effect can be reduced at the expense of higher computational effort [67], this does not seem necessary, since these stagnations occur in a regime in which the Gaussian modeling becomes appropriate and analytic solutions for integrals are available (Sec. II F 2 b).

One can see from Fig. 4 — especially in the inset — that in all cases where there is data for different values of N for a given infidelity, the optimization with the fewest repetitions performs best. It thus always seems to be preferable to explore many points of the control landscape with few repetitions rather than trying to resolve the control landscape accurately for few points. Reducing the number of repetitions of the same measurement, however, also implies additional computational cost: the algorithm scales as $\mathcal{O}(M^3)$ [50] in terms of the number of iterations M , and at a constant number $N_r \propto MN$ of measurements, a reduction in N implies an increase in M . That is, whereas

fundamentally optimizations with the fewest repetitions seem to perform best, it can become necessary in practice to increase the number of repetitions N .

3. Adaptive strategy

Given that working with few repetitions of the same measurement (i.e., small values of N) is beneficial for fast convergence, but that this expensive benefit tails off at high fidelities, one may use a variable number of N for an easily implementable reduction of computational effort. We find it most practical to start the optimization with few shot measurements (e.g., $N \leq 10$) in order to rapidly identify good domains in parameter space. As the search circles in on a good solution, one can restrict the exploration to a smaller domain; dropping the data outside this domain reduces the numerical effort, and the number of repetitions N can be increased as the search approaches a high-fidelity solution.

The black curve in Fig. 5 depicts the convergence of this strategy. The initial 150 steps of the optimization are performed with $N = 5$ repetitions and binomial modeling of the noise, corresponding to the first three data points of the curve. After these initial steps the parameter space is reduced around the 75 best parameters probed so far, and the number of repetitions is increased. The next three points, with N_r close to 10^6 runs of the circuit, correspond to an additional 500 steps in the optimization with Gaussian modeling and an increased number of $N = 100, 250$, and 500 repetitions. Finally, the two remaining points of the curve, with $N_r > 10^8$ runs, correspond to an additional stage of optimization with Gaussian modeling starting with the solutions obtained with $N = 500$ repetitions. For this round, the parameter space is reduced once more, and N is further increased to the values $N = 50\,000$ and $N = 500\,000$, respectively. Overall, the adaptive strategy quickly locates good control parameters (infidelities below 1% after 5000 runs) and subsequently is able to refine them to reach infidelities as low as $\mathcal{I} \approx 10^{-8}$.

The remaining five data sets (shown in color in Fig. 5) correspond to optimizations performed using SPSA with a fixed number of repetitions N . Although SPSA also shows convergence towards low infidelities, this convergence is substantially slower than with Bayesian optimization: comparison of Figs. 4 and 5 for any given number of repetitions N indicates dramatically faster convergence of the optimizations with binomial modeling than with SPSA, both during the early stage of the optimization when the landscape is being explored broadly and later when solutions are being refined close to the optimum. Also, the adaptive framework that enables Bayesian optimization with many runs N_r clearly outperforms SPSA at any given number of runs N_r .

In addition, the Bayesian optimization with binomial modeling is also rather robust to statistical noise in

FIG. 5. Comparison of the adaptive strategy (black) explained in Sec. III B 3 with optimizations performed using SPSA with a fixed number of repetitions N of each measurement given in legend. The adaptive method converges systematically faster than any of the other cases.

the measurement results, whereas the interquartile range (shaded regions of Fig. 5) shows that optimizations based on SPSA are subject to very large fluctuations in performance. In fact, a significant fraction (between 25% and 50%) of all optimizations does not get close to the achievable fidelities, but gets trapped in local minima with an infidelity of about 50%.

4. Additional noise

Last but not least, we can demonstrate that the present control algorithm is not just able to cope with measurement noise, but that it can also be used to find good controls in the presence of additional sources of noise. To this end, we can consider the noisy unitaries $N_\varepsilon = \sqrt{1 - \varepsilon^2} \mathbb{1} + \varepsilon \vec{n} \vec{\sigma}$, where $\vec{\sigma}$ is the vector of the three Pauli matrices.

In each individual run of the gate sequence, ε is drawn from a Gaussian distribution with vanishing mean and width σ_N , and the vector $\vec{n}$ of unit length is drawn from a uniform distribution following the Haar measure. In addition, we consider a finite probability of readout error p_{ro} for each of the three detectors to yield the wrong result, i.e., “0” instead of “1” and vice versa.

In the presence of this noise with no correlations between the different noisy elements or between consecutive runs, it is fundamentally impossible to reach perfect fidelities. The present goal therefore cannot be to aim at unit fidelity in a noisy circuit, but to demonstrate the ability of Bayesian optimization to find good solutions with a noisy circuit. The infidelities in Fig. 6 therefore correspond to what is achieved with a noiseless circuit but for a

FIG. 6. Convergence to low infidelities in the case of noisy circuits. The different colors correspond to different noise levels as specified in the legend where σ_N denotes the amplitude of the noisy unitaries and p_{ro} the probability of a readout error. The noiseless case (black curve) is provided as a reference.

control solution that has been optimized in the presence of noise.

Similarly to the other figures, Fig. 6 depicts a decrease in the median infidelity and quartiles as a function of the total number of runs of the circuit for different values of σ_N and p_{ro} , using the adaptive strategy described in the previous section with an initial number of $N = 5$ repetitions. Even though the control algorithm is entirely agnostic to the nature of the additional noise, there is rapid progress towards high fidelities. Most cases feature a rapid drop in infidelity after a few thousand runs, with the $\sigma_N = 5\%$ and $p_{ro} = 10\%$ cases being the exception. Even in the cases of very substantial additional noise, however, the optimization manages to converge well to low infidelities.

These results can be compared to the convergence of the SPSA algorithm in the presence of noise shown in Fig. 7 for the single-shot (i.e., $N = 1$) measurements case. Remarkably, noise can help in the convergence of the algorithm, especially for a low total number of runs ($N_r < 10000$). The large interquartile intervals depicted in Fig. 7, however, indicate that more than 25% of the optimizations ended up trapped in local minima with high infidelity ($\mathcal{I} \approx 50\%$), whereas such unsuccessful instances of an optimization with binomial Bayesian optimization are very rare.

For small levels of unitary noise ($\sigma_N = 0.1\%$), a typical instance of a SPSA optimization converges an order of magnitude slower than with binomial Bayesian optimization. For higher noise levels, this advantage is slightly reduced as statistical noise is not the main source of error anymore. Still, in all cases binomial Bayesian optimization converges a factor of 2 to 5 faster than SPSA.

FIG. 7. Convergence of SPSA in the presence of noise. All the results are obtained with single-shot measurements ($N = 1$), since this gives the best results in the noiseless case.

C. Quantum many-body states

Having analyzed the performance of the present approach with the example of a GHZ state, we can show with the example of a Bose-Hubbard system that the findings in Sec. III B are not specific to systems of qubits, but that similar performance can be expected for control of a broad range of quantum systems.

1. Model

We consider a system of N_b bosons in a one-dimensional arrangement with L sites and periodic boundary conditions. The bosonic creation and annihilation operator at site i of the chain is denoted $\hat{b}_i^\dagger$ and $\hat{b}_i$, and the local number operator reads $\hat{n}_i = \hat{b}_i^\dagger \hat{b}_i$. The physics is described by the Bose-Hubbard model

$$H(t) = -J(t)\hat{T} + U(t)\hat{V}, \quad (29)$$

with tunneling

$$\hat{T} = \sum_i \hat{b}_{i+1}^\dagger \hat{b}_i + \hat{b}_{i-1}^\dagger \hat{b}_i, \quad (30)$$

on-site interaction

$$\hat{V} = \frac{1}{2} \sum_i \hat{n}_i(\hat{n}_i - 1), \quad (31)$$

as well as tunneling rate $J(t)$ and interaction strength $U(t)$ that can be tuned in actual optical lattice experiments [68,69]. We consider the transition from a noninteracting system ($U/J \rightarrow 0$) with the superfluid ground state as the initial state to a Mott-insulating phase with $J/U \rightarrow 0$, since the preparation of many-body quantum states with

well-defined particle number per site is a prerequisite in many applications with neutral atoms [70–73].

In order to account for the fact that the experimental values of the tunneling rate and interaction strength are bounded, while perfect, adiabatic protocols in finite time are obtained in the limit $J, U \rightarrow \infty$, we parameterize $J(t)$ and $U(t)$ as

$$J(t) = 1 - \Gamma(t), \quad (32)$$

$$U(t) = \Gamma(t), \quad (33)$$

with $\Gamma(t=0) = 0$ and $\Gamma(t=T) = 1$.

The goal of finding a time-dependent control function $\Gamma(t)$ such that the system ends up in the Mott-insulating ground state becomes increasingly difficult with decreasing sweep time T , and the quantum speed limit [12] asserts that it cannot be achieved if T is below the threshold of π/Δ with the minimal lowest energy gap Δ of the system Hamiltonian during the sweep. In the following, we choose the specific value of $T = (3\pi)/(2\Delta)$, which is substantially shorter than the time scale required for an adiabatic transition. An optimal solution under this time constraint will typically not be a slight deformation of a solution that is optimal when close to the adiabatic regime. Still, this duration is sufficiently long to find a good control pulse $\Gamma(t)$ within the additional restrictions of an explicit parametrization.

Because of the rapid growth of Hilbert space dimensions that limits the accessible range of exact diagonalization, the following example is limited to a system with $L = 5$ sites and $N_b = 5$ bosonic particles, which corresponds to a 16-dimensional Hilbert space, taking into account selection rules resulting from the system's translational and inversion symmetries.

2. Control function

To be amenable to optimization, the time-dependent control $\Gamma(t)$ is parameterized in terms of S parameters θ_i corresponding to points equally spaced in time, i.e.,

$$\Gamma(t_i) = \theta_i \quad \text{with} \quad t_i = \frac{i}{S+1}T. \quad (34)$$

The parameters θ_i for $i = 1, \dots, S$ are bound to the interval $[0, 1]$, and the values of the end points θ_0 and θ_{S+1} are determined by the above boundary conditions, i.e., $\theta_0 = 0$ and $\theta_{S+1} = 1$. A continuous control function $\Gamma(t)$ is obtained uniquely in terms of a cubic spline interpolation between the points θ_i .

3. Control target

In principle, one could select a state fidelity as the control target. Given the rapid growth of the Hilbert space with particle number N_b and site number L , however, the vast

majority of quantum states would yield a negligibly small fidelity, and the state fidelity will fail to identify progress in the optimization.

With scalability in mind, we therefore employ a different control target, defined in terms of the local particle numbers that can be experimentally observed, for example, with atomic gas microscopes [74]. Denoting the probability to observe exactly one atom at site i by p_i , we strive to minimize the average filling error

$$\mathcal{E} = 1 - \frac{1}{L} \sum_i p_i. \quad (35)$$

Similarly to the infidelity, it yields the optimal value only for the target Mott-insulating state. In contrast to the state fidelity, however, states with an atom distribution similar to the Mott-insulating state yield a close to optimal value. The average filling error thus results in smoother control landscapes that allow algorithms to identify a reasonably broad peak around the target state, rather than an exponentially narrow peak, as would be the case with the state fidelity. Moreover, the average filling error is less sensitive to readout errors. Any of such readout error would cause the fidelity to drop to zero, while it will only result in a smaller increase, inversely proportional to the system size, in the average filling error.

4. Results

Similarly to the results presented in the previous example, we start with a comparison between Bayesian optimizations with binomial and Gaussian modeling. In Fig. 8 we depict the result of this comparison in terms of the filling error $\mathcal{E}$ for different fixed numbers of repetitions per measurement, with all optimizations limited to $M = 2000$ iterative steps.

Reaching subpercent filling errors is not possible for the fast protocol and the small number of control parameters ($S = 5$). In order to gauge the performance of the optimizations with a finite number of repetitions, we compare them to optimizations without shot noise, i.e., perfectly resolved data. In the latter case all the optimizations converge reliably to the same solution with a filling error $\mathcal{E}_{\min} \approx 1.6\%$ that is shown in Fig. 8 by a horizontal red dashed line.

Similarly to Sec. III B, binomial modeling (filled symbols) shows significantly better performance than its Gaussian counterpart (open symbols). In particular, with 25 or fewer repetitions, the Gaussian version failed to improve beyond the mediocre filling error value of $\mathcal{E} = 10\%$, while the binomial ones converge steadily towards the supposed minimum $\mathcal{E}_{\min}$. Only with more frequent repetitions ($N = 50$ and $N = 100$ in Fig. 8) does the Gaussian version manage to converge towards the low filling error $\mathcal{E}_{\min}$.

Results of the adaptive strategy described in Sec. III B 3 are presented in Fig. 9 (black curve). The first four points

FIG. 8. Convergence of the filling error $\mathcal{E}$ of a bosonic chain as a function of the total number of runs N_r . Results obtained with binomial and Gaussian modeling are depicted with filled and open symbols, respectively. All optimizations are performed with a fixed number of repetitions N of each measurement (specified in the legend) and a total number of iteration steps M ranging from 250 to 2000. The dashed red line shows the result of the best solution found with $\mathcal{E}_{\min} \approx 1.6\%$.

of the curve pertain to the first stage with a small number of repetitions (here $N = 10$). Binomial modeling is used for up to $M = 1500$ steps. After this, the parameter space is reduced (similarly to the example of GHZ states in Fig. 5) and the number of repetitions is increased. Results of this second round for 1000 extra iteration steps and several choices for the increased number of repetitions, $N = 50, 100, 200$, are shown as the remaining three points of the curve. This strategy converges towards the supposed minimum $\mathcal{E}_{\min}$ after a total number of $N_r \approx 10^5$ runs.

Results obtained with the SPSA method for different numbers of fixed repetitions are depicted in color in Fig. 9. Consistent with the findings for GHZ states, depicted in Fig. 5, the convergence with SPSA is substantially slower than the convergence with Bayesian optimization and binomial modeling. The optimizations with SPSA using $N = 10$ and $N = 100$ repetitions managed to converge close to $\mathcal{E}_{\min}$, but did so after substantially more runs than with the adaptive method. The optimizations with SPSA using $N = 1$ and $N = 5$ repetitions even failed to converge to $\mathcal{E}_{\min}$ within 2.5×10^5 iterations.

Figure 10 gives evidence of the robustness of the framework against fluctuations in the number of trapped atoms. It shows (in purple) convergence to low filling errors assuming a noise model with a 10% chance for the system to be initialized with four atoms or six atoms, respectively. The filling error resulting from the best control found (obtained with perfect data without measurement noise) is depicted with a dashed line for reference. Corresponding

FIG. 9. Optimization results for the adaptive strategy (Sec. III B 3) and SPSA with different numbers of repetitions N . The dashed red line indicates the lowest obtained filling error.

data for the case without fluctuations in the particle number are included for comparison (in black). As one can see, in both cases, there is very reliable convergence to an optimal solution.

Since we have consistently found that Bayesian optimization with binomial modeling outperforms all tested alternative approaches, it seems in order to stress the aspects of this method that gives rise to this preferred behavior. A general aspect of Bayesian optimization is that at any step in the iterative procedure the full set of accumulated data is used in order to identify the next control parameters to probe, whereas gradient methods employ

FIG. 10. Convergence of optimizations in the presence and absence of fluctuations in the particle number N_b (purple and black). In both cases the final filling errors are very close to the best filling errors obtained with simulations assuming perfectly accurate measurements (dashed lines).

only data in the vicinity of one point of the control landscape. As such, the scarce resource “data” is used in a more sustainable fashion and thus results in more accurate estimates of the control landscape. We explicitly verified that estimates of the control landscape based on measurement with few repetitions are substantially more accurate with Bayesian inference than with estimates in terms of relative frequencies of measurement outcomes. This effect is particularly pronounced towards the maxima of the control landscape where probabilities of measurement results are close to the extreme values of “0” and “1”.

D. Optimization on a NISQ device

After verification of the optimization’s convergence in theoretical simulations, we use the task of state preparation on publicly available NISQ devices, namely the IBM Q Experience quantum chips [51], as final demonstration.

The aim is to prepare the single-qubit target state

$$|\Psi\rangle = \cos\left(\frac{\pi}{8}\right)|0\rangle + e^{-i\pi/4}\sin\left(\frac{\pi}{8}\right)|1\rangle, \quad (36)$$

using a parameterized circuit, composed of a rotation around the x axis, $R_x(\theta_1) = \exp[-i(\theta_1/2)\sigma_x]$, followed by a rotation around the z axis, $R_z(\theta_2) = \exp[-i(\theta_2/2)\sigma_z]$.

Similarly to Eq. (27), the fidelity of a state ρ with respect to the target state is recast in terms of experimental observables as

$$F(\rho) = \frac{1}{2}[(1 - \sqrt{2}) + \langle P_x \rangle - \langle P_y \rangle + \sqrt{2}\langle P_z \rangle], \quad (37)$$

where P_i denotes the projector onto the eigenspace with positive eigenvalue of the Pauli operator σ_i , and $\langle P_i \rangle$ is shorthand notation for $\text{tr}(\rho P_i)$.

The control problem defined by the parameterized circuit and the target state of Eq. (36) admits a solution $\theta_{\text{id}} = (\pi/4, \pi/4)$ with unit fidelity in an idealized noise-free implementation. However, due to the limited accuracy of the gates and readout in present NISQ devices, one cannot expect to reach close-to-unit fidelities. To gauge what is realistically achievable, we estimate fidelities with parameters θ_{id} on the quantum chips. Each estimate is obtained by taking $N = 200\,000$ measurement repetitions per observable with a run time between 10 and 15 min per estimate. The overall median obtained for these estimates taken at different points in time is 98.2%, consistent with the errors in readout reported on the IBM Q platform [51], which range between 1% and 5% as the numbers are updated. The standard deviation of these fidelities is found to be 0.92%, indicating a significant amount of fluctuation in the device’s characteristics over time.

The time cost to resolve a fidelity with high accuracy indicates the burden of optimizations with accurately resolved observations. The following optimizations are therefore based on observations with low statistics only.

TABLE I. State fidelities [Eq. (37)] resultant from optimizations with binomial Bayesian optimization (BO), Gaussian Bayesian optimization, and SPSA on a NISQ device. The fidelities are medians over several repetitions, and the values in brackets denote the uncertainties of these medians due to the finite number of repetitions.

	$N_r = 300$		$N_r = 1500$
	$N = 1$	$N = 5$	$N = 5$
Binomial BO	0.966 (0.011)	0.950 (0.017)	0.978 (0.002)
Gaussian BO	0.852 (0.087)	0.943 (0.020)	0.958 (0.034)
SPSA	0.943 (0.013)	0.933 (0.014)	0.970 (0.002)

Even though monitoring convergence of an algorithm towards high fidelities might be insightful, the excessive burden of accurate fidelity estimates during the course of the optimization is prohibitive. In contrast to the detailed numerical investigations of the previous sections, the performance of the various control techniques will be characterized in terms of final fidelities only, and those are recorded with high statistical accuracy based on $N = 200\,000$ repetitions per observable.

Table I depicts the median of fidelities over several repetitions of optimization based on binomial Bayesian optimization, Gaussian Bayesian optimization, and SPSA with $N_r = 300$ and $N_r = 1500$ runs. The approximate gradients in the optimizations with SPSA are estimated in terms of finite differences of decreasing magnitude as the optimization converges. We start with a magnitude of 0.63 that is decreased to a value of 0.42. New values of the control parameters are chosen with this estimated gradient and an additional scalar factor, i.e., step size. This step size is gradually decreased from an initial value of 1.0 to a final value of 0.1 as the optimization converges. We have explored various different choices for these parameters in numerical simulations, and this choice has proven to yield the best results.

For each optimization, the initial random parameters are drawn uniformly in the interval $[-\pi/4, 3\pi/4]$, i.e., an interval of width π around the parameters θ_{id} . Besides the changes in the initialization, the different repetitions are performed identically.

Since the central feature of the optimization procedure is that decisions on control parameters are taken based on experimental observations, it is not possible to submit a complete optimization problem to the queuing system that realizes the interface to the quantum device. It is rather necessary to submit a new job for every new set of control parameters to be tested. This resulted in durations ranging between 15 min and 5 h for a complete optimization, with the wait time in the queue being the dominant contribution. Over a period of two weeks, each optimization was repeated 5 times, except for the case with $N = 5$ and $N_r = 300$ (second column of Table I) for which they were repeated 15 times because of larger fluctuations.

TABLE II. Results of the five individual optimization runs for the case involving a total number $N_r = 1500$ runs of the circuit and $N = 5$ shots. Both the fidelity based on the parameters θ_{opt} optimized using the binomial version of Bayesian optimization (first row) and also based on the solution θ_{id} of the control problem for a noise-free implementation (second row) are estimated on the quantum device at the end of each run.

	Fidelity (%)				
	θ_{opt}	θ_{id}	θ_{opt}	θ_{id}	θ_{opt}
θ_{opt}	97.49	97.06	97.82	98.65	98.07
θ_{id}	96.85	97.83	97.10	98.76	98.72

Similarly to the discussions in the previous sections, the fidelities in Table I are medians over these repetitions. In addition to variations due to the statistical nature of quantum mechanics and initializations of the search for optimal solutions, there are also variations in the properties of the NISQ devices due to drift and recalibration. The resultant inaccuracy in the estimate of the state fidelities are depicted in brackets in Table I.

As one can see from Table I, the binomial version of Bayesian optimization outperforms both SPSA and Gaussian Bayesian optimization in all cases. Gaussian Bayesian optimization fails to reach high fidelities in the single-shot case ($N = 1$), but both the binomial Bayesian optimization and SPSA reach high fidelities. Also, the fluctuations in the fidelities obtained with the latter two control algorithms are lower than in the case of Gaussian Bayesian optimization. Binomial Bayesian optimization thus clearly outperforms Gaussian Bayesian optimization as expected. Furthermore, binomial Bayesian optimization also outperforms SPSA. In the case of fewer runs of the circuits ($N_r = 300$), this advantage is of the order of several percent and only after optimizations with a larger number of runs ($N_r = 1500$) is the gap between the fidelities obtained with the binomial Bayesian optimization and SPSA smaller; this highlights the fact that optimizations with binomial Bayesian optimization converge more quickly, making this approach more resource efficient.

Finally, the details of the individual runs of optimizations for the case $N_r = 1500$ runs and $N = 5$ shots with binomial Bayesian optimization (corresponding to the results of the first row, last column of Table I) are presented in Table II. For each of the five optimization runs, the fidelity of the optimized parameters is reported (first row). Furthermore, at the end of each optimization, we also estimate the fidelity with parameters θ_{id} (second row).

In each individual run the optimization converged to a fidelity that is comparable to the corresponding fidelity obtained with θ_{id} . The deviations between the fidelities obtained with the optimized parameters θ_{opt} and with θ_{id} are in the subpercent range, and thus under the noise level of the readout. The total number $N_r = 1500$ of runs employed in the present example is comparable to the

number of repetitions that is typically used to determine the expectation value of an observable in a single step of optimization [30,32,53,55]. With the equivalent of a single step in a more conventional framework, the present control algorithm has thus found a solution that is limited mostly by the device imperfections.

IV. CONCLUSIONS

While the description of measurement outcomes in terms of their full probability distribution is extensively used in the context of characterization of quantum systems [75,76], quantum estimation, and metrology [77–79], it has hardly found applicability in the context of quantum optimal control. The improvement resultant from accurate probabilistic modeling found here highlights that tools from statistical analysis can also have similar impact on quantum optimal control.

The ability to find close-to-optimal solutions with limited experimental data enabled by such statistical methodology can advance technological development and precision experiments on a wide range of physical systems. It offers a very resource-efficient pathway towards the optimal use of currently existing quantum hardware with 10 to 100 qubits; the size of these systems makes theoretical modeling prohibitively expensive, and the noisy character of the individual qubits calls for well-designed control sequences that prevent rapid accumulation of errors. Given the availability of cheap resources for *classical* computation, it is essential to use them as much as possible to support the limited capabilities of near-term quantum hardware. The present framework contributes directly to this goal in that it allows us to find optimal uses of quantum systems in terms of limited experimental data, but at the expense of increased computational overhead compared to other control algorithms. This extra computational effort is found to be around 5 times more than a traditional Bayesian optimization. For example, 1000 steps of optimizations in Sec. III C took about 2 h for Bayesian optimization with binomial modelling, whereas the corresponding optimization with Gaussian finished within half an hour. This trade-off between the use of cheap classical computational resources versus expensive quantum mechanical ones can be tilted further in favor of one or the other. Resorting to more accurate estimates of the landscape (see, e.g., Ref. [67] for a review) is likely to increase the efficiency with the quantum hardware at the expense of classical computations. On the other hand, the field of probabilistic machine learning [80] provides fast approximate methods that could be incorporated into the framework. For example, probabilistic neural networks [81–83] and stochastic variational inference techniques [84], which could be used for both the Gaussian and binomial versions of Bayesian optimization, hold the promise of increased efficiency in classical calculations, but this might imply

slightly slower convergence towards the optimal control solution.

The applicability of the proposed methodology is by no means restricted to state preparation, but includes, for example, the optimization of quantum gates in the presence of uncharacterized noise, or the direct realization of few qubit gates avoiding decomposition into more elementary gates in order to realize more complex quantum algorithms within limited coherence times. Whereas optimization of gate fidelities will likely remain limited to systems comprised of few qubits, there are also direct applications with large qubit registers. One example is the variational quantum eigensolver [48,85] aimed at finding the quantum state that minimizes an energy functional; the present method allows one to perform the optimization over a wide range of states without the need to accurately estimate their energy or the gradient of the landscape [86,87] during the search for the optimal state. With the identification of ground states as a promising route towards the solutions of many classical problems such as the traveling salesman problem, the present methodology holds the potential to substantially advance the practical value of quantum systems for real-life applications.

The improvement in performance shown here for the case of shot noise in projective measurements is also not necessarily limited to this specific type of noise, but similar improvement should be expected for different types of noise, such as environmental noise or noise in the actual control fields if their spectral properties or temporal correlations are taken into account appropriately.

With probabilistic modeling and Bayesian optimization being active fields in mathematics and computer science, one should expect that the applicability and performance of similar methodologies will rapidly increase beyond what has been demonstrated in this work.

ACKNOWLEDGMENT

We are indebted to stimulating discussions with Annie Pichery, Dominic Lennon, Robert Löw, Hongzheng Zhao, Jake Lishman, Ulrich Schneider, and Zoë Holmes. The project Theory-Blind Quantum Control *TheBlinQC* has received funding from the QuantERA ERA-NET Cofund in Quantum Technologies implemented within the European Union’s Horizon 2020 Programme and from the EPSRC under Grant No. EP/R044082/1. This work is supported through a studentship in the Quantum Systems Engineering Skills and Training Hub at Imperial College London funded by EPSRC (Grant No. EP/P510257/1).

-
- [1] H. Katori, Optical lattice clocks and quantum metrology, *Nat. Photonics* **5**, 203 (2011).
 - [2] C. Monroe and J. Kim, Scaling the ion trap quantum processor, *Science* **339**, 1164 (2013).

- [3] R. Barends, *et al.*, Superconducting quantum circuits at the surface code threshold for fault tolerance, *Nature* **508**, 500 (2014).
- [4] J. M. Gambetta, J. M. Chow, and M. Steffen, Building logical qubits in a superconducting quantum computing system, *npj Quantum Inf.* **3**, 2 (2017).
- [5] R. S. Judson and H. Rabitz, Teaching Lasers to Control Molecules, *Phys. Rev. Lett.* **68**, 1500 (1992).
- [6] C. Ferrie and O. Moussa, Robust and efficient in situ quantum control, *Phys. Rev. A* **91**, 052306 (2015).
- [7] J. Li, X. Yang, X. Peng, and C.-P. Sun, Hybrid Quantum-Classical Approach to Quantum Optimal Control, *Phys. Rev. Lett.* **118**, 150503 (2017).
- [8] D. Lu, K. Li, J. Li, H. Katiyar, A. J. Park, G. Feng, T. Xin, H. Li, G. Long, A. Brodutch, J. Baugh, B. Zeng, and R. Laflamme, Enhancing quantum control by bootstrapping a quantum processor of 12 qubits, *npj Quantum Inf.* **3**, 45 (2017).
- [9] B. Dive, A. Pitchford, F. Mintert, and D. Burgarth, In situ upgrade of quantum simulators to universal computers, *Quantum* **2**, 80 (2018).
- [10] G. Feng, F. H. Cho, H. Katiyar, J. Li, D. Lu, J. Baugh, and R. Laflamme, Gradient-based closed-loop quantum optimal control in a solid-state two-qubit system, *Phys. Rev. A* **98**, 052341 (2018).
- [11] S. Rosi, A. Bernard, N. Fabbri, L. Fallani, C. Fort, M. Inguscio, T. Calarco, and S. Montangero, Fast closed-loop optimal control of ultracold atoms in an optical lattice, *Phys. Rev. A* **88**, 021601 (2013).
- [12] S. van Frank, M. Bonneau, J. Schmiedmayer, S. Hild, C. Gross, M. Cheneau, I. Bloch, T. Pichler, A. Negretti, T. Calarco, and S. Montangero, Optimal control of complex atomic quantum systems, *Sci. Rep.* **6**, 34187 (2016).
- [13] J. Kelly, *et al.*, Optimal Quantum Control Using Randomized Benchmarking, *Phys. Rev. Lett.* **112**, 240504 (2014).
- [14] F. Poggiali, P. Cappellaro, and N. Fabbri, Optimal Control for One-Qubit Quantum Sensing, *Phys. Rev. X* **8**, 021059 (2018).
- [15] D. J. Egger and F. K. Wilhelm, Adaptive Hybrid Optimal Quantum Control for Imprecisely Characterized Systems, *Phys. Rev. Lett.* **112**, 240503 (2014).
- [16] E. Zahedinejad, S. Schirmer, and B. C. Sanders, Evolutionary algorithms for hard quantum control, *Phys. Rev. A* **90**, 032310 (2014).
- [17] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, Reinforcement Learning in Different Phases of Quantum Control, *Phys. Rev. X* **8**, 031086 (2018).
- [18] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, *Nat. Commun.* **9**, 4812 (2018).
- [19] A. Assion, T. Baumert, M. Bergt, T. Brixner, B. Kiefer, V. Seyfried, M. Strehle, and G. Gerber, Control of chemical reactions by feedback-optimized phase-shaped femtosecond laser pulses, *Science* **282**, 919 (1998).
- [20] P. Palittapongrampim, P. Wittek, E. Zahedinejad, S. Vedaie, and B. C. Sanders, Learning in quantum control: High-dimensional global optimization for noisy quantum dynamics, *Neurocomputing* **268**, 116 (2017).
- [21] A. A. Melnikov, H. Poulsen Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H. J. Briegel, Active learning machine learns to create new quantum experiments, *Proc. Natl. Acad. Sci. USA* **115**, 1221 (2018).
- [22] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, Reinforcement Learning with Neural Networks for Quantum Feedback, *Phys. Rev. X* **8**, 031084 (2018).
- [23] M. August and J. M. Hernández-Lobato, in *International Conference on High Performance Computing* (Springer, 2018), p. 591.
- [24] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven, Universal quantum control through deep reinforcement learning, *npj Quantum Inf.* **5**, 33 (2019).
- [25] E. Brochu, V. M. Cora, and N. De Freitas, A tutorial on Bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning, arXiv:1012.2599.
- [26] J. Snoek, H. Larochelle, and R. P. Adams, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Red Hook, 2012), p. 2951.
- [27] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, Taking the human out of the loop: A review of Bayesian optimization, *Proc. IEEE* **104**, 148 (2015).
- [28] P. I. Frazier, A tutorial on Bayesian optimization, arXiv:1807.02811.
- [29] P. B. Wigley, P. J. Everitt, A. van den Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn, I. R. Petersen, A. N. Luiten, J. J. Hope, N. P. Robins, and M. R. Hush, Fast machine-learning online optimization of ultra-cold-atom experiments, *Sci. Rep.* **6**, 25890 (2016).
- [30] J. S. Otterbach, *et al.*, Unsupervised machine learning on a hybrid quantum computer, arXiv:1712.05771.
- [31] B. M. Henson, D. K. Shin, K. F. Thomas, J. A. Ross, M. R. Hush, S. S. Hodgman, and A. G. Truscott, Approaching the adiabatic timescale with machine learning, *Proc. Nat. Acad. Sci. USA* **115**, 13216 (2018).
- [32] D. Zhu, N. M. Linke, M. Benedetti, K. A. Landsman, N. H. Nguyen, C. H. Alderete, A. Perdomo-Ortiz, N. Korda, A. Garfoot, C. Breche, L. Egan, O. Perdomo, and C. Monroe, Training of quantum circuits on a hybrid quantum computer, *Sci. Adv.* **5**, eaaw9918 (2019).
- [33] I. Nakamura, A. Kanemura, T. Nakaso, R. Yamamoto, and T. Fukuhara, Non-standard trajectories found by machine learning for evaporative cooling of 87 Rb atoms, *Optics express* **27**, 20435 (2019).
- [34] R. Mukherjee, F. Sauvage, H. Xie, R. Loew, and F. Mintert, Preparation of ordered states in ultra-cold gases using Bayesian optimization, *New J. Phys.* **22**, 075001 (2020).
- [35] A. P. Peirce, M. A. Dahleh, and H. Rabitz, Optimal control of quantum-mechanical systems: Existence, numerical approximation, and applications, *Phys. Rev. A* **37**, 4950 (1988).
- [36] J. Werschnik and E. K. U. Gross, Quantum optimal control theory, *J. Phys. B* **40**, R175 (2007).
- [37] S. J. Glaser, U. Boscain, T. Calarco, C. P. Koch, W. Köckenberger, R. Kosloff, I. Kuprov, B. Luy, S. Schirmer, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, Training Schrödinger's cat: Quantum optimal control, *Eur. Phys. J. D* **69**, 279 (2015).

- [38] T. Caneva, T. Calarco, and S. Montangero, Chopped random-basis quantum optimization, *Phys. Rev. A* **84**, 022326 (2011).
- [39] F. Dolde, V. Bergholm, Y. Wang, I. Jakobi, B. Naydenov, S. Pezzagna, J. Meijer, F. Jelezko, P. Neumann, T. Schulte-Herbrüggen, J. Biamonte, and J. Wrachtrup, High-fidelity spin entanglement using optimal control, *Nat. Commun.* **5**, 3371 (2014).
- [40] L. Bianchi, N. Pancotti, and S. Bose, Quantum gate learning in qubit networks: Toffoli gate without time-dependent control, *npj Quantum Inf.* **2**, 16019 (2016).
- [41] F. Platzter, F. Mintert, and A. Buchleitner, Optimal Dynamical Control of Many-Body Entanglement, *Phys. Rev. Lett.* **105**, 020501 (2010).
- [42] J. Cui and F. Mintert, Robust control of long-distance entanglement in disordered spin chains, *New J. Phys.* **17**, 093014 (2015).
- [43] T. Gefen, F. Jelezko, and A. Retzker, Control methods for improved fisher information with quantum sensing, *Phys. Rev. A* **96**, 032310 (2017).
- [44] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, *J. Mag. Res.* **172**, 296 (2005).
- [45] T. Nöbauer, A. Angerer, B. Bartels, M. Trupke, S. Rotter, J. Schmiedmayer, F. Mintert, and J. Majer, Smooth Optimal Quantum Control for Robust Solid-State Spin Magnetometry, *Phys. Rev. Lett.* **115**, 190801 (2015).
- [46] S. Machnes, E. Assémat, D. Tannor, and F. K. Wilhelm, Tunable, Flexible, and Efficient Optimization of Control Pulses for Practical Qubits, *Phys. Rev. Lett.* **120**, 150401 (2018).
- [47] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. T. Wang, S. Ebadi, H. Bernien, A. S. Zibrov, H. Pichler, S. Choi, J. Cui, M. Rossignolo, P. Rembold, S. Montangero, T. Calarco, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Generation and manipulation of Schrödinger cat states in Rydberg atom arrays, *Science* **365**, 570 (2019).
- [48] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, A variational eigenvalue solver on a photonic quantum processor, *Nat. Commun.* **5**, 4213 (2014).
- [49] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, arXiv:1411.4028.
- [50] C. K. Williams and C. E. Rasmussen, *Gaussian Processes for Machine Learning* (MIT Press Cambridge, MA, 2006), Vol. 2.
- [51] IBM, The IBM quantum experience, <http://www.research.ibm.com/quantum> (2016).
- [52] J. C. Spall, Multivariate stochastic approximation using a simultaneous perturbation gradient approximation, *IEEE Trans. Automat. Contr.* **37**, 332 (1992).
- [53] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets, *Nature* **549**, 242 (2017).
- [54] R. J. Chapman, C. Ferrie, and A. Peruzzo, Experimental Demonstration of Self-Guided Quantum Tomography, *Phys. Rev. Lett.* **117**, 040402 (2016).
- [55] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature* **567**, 209 (2019).
- [56] J. Johansson, P. Nation, and F. Nori, Qutip 2: A python framework for the dynamics of open quantum systems, *Comput. Phys. Commun.* **184**, 1234 (2013).
- [57] P. Weinberg and M. Bukov, QuSpin: A python package for dynamics and exact diagonalisation of quantum many body systems part I: Spin chains, *SciPost Phys.* **2**, 003 (2017).
- [58] GPy, GPy: A gaussian process framework in python, <http://github.com/SheffieldML/GPy> (2012).
- [59] The GPyOpt authors, Gpyopt: A Bayesian optimization framework in python, <http://github.com/SheffieldML/GPyOpt> (2016).
- [60] IBM, Qiskit: An open-source framework for quantum computing.
- [61] C. F. Roos, M. Riebe, H. Häffner, W. Hänsel, J. Benhelm, G. P. T. Lancaster, C. Becher, F. Schmidt-Kaler, and R. Blatt, Control and measurement of three-qubit entangled states, *Science* **304**, 1478 (2004).
- [62] D. Leibfried, M. D. Barrett, T. Schaetz, J. Britton, J. Chiaverini, W. M. Itano, J. D. Jost, C. Langer, and D. J. Wineland, Toward Heisenberg-limited spectroscopy with multiparticle entangled states, *Science* **304**, 1476 (2004).
- [63] K. Wright, *et al.*, Benchmarking an 11-qubit quantum computer, *Nat. Commun.* **10**, 5464 (2019).
- [64] L. DiCarlo, M. D. Reed, L. Sun, B. R. Johnson, J. M. Chow, J. M. Gambetta, L. Frunzio, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, Preparation and measurement of three-qubit entanglement in a superconducting circuit, *Nature* **467**, 574 (2010).
- [65] M. Neeley, R. C. Bialczak, M. Lenander, E. Lucero, M. Mariantoni, A. D. O'Connell, D. Sank, H. Wang, M. Weides, J. Wenner, Y. Yin, T. Yamamoto, A. N. Cleland, and J. M. Martinis, Generation of three-qubit entangled states using superconducting phase qubits, *Nature* **467**, 570 (2010).
- [66] F. Arute, *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [67] H. Nickisch and C. Rasmussen, Approximations for binary gaussian process classification, *J. Mach. Learn. Res.* **9**, 2035 (2008).
- [68] M. Greiner, O. Mandel, T. Esslinger, T. W. Hänsch, and I. Bloch, Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms, *Nature* **415**, 39 (2002).
- [69] M. Schreiber, S. S. Hodgman, P. Bordia, H. P. Lüschen, M. H. Fischer, R. Vosk, E. Altman, U. Schneider, and I. Bloch, Observation of many-body localization of interacting fermions in a quasirandom optical lattice, *Science* **349**, 842 (2015).
- [70] H. Labuhn, D. Barredo, S. Ravets, S. de Léséleuc, T. Macri, T. Lahaye, and A. Browaeys, Tunable two-dimensional arrays of single Rydberg atoms for realizing quantum Ising models, *Nature* **534**, 667 (2016).
- [71] D. S. Weiss and M. Saffman, Quantum computing with neutral atoms, *Phys. Today* **70**, 44 (2017).

- [72] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletic, and M. D. Lukin, Probing many-body dynamics on a 51-atom quantum simulator, *Nature* **551**, 579 (2017).
- [73] T. L. Nguyen, J. M. Raimond, C. Sayrin, R. Cortiñas, T. Cantat-Moltrecht, F. Assemat, I. Dotsenko, S. Gleyzes, S. Haroche, G. Roux, T. Jolicoeur, and M. Brune, Towards Quantum Simulation with Circular Rydberg Atoms, *Phys. Rev. X* **8**, 011032 (2018).
- [74] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, A quantum gas microscope for detecting single atoms in a Hubbard-regime optical lattice, *Nature* **462**, 74 (2009).
- [75] A. Sergeevich, A. Chandran, J. Combes, S. D. Bartlett, and H. M. Wiseman, Characterization of a qubit hamiltonian using adaptive measurements in a fixed basis, *Phys. Rev. A* **84**, 052315 (2011).
- [76] 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).
- [77] A. S. Holevo, *Probabilistic and Statistical Aspects of Quantum Theory* (Springer Science & Business Media, Berlin, 2011), Vol. 1.
- [78] V. Giovannetti, S. Lloyd, and L. Maccone, Advances in quantum metrology, *Nat. Photonics* **5**, 222 (2011).
- [79] C. L. Degen, F. Reinhard, and P. Cappellaro, Quantum sensing, *Rev. Mod. Phys.* **89**, 035002 (2017).
- [80] Z. Ghahramani, Probabilistic machine learning and artificial intelligence, *Nature (London)* **521**, 452 EP (2015).
- [81] J. T. Springenberg, A. Klein, S. Falkner, and F. Hutter, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Red Hook, 2016), p. 4134.
- [82] B. Lakshminarayanan, A. Pritzel, and C. Blundell, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Red Hook, 2017), p. 6402.
- [83] A. D. Tranter, H. J. Slatyer, M. R. Hush, A. C. Leung, J. L. Everett, K. V. Paul, P. Vernaz-Gris, P. K. Lam, B. C. Buchler, and G. T. Campbell, Multiparameter optimisation of a magneto-optical trap using deep learning, *Nat. Commun.* **9**, 4360 (2018).
- [84] J. Hensman, A. Matthews, and Z. Ghahramani, Scalable variational gaussian process classification, *Proc. Mach. Learn. Res.* **38** (351–2015),
- [85] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, The theory of variational hybrid quantum-classical algorithms, *New J. Phys.* **18**, 023023 (2016).
- [86] J. M. Kübler, A. Arrasmith, L. Cincio, and P. J. Coles, An adaptive optimizer for measurement-frugal variational algorithms, *Quantum* **4**, 263 (2020).
- [87] R. Sweke, F. Wilde, J. Meyer, M. Schuld, P. K. Fährmann, B. Meynard-Piganeau, and J. Eisert, Stochastic gradient descent for hybrid quantum-classical optimization, *Quantum* **4**, 314 (2020).