


Continuous Quantum Gate Sets and Pulse-Class Meta-Optimization

Francesco Preti^{1,2,*}, Tommaso Calarco^{1,2} and Felix Motzoi^{1,†}

¹*Forschungszentrum Jülich, Institute of Quantum Control (PGI-8), Jülich D-52425, Germany*

²*Institute for Theoretical Physics, University of Cologne, Köln, D-50937, Germany*

 (Received 11 May 2022; revised 20 August 2022; accepted 6 September 2022; published 25 October 2022)

Reduction of the circuit depth of quantum circuits is a crucial bottleneck to enabling quantum technology. This depth is inversely proportional to the number of available quantum gates that have been synthesized. Moreover, quantum gate-synthesis and control problems exhibit a vast range of external parameter dependencies, both physical and application specific. In this paper, we address the possibility of learning families of optimal-control pulses that depend adaptively on various parameters, in order to obtain a global optimal mapping from the space of potential parameter values to the control space and hence to produce continuous classes of gates. Our proposed method is tested on different experimentally relevant quantum gates and proves capable of producing high-fidelity pulses even in the presence of multiple variables or uncertain parameters with wide ranges.

DOI: [10.1103/PRXQuantum.3.040311](https://doi.org/10.1103/PRXQuantum.3.040311)

I. INTRODUCTION

The standard view of quantum computation [1] uses the classical-computing abstraction of a subdivision into a finite set of gates, measurements, and state reset tasks. This paradigm has a number of benefits: notably, it permits formal derivation of universal computation [2,3], that is, the composition of a quantum circuit for any desired unitary operation, as well as error-correcting codes [4,5], where specific error syndromes can be measured and corrected. On the other hand, this abstraction abandons the essential analog character of quantum devices, from which they have the most to gain or lose in terms of their expressibility or fragility, respectively. That is, the power of quantum processors depends strongly on the number of usable gates available to them.

Practically speaking, the usage of discrete gate sets falls short in at least four important respects. First, the analog character is a more complete (and therefore efficient) description of variability between different qubits, which is inevitable, for instance, in solid-state qubits [6]. The use of qubit-agnostic gate sets as the computational primitive means that each qubit must be individually optimized and calibrated to yield each such gate [6], where parametric

description of the gates would naturally capture the variations. Second, these variations can occur for a given qubit as a function of time [7], e.g., due to time-dependent noise, and require complete recalibrations where, typically, drift involves only a single parameter. This is very taxing on the routines for error suppression, mitigation, and correction.

Third, the complexity arising from parameter variations in space and time is exacerbated by the subsequent circuit complexity in composing useful circuits. It is well known that while discrete gate sets can be universal, the required number of discrete constituent gates can be very large even for a simple circuit [3,8,9]. Allowing for analog parameter tuning can dramatically increase the controllability of the system and thereby necessarily reduce the depth of the quantum circuit for arbitrary tasks (see, e.g., Appendix C). Since errors accrue with circuit depth (notably due to decoherence), this increase in the circuit success probability may be highly beneficial to both short-term and long-term (i.e., fault-tolerant) approaches.

Fourth, a further optimization layer is currently ubiquitous in noisy intermediate-state quantum (NISQ) algorithms [10]. Such optimizable circuits include adiabatic [11,12], annealing [13,14], and variational [15,16] quantum algorithms. The common denominator in these approaches is the circuit being treated as a black box, with a set of analog parameters acting as knobs to tune as inputs for the respective algorithm. These analog inputs act as terms in the Hamiltonian and thus may generate various quantum continuous gate sets. These cannot realistically be compiled with digital gates and, moreover, to ensure that the gate set is correctly specified requires a formal approach for their general construction.

*f.preti@fz-juelich.de

†f.motzoi@fz-juelich.de

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.

Our contribution, in this work, is to present a unified framework to efficiently describe and optimize continuous quantum gate sets in these scenarios. This framework allows learning of the parametric gates that can be tuned to very high fidelity across a large number and wide range of parameter values. We refer to our method as single-optimization multiple-application (SOMA) quantum gate synthesis. This kind of learning can be understood as an instance of meta-optimization [17,18] or adaptive-trajectory learning [19–25] and can be related to more recent uses of neural networks in quantum physics [9,26–30]. That is, we find a solution for an optimizer that itself provides solutions to specific problem instances or, with a different formulation, automatically discovers heuristics to construct solutions for a specific optimization problem. In particular, we see that our algorithms are able to synthesize heuristics for general Hamiltonians.

We break down the problem of generalized gate synthesis into the following components. We parametrize our quantum gate set using continuous indices that represent either physical system parameters or desired angles of a continuous Lie group. We then present two different machine-learning methods for obtaining continuous control parametrizations that generate the indexed gate set. The first is a supervised approach where traditional optimal-control theory is used to generate an operational data set from which a generalized gate-set recipe can be trained. The second is an unsupervised approach using back propagation, from which the continuous gate set can be incrementally learned over the entire training population. Finally, we show that our approaches encompass the various situations discussed above, including general solutions for gates given generic physical architectures with wide parameter ranges, noise-adaptive optimal-control theory, and compilation of a Lie group instead of a single element.

The paper is organized as follows. In Sec. II, we introduce the notation for supervised and unsupervised training of parametrized pulses. In Sec. III, we discuss the results obtained by applying these methods to single-qubit and two-qubit gates in the presence of leakage, showing that they show similarities with known analytical solution families. Furthermore, we also investigate how our methods perform compared to other existing algorithms. Finally, in Sec. IV we analyze the dependence on the parameter variability range, the training-data-set size and batch size, and the system and network size, and again compare our algorithms to other numerical robust approaches. We summarize our conclusions in Sec. V.

II. CONTINUOUS GATE-SET LEARNING

A. Definitions

We define a continuous gate set in terms of some continuous sets or distributions of n system parameters

s_1, s_2, \dots, s_n and m gate specifications $\theta_1, \theta_2, \dots, \theta_m$. An element of such a gate set is a unitary transformation between two Hilbert spaces H_1 and H_2 :

$$U(s_1, s_2, \dots, s_n, \theta_1, \theta_2, \dots, \theta_m) : H_1 \mapsto H_2. \quad (1)$$

Obtaining a single element of this set is a well-known problem in quantum information. Depending on whether the unitary is synthesized from discrete or analog dynamics, its composition is referred to as circuit compilation [2,3,37,38] or optimal-control theory [39–44], respectively.

In Fig. 1(a), we see an example of a generic circuit acting on different qubits. Each qubit in space (and time) will have different values of the common system parameters $\{s_k\}$. In addition, the different unitaries in the gate set $\{U_i\}$, each additionally characterized by rotation angles $\{\theta_{ij}\}$, may vary both throughout the circuit and in iterated uses of the circuit (e.g., in NISQ algorithms). For compactness, we now regroup the continuous indices into a common array of indices $\vec{\lambda} = (s_1, s_2, \dots, s_n, \theta_1, \theta_2, \dots, \theta_m)$.

Such a generalization of the gate-synthesis problem can be framed as solving for the inverse function of the general dynamics given in Eq. (1), that is, for

$$g : \vec{\lambda} \mapsto \mathbf{u}(t), \quad (2)$$

where $\mathbf{u}(t)$ is a wave-form function in the standard case of optimal-control theory but can also be thought of as a discrete sequence of hard pulses, as in NMR applications, or of unitary gates in circuit compilation. Importantly, the optimal $\mathbf{u}(t)$ changes for each parametrization $\vec{\lambda}$ of the unitary, which can be generated from, e.g., the Schrödinger equation

$$\dot{U}(\vec{\lambda}, t) = -iH(\vec{\lambda}, \mathbf{u}(t))U(\vec{\lambda}, \mathbf{u}(t)), \quad (3)$$

or other equations of motion defining the system. This task can be cast as an instance of meta-optimization [45] or trajectory learning for control [24,46,47]. We formalize the meta-optimization problem as a list of problem-parameter vectors $\vec{\lambda}_1, \dots, \vec{\lambda}_L \in V_\lambda \subset \mathbb{R}^D$ with figures of merit $F_1(\mathbf{x}) = F(\mathbf{x}, \vec{\lambda}_1), \dots, F_L(\mathbf{x}) = F(\mathbf{x}, \vec{\lambda}_L)$ and initial guess \mathbf{x}_0 , the physical parameters of which vary somewhat from each other, such that $\vec{\lambda}_i \sim \pi(\vec{\lambda}|\mathbf{v})$ for $0 \leq i \leq L$ and $\mathbf{v} \in V_\lambda$, is drawn from a parameter distribution. Throughout the paper we assume, without loss of generality, π to be a multidimensional uniform distribution, such that $\pi(\vec{\lambda}|\mathbf{v}) = U(\vec{\lambda}_{\min}, \vec{\lambda}_{\max})$, with $\mathbf{v} = (\vec{\lambda}_{\min}, \vec{\lambda}_{\max})$ defining the parameter space V_λ .

The objective is to find optimal parameters \mathbf{w}^* ,

$$\mathbf{w}^* = \underset{\mathbf{w}}{\operatorname{argmin}} \{1 - F_i(\mathbf{w}) \mid \forall i = 1, \dots, L\}, \quad (4)$$

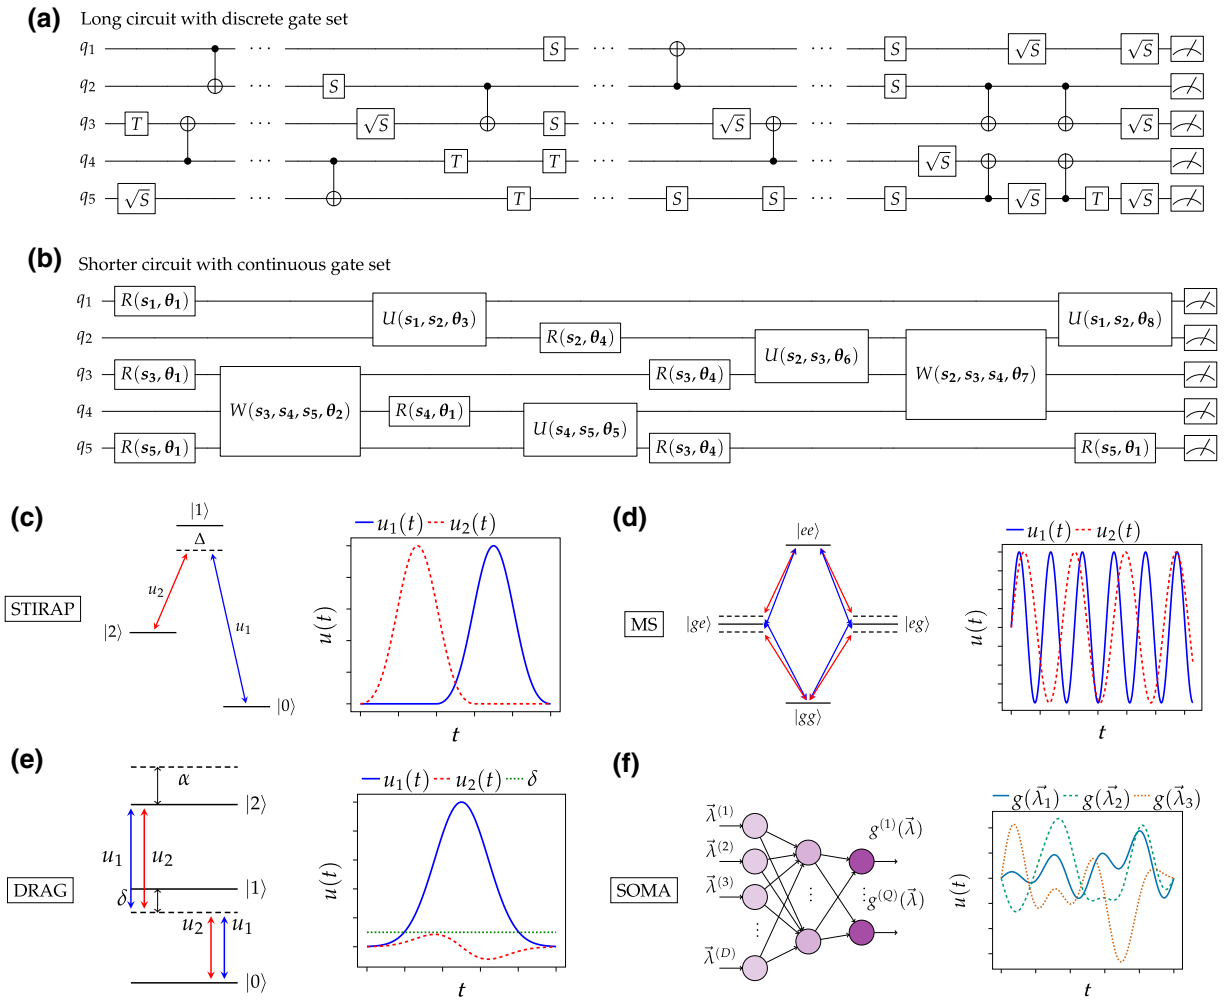


FIG. 1. An explanatory diagram of the concepts discussed in Sec. II. (a) A general quantum circuit containing a long sequence of discrete unitaries (Clifford + T set), which do not exhibit any dependence on continuous parameters. (b) A general quantum circuit containing different unitaries with different continuous parameters s_1, s_2, \dots, s_5 for different qubits q_1, q_2, \dots, q_5 and angles $\theta_1, \theta_2, \dots, \theta_8$ parametrizing each gate. R , U , and W represent an analog single-, two- and three-qubit gate, respectively. Here, the same unitary parametrization is capable of representing all the different gates needed in the circuit. We show some notable analytical solutions used to engineer specific gate operations, which are usually implemented due to their adaptive character and simplicity [see Eqs. (5), (6), and (7)]. (c) STIRAP [31,32]. (d) The Mølmer-Sørensen gate [33,34]. (e) DRAG [35,36]. (f) The method that we propose, SOMA, does not make strict assumptions on how the pulse depends on the problem parameters but, rather, discovers it more generally through training.

which allow for simultaneous optimization of all the systems considered within the range of sampled parameters.

B. Analytical adaptive control

The most straightforward way to generate classes of solutions to quantum gates has been the development of analytical solutions for particular quantum systems.

They have in common the knowledge of the relevant state of the system at all times during the evolution. In particular, analytical knowledge of the eigenvalues [48] allows reverse engineering of the pulses to provide (near) exact solutions for the desired states.

Figures 1(c)–1(f) shows several celebrated examples where such general classes of solutions have been found. We quickly review some of their main features. Given a trial pulse shape such as a Gaussian envelope $p(t, t_1, t_2, \theta) = \mathcal{A}e^{-(t-t_2/2+t_1/2)^2/\sigma^2}$, where θ denotes the area under the curve, the following dynamical solutions have been found.

Figure 1(c) shows the stimulated-rapid-adiabatic-passage (STIRAP) solution for transferring population between disconnected states $|0\rangle$ and $|2\rangle$, while avoiding any (nonvanishing) temporary population in the intermediary connecting state $|1\rangle$. The pulse shaping is given

by

$$\begin{aligned} \text{STIRAP}(|0\rangle \rightarrow |2\rangle) : (\Omega_1, \Omega_2, \Delta, \theta, T) &\mapsto \mathbf{u}(t), \\ u_1(t) &= \Omega_1 e^{i\theta} p(t, T/3, T, \pi) e^{i\Delta t} |0\rangle\langle 1|, \\ u_2(t) &= \Omega_2 p(t, 0, 2T/3, \pi) e^{-i\Delta t} |1\rangle\langle 2|, \end{aligned} \quad (5)$$

where the order of the pulses u_1 and u_2 is famously counterintuitive [31,32].

Figure 1(d) shows the Mølmer-Sørensen (MS) method for trapped-ion gates [33,34]. The required laser pulses are given by

$$\begin{aligned} \text{MS}(\sigma_x \otimes \sigma_x \otimes \dots \otimes \sigma_x) : (\Omega_1, \Omega_2, \delta, \theta_1, \theta_2) &\mapsto \mathbf{u}(t), \\ u_1(t) &= \Omega_1 p(t, 0, T, \theta_1) e^{i\theta_2} e^{i\delta t} |g, n\rangle\langle e, n-1|, \\ u_2(t) &= \Omega_2 p(t, 0, T, \theta_1) e^{-i\delta t} |g, n\rangle\langle e, n+1|, \end{aligned} \quad (6)$$

where the first state index denotes the state of the relevant qubit and second index denotes the phonon occupation number. Because of the symmetry of the gate, it can in theory be used on an arbitrarily large number of qubits, i.e., it is a collective gate [38,49].

Figure 1(e) shows the derivative-removal-for-adiabatic-gate (DRAG) solution, for leakage suppression in multi-level systems [35,36]. This pulse profile is given by

$$\begin{aligned} \text{DRAG}(|0\rangle \leftrightarrow |1\rangle) : (\Omega_1, \Omega_2, \alpha, \delta, \theta_1, \theta_2) &\mapsto \mathbf{u}(t), \\ u_1(t) &= e^{i\theta_2} e^{i\delta t} p(t, 0, T, \theta_1) (\Omega_1 |0\rangle\langle 1| + \Omega_2 |1\rangle\langle 2|), \\ u_2(t) &= ie^{i\theta_2} e^{i\delta t} \partial_t p(t, 0, T, \theta_1) / \alpha (\Omega_1 |0\rangle\langle 1| \\ &\quad + \Omega_2 |1\rangle\langle 2|). \end{aligned} \quad (7)$$

We see, with these families of solutions, the common trend that they allow for different known system parameters or for different rotation or phase angles. Naturally, this is just a representative set but where the equations of motion are integrable, such solutions are numerous in the literature.

However, there are a few evident concerns about finding such analytical solutions. First, it is a labor-intensive task with no guaranteed result, where even particular solutions do not preclude the possibility that a more systematic search would produce better results. Second, it is important to emphasize that these are solutions to idealized models and in practice the more accurate physical models are not exactly solved by these ansätze. Finally, most physical systems have to date not been able to find general analytical solutions beyond qubits, qutrits, and highly symmetric systems, both because of the larger state space and the larger parameter space. This limits their viability for quantum computing, which requires much larger Hilbert spaces.

C. Gradient-ascent pulse engineering

Gradient-ascent pulse engineering (GRAPE) [41] is a method originally developed in the context of quantum chemistry for the optimization of dynamical evolution of NMR systems. The algorithm assumes the unitary dynamics $U(t) = e^{-\int_0^t iH(\tau)d\tau}$, governed by a Hamiltonian of type

$$H(t) = H_0 + \sum_{m=1}^M u_m(t) H_m, \quad (8)$$

where H_0 is a time-independent drift Hamiltonian and H_m , with $m = 1, \dots, M$, are different suitable control Hamiltonians with corresponding control fields $u_m(t)$. In simulations, $U(t)$ is often computed through different types of Trotterization [52]. GRAPE provides us with an efficient gradient of the merit function with respect to the control-pulse values. The merit function is usually given by the gate fidelity:

$$F = \frac{1}{d^2} |\text{Tr}\{UG^\dagger\}|^2, \quad (9)$$

where the normalization factor d corresponds to the dimension of the Hilbert space and G is a target unitary, which we would like to generate using the unitary dynamics. We assume a Trotter-like unitary evolution of the system of type

$$U(T) = \prod_{i=1}^{N_{\text{evo}}} U_j(t_j, t_{j-1}), \quad (10)$$

where $U_j(t_j, t_{j-1}) = e^{-iHdt}$, with $dt = T/N_{\text{evo}} = t_j - t_{j-1} \forall j = 1, \dots, N_{\text{evo}}$, in which N_{evo} defines the number of time steps used in the Trotterization. In particular, considering Eq. (8), the unitary step U_j reads

$$U_j = \exp \left\{ -idt \left(H_0 + \sum_k u_k(j) H_k \right) \right\}. \quad (11)$$

The gradient of the fidelity can be computed iteratively starting from the so-called propagated optimal state [41],

$$O_j = U_{j+1} \dots U_{N_{\text{evo}}} G U_1 \dots U_{j-1}, \quad (12)$$

so that the gradient approximately results in

$$\frac{\partial F}{\partial u_k(j)} \approx \frac{2}{d^2} \text{Re} \left\{ \text{Tr} \left\{ UG^\dagger \right\} \text{Tr} \left\{ idt H_k O_j \right\} \right\}. \quad (13)$$

This approach, however, cannot directly account for variations of the underlying dynamics due to, e.g., stochastic noise [53], field inhomogeneity [54], or Hamiltonian

uncertainties [55] and the optimal pulses output by following the native gradient direction can prove significantly worse than expected if some of the underlying problem parameters vary. A possible way around this is to switch to a robust-control approach, in which the cost function given in Eq. (9) accounts for parameter shifts. A simple way [41] is to use an average fidelity over the parameter space sampled with quasi-Monte Carlo,

$$\bar{F} = \frac{1}{L} \sum_{l=1}^L F(\mathbf{w}, \lambda_l), \quad (14)$$

where L is the number of samples.

D. Robust control

Similar to the adaptive solutions using analytical methods, solutions for controls robust to uncertainty in parameters have been found both by analytical and numerical means [27,40,53–59]. Since this involves only a single solution and not a family thereof, this has been the preferred method for parameter variability in quantum gate sets, as they are easier to design.

Robust solutions are generally defined slightly differently from the adaptive solutions, using the figure of merit

$$\mathbf{w}^* = \underset{\mathbf{w}}{\operatorname{argmin}} \left(1 - \frac{1}{L} \sum_{i=1}^L F_i(\mathbf{w}) \right). \quad (15)$$

This cost function is more tractable from an optimization point of view because we can simply account for an ensemble of individual cost functions by taking the average of the cost functions as the objective of the optimization.

Robust solutions have also been found using analytical methods. In particular, a common pulse sequence for gates with robustness to amplitude noise is the broadband 1 (BB1) sequence

$$\begin{aligned} \text{BB1}(|0\rangle \leftrightarrow |1\rangle) : (\Omega, \Delta\Omega, \theta_1, \theta_2) &\mapsto \mathbf{u}(t), \\ u_1(t) &= \Omega p(t, 0, T/4, \theta_1) e^{i\phi} |0\rangle\langle 1|, \\ u_2(t) &= \Omega p(t, T/4, T/2, \pi) e^{i\theta_2 + i \cos^{-1}(-\theta_1/4\pi)} |0\rangle\langle 1|, \\ u_3(t) &= \Omega p(t, T/2, 3T/4, 2\pi) \\ &\quad \times e^{i\theta_2 + i3 \cos^{-1}(-\theta_1/4\pi)} |0\rangle\langle 1|, \\ u_4(t) &= \Omega p(t, 3T/4, T, \pi) e^{i\theta_2 + i \cos^{-1}(-\theta_1/4\pi)} |0\rangle\langle 1|, \end{aligned} \quad (16)$$

which is independent of offsets in Rabi frequency $\Delta\Omega$ up to sixth order, $1 - F = O(\Delta\Omega^6)$ [60].

Likewise, when applying gates with unknown frequency offsets, the compensation for off-resonance with a pulse

sequence (CORPSE) method [61] given by

$$\begin{aligned} \text{CORPSE}(|0\rangle \leftrightarrow |1\rangle) : (\Omega, \Delta, \theta_1, \theta_2) &\mapsto \mathbf{u}(t), \\ u_1(t) &= \Omega e^{i\theta_2} p \left(t, 0, \frac{T}{3}, 2\pi + \frac{\theta_1}{2} - \sin^{-1} \left(\frac{\beta}{2} \right) \right) |0\rangle\langle 1|, \\ u_2(t) &= \Omega e^{i\theta_2} p \left(t, \frac{T}{3}, \frac{2T}{3}, 2\pi - 2 \sin^{-1} \left(\frac{\beta}{2} \right) \right) |0\rangle\langle 1|, \\ u_3(t) &= \Omega e^{i\theta_2} p \left(t, \frac{2T}{3}, T, \frac{\theta_1}{2} - \sin^{-1} \left(\frac{\beta}{2} \right) \right) |0\rangle\langle 1|, \end{aligned} \quad (17)$$

with $\beta = \sin(\theta/2)$, is robust to the exact value of Δ , for small enough Δ .

Robustness has found widespread use in quantum computation, where fabrication uncertainty, use of ensemble systems, and noise have made control challenging. The use of robust control is especially important where small deviations occur over time scales roughly on par with gate durations.

Nevertheless, if deviations are not small, if they are over a much longer (or much shorter) time scale, or if very high fidelity is sought after, then typically they have limited value. This is especially the case where variability occurs as a result of design uncertainty or slow parameter drift or when continuous gate sets are needed as for NISQ algorithms. To understand why maximum fidelities suffer as a result of improved robustness, note that a longer pulse sequence (with multiple pulses) will necessarily incur more decoherence. Thus while pulses such as BB1 and CORPSE will reduce drift error, the overall fidelity will not be as high as a single pulse at a fraction of the duration could have yielded. We also show this quantitatively in the following sections.

E. Supervised training method: SOMA SL

Rather than constructively producing such classes or relying solely on optimal-control theoretic methods, here we pursue the approach of machine learning a functional approximation to the general solutions. Function approximators are mathematical objects capable of reproducing arbitrary functions using families of functions [62]. They are normally identified with neural networks and find extensive application in machine learning, data analysis, etc.

For the supervised approach, which is sometimes referred to as trajectory learning in the robotic control literature [47], we employ an optimizer to find the corresponding minima for a set of problems and a regressor $g : \mathbb{R}^D \mapsto \mathbb{R}^Q$ that maps the problem-parameter space to the space of solutions to the given problem. We refer to this approach as SOMA with supervised learning (SOMA SL). A sketch of the algorithm is provided in Fig. 2(a). Starting from a seed problem with solution \mathbf{x}_0^* , generated

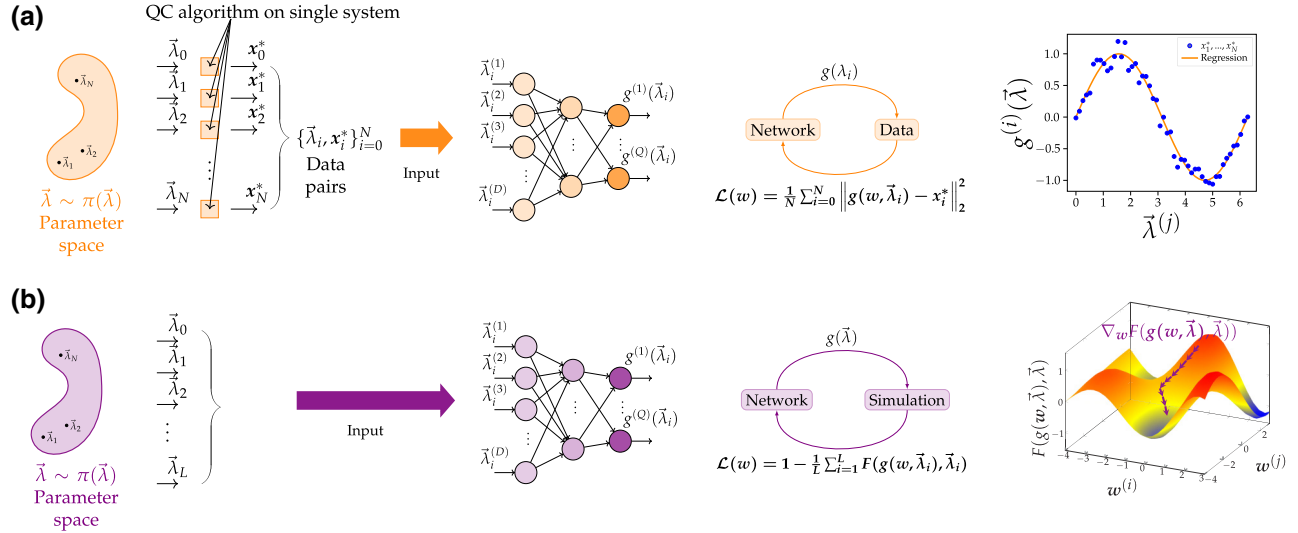


FIG. 2. The SOMA gate-synthesis method represented in its two variants. (a) In SOMA SL, we first optimize N different QOC problems with different problem parameters $\bar{\lambda}_1, \bar{\lambda}_2, \dots, \bar{\lambda}_N$ using an optimal quantum control algorithm such as in Refs. [41,42,50,51] to obtain optima $\mathbf{x}_1^*, \mathbf{x}_2^*, \dots, \mathbf{x}_N^*$ and then using a function approximator to learn the mapping $g: \mathbb{R}^D \mapsto \mathbb{R}^Q, \bar{\lambda} \rightarrow \mathbf{x}^*(\bar{\lambda})$ between the problem parameters and the optimal pulses. (b) In SOMA BP, we sample L QOC problems $\bar{\lambda}_1, \bar{\lambda}_2, \dots, \bar{\lambda}_L$ and train the function approximator to minimize the average infidelity of the ensemble of problems using back propagation, without generating optimal solutions for a single problem with a standard quantum control method.

previously, for N different optimal quantum control problems parametrized by $\bar{\lambda}_1, \dots, \bar{\lambda}_N$, we generate N solutions $\mathbf{x}_1^*, \dots, \mathbf{x}_N^*$. Then we train the neural network to find the best nonlinear mapping between the original parameters and the solutions. Training is performed via a standard mean-squared error (MSE) loss:

$$\mathcal{L}(\mathbf{w}) = \sum_{i=1}^N \|\mathbf{z}_i^* - g(\mathbf{w}, \bar{\lambda}_i)\|_2^2, \quad (18)$$

$$\mathbf{z}_i^* = \frac{\mathbf{x}_i^* - \bar{\mathbf{x}}}{\sigma_{\mathbf{x}}}, \quad (19)$$

where $\bar{\mathbf{x}}$ is the mean value of the generated data, $\sigma_{\mathbf{x}}$ is its standard deviation, and \mathbf{z}_i^* are the normalized data.

F. Direct-training method with back propagation: SOMA BP

An adaptive trajectory is a solution that depends on a specific parameter of the physical system, which the optimizer does not control, although it can make use of it. Usually, for many robotics applications, the system does not have an analytical model, thereby preventing direct-learning strategies. However, for quantum dynamics, we show how we can use the model to more directly train the function approximator. We refer to this approach as SOMA with back propagation (SOMA BP). A sketch of the algorithm is provided in Fig. 2(b).

Prime examples of adaptive trajectories are the analytical pulses in Sec. II B. For instance, by using a frequency-dependent solution, DRAG eliminates the leakage inside a qutrit. Moreover, this solution parametrized by the function approximator, g , directly depends on the physical system values and can therefore be tuned if these are shifted.

The cost function for an ensemble of L quantum-optimal-control (QOC) problems defined by problem parameters $\bar{\lambda}_1, \dots, \bar{\lambda}_L$ is given by

$$\mathcal{L}(\mathbf{w}) = 1 - \frac{1}{L} \sum_{l=1}^L F(g(\mathbf{w}, \bar{\lambda}_l), \bar{\lambda}_l), \quad (20)$$

where F , as before, is the figure of merit, e.g., the overlap fidelity of the operation with the target quantum gate [64].

By parametrizing the solution in terms of a neural network that depends on the gate parameters, gradient-based optimization algorithms can be used to train the network directly off the above cost function. In essence, the usual back propagation of neural networks matches naturally with gradient-descent optimal-control methods such as GRAPE [41]. Thus, while optimizing in the fidelity landscape of the controls, our algorithm is able to simultaneously train the network to adapt to the extraneous system and gate parameters. This method can also be used in combination with a more standard robust-GRAPE approach. In this case, those parameters the calibration and control of which proves difficult can be excluded from the network input.

G. Experimental adaptation

Direct experimental application of SOMA is possible both for model-based and model-free implementations. For these purposes, one must have access to a controlled distribution of $\vec{\lambda}$ values, corresponding to gate parameters, pulse parameters, and system parameters. For *in situ* optimization of the gates, one further requires access to an experimental cost function to ascertain (with low noise and bias) the merit of the pulse sequence. For model-free control learning, the system parameters should still be indexed in some way, e.g., by performing characterization, by using a proxy such as other known characteristics, or by externally tuning parameters (e.g., the qubit frequency via magnetic or Stark shifts). For model-based approaches, one can vary the parameters $\vec{\lambda}$ in software to map the solution space of the continuous gate sets. Thus, in both cases, provided that there is known variation in some parameters, then one can index them, e.g., discretely in space or (slowly varying) continuously in time, as shown in Figs. 3(a) and 3(b), respectively.

Whichever parameters for $\vec{\lambda}$ one chooses for the experiment, the task then becomes to learn the neural-network weights for the maximal performance on the relevant device and gate defined uniquely by $\vec{\lambda}$. Two different approaches are shown in Figs. 3(c) and 3(d). When an accurate model for the generators of the dynamics is known, then Fig. 3(d) is a natural choice whereby offline

(i.e., open-loop) optimization of the simulated gates is first brought up and only once the solution class has been learned *in situ* (i.e., closed-loop) is control learning performed. In this secondary step, one can reoptimize either over the space of solutions (fine tuning optimal $\vec{\lambda}^*$) or over the space of physical controls (fine tuning optimal \vec{x}^*).

Closed-loop optimization directly on the experiment can be performed a number of ways, including numerical and parameter-shift approximations of the gradient, Nelder-Mead [65], or evolutionary algorithms [66]. In the latter case, for example, Monte Carlo gradient sampling can be used to estimate the update direction [67]:

$$\nabla_{\mathbf{x}} F(\mathbf{x}) = \frac{1}{\tilde{N}} \sum_{i=1}^{\tilde{N}} F(\mathbf{x} + \epsilon_i \sigma) \epsilon_i, \quad (21)$$

where $\epsilon_i \sim \mathcal{N}(\mathbf{0}, I_Q)$, $i = 1, \dots, \tilde{N}$ is a normally distributed stochastic variable sampled \tilde{N} times, \mathbf{x} is the network output representing the pulse, and σ is the standard deviation of the sampled pulses. This method is often referred to as natural evolution strategy [67]. Figure 3(c) shows how the experimental gradient of the cost function can be then back propagated via the optimizer (similarly to the GRAPE implementation above) in order to update the network weights \mathbf{w}^* directly, when the different $\vec{\lambda}$ can be sampled simultaneously.

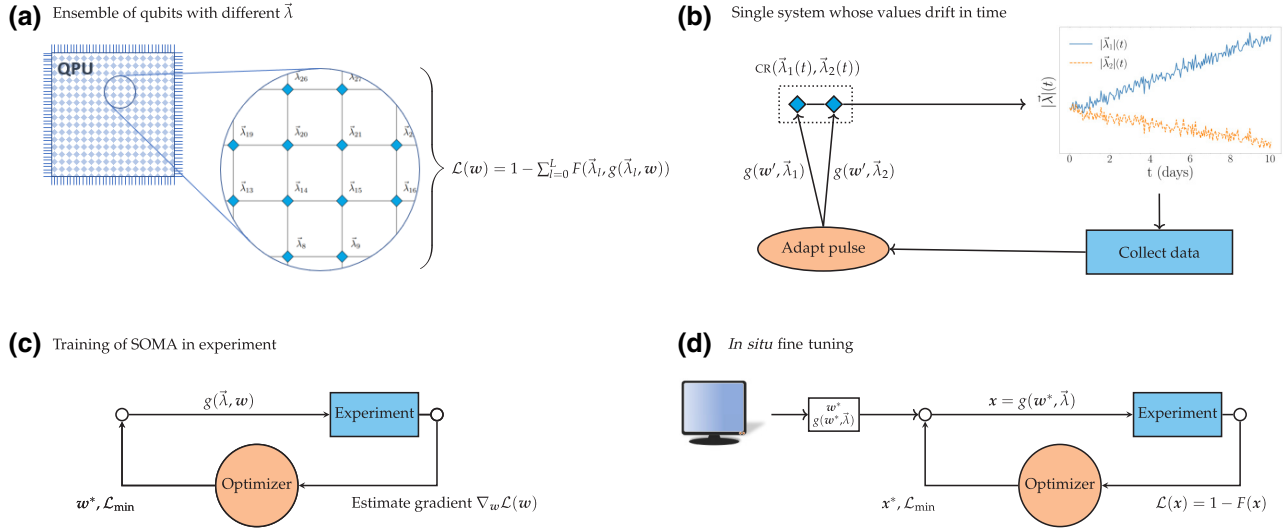


FIG. 3. Four diagrams describing two possible use cases of SOMA and two possible experimental implementations of SOMA BP, respectively. (a) We consider a chip with several qubits, each one with its own Hamiltonian parameter $\vec{\lambda}$. We assume parameter variations $\delta\vec{\lambda}$ to be so large that robust pulses are generally ineffective. (b) In the second use case, we consider a single system, the parameters of which vary with time. Pulses are trained on the average cost function over ensembles of qubits. (c) The approximator is trained directly on the experimental setting by using a gradient estimator, in this case a policy gradient [63], which usually requires large numbers of samples to be drawn from the system. (d) The regressor is first trained first in open-loop simulation and then used as an ansatz for a closed-loop optimization, which leaves the neural-network parameter untouched but modifies the output amplitude parameter (or, alternatively and if possible, the input parameters) to maximize the fidelity for specific experimental configurations. The optimization routine can be freely chosen among gradient-free algorithms, such as Nelder-Mead in Ref. [43]. This method can be a viable option if the experimental setting can be simulated with sufficient precision.

III. RESULTS

We test our methods and compare to previous optimal-control theoretic approaches. For this purpose, we train our solution networks to learn how to perform continuous gate sets for both single- and two-qubit operations. As a figure of merit of the QOC problems, we choose the gate fidelity defined in Eq. (9). While the gradients of the fidelity with respect to the pulse parameters x can be computed using GRAPE [41,68] (see Sec. II C and Appendix A), in the context of these simulations the gradient can also be obtained through automatic differentiation [69,70]. To simulate the quantum system, we use a second-order Magnus propagator as derived in Refs. [71,72] and which is compatible with analytical or automatic differentiation [72]. For the optimization of all the parameters, we employ the algorithm L-BFGS-B [73].

A. Single-qubit gates

Single-qubit gates are the fundamental building blocks of quantum circuits. Their most general form—e.g., as they are implemented in the IBMQ compiler [74]—is given by

$$R(\theta_1, \theta_2, \theta_3) = \begin{bmatrix} \cos(\theta_1/2) & -e^{i\theta_2} \sin(\theta_1/2) \\ e^{i\theta_3} \sin(\theta_1/2) & e^{i(\theta_3+\theta_2)} \cos(\theta_1/2) \end{bmatrix}. \quad (22)$$

By choosing the vector parameter θ appropriately, one can construct arbitrary single-qubit unitaries. Therefore, in any optimal quantum control problem, this triple can be considered as a vector of problem parameters, since they define the entire class of QOC problems, the goal of which is the optimization of arbitrary single-qubit gates.

In the following section, we consider an ensemble of QOC problems defined by parameters of the unitary target gate, Hamiltonian parameters, parameters of the control fields, and the evolution time T . To simplify the problem, we consider a target gate of type

$$R_1(\theta) = \begin{bmatrix} \cos(\theta/2) & \sin(\theta/2) \\ \sin(\theta/2) & -\cos(\theta/2) \end{bmatrix} \quad (23)$$

by setting $\theta_1 = \theta$, $\theta_2 = \pi$ and $\theta_3 = 0$ in Eq. (22). For $\theta = \pi/2$, the gate is the H gate, whereas for $\theta = \pi$ it produces the X gate.

The second family of unitaries that we consider can be obtained by setting $\theta_1 = \pi$ and $\theta_2 = \theta_3 = \theta - \pi$ in Eq. (22). The resulting family of gates,

$$R_2(\theta) = \begin{bmatrix} 0 & e^{i(\theta-\pi)} \\ e^{-i(\theta-\pi)} & 0 \end{bmatrix}, \quad (24)$$

generates, among other unitaries, the X gate for $\theta = \pi$ and the Y gate for $\theta = \pi/2$. We would like to point out that a combination of single-qubit unitary gates as in Eqs.

(24) and (23) is sufficient to generate arbitrary single-qubit unitaries [1].

For single-qubit simulations, we consider the Hamiltonian of a superconducting transmon qubit [75]. This system can be effectively reduced to a qutrit Hamiltonian [75,76], where the $|0\rangle$ and $|1\rangle$ levels provide the computational subspace and the $|2\rangle$ level represents the leakage. The drift Hamiltonian for our system [35] reads

$$H_d = \omega_d \hat{n} + \alpha \hat{\Pi}_2, \quad (25)$$

$$\hat{\Pi}_j = |j\rangle \langle j|, \quad (26)$$

$$\hat{n} = \sum_j j \hat{\Pi}_j, \quad (27)$$

where ω_d is the qubit frequency and α is the anharmonicity. Furthermore, we consider a control Hamiltonian of type

$$H_c(t) = \Omega(t) e^{i\phi + i\omega t} \hat{\sigma}_+ + \Omega(t)^* e^{-i\phi - i\omega t} \hat{\sigma}_-, \quad (28)$$

where ω is the driving frequency and ϕ represents a time-independent phase shift between the raising $\hat{\sigma}_+$ and the lowering $\hat{\sigma}_-$ operators, which can be related to the rotating-wave approximation (RWA) [68,77] and which is, in this model, the only problem parameter influencing the control fields.

Computing the RWA with detuning δ allows us to rewrite the drift Hamiltonian as

$$H_d = \delta \hat{\Pi}_1 + (\alpha - 2\delta) \hat{\Pi}_2, \quad (29)$$

$$H_c(t) = u_1(t) \hat{X}(\phi) + u_2(t) \hat{Y}(\phi), \quad (30)$$

where $\delta = \omega_d - \omega$, $\hat{X}(\phi) = e^{i\phi} \hat{\sigma}_+ + e^{-i\phi} \hat{\sigma}_-$, and $i\hat{Y}(\phi) = e^{i\phi} \hat{\sigma}_+ - e^{-i\phi} \hat{\sigma}_-$.

Input $\mathbf{w}, \vec{\lambda}_0, \vec{\lambda}_{\min}, \vec{\lambda}_{\max}$, optimizer OPT

Output \mathbf{w}^*

- 1: $\mathbf{x}^* = OPT(F(\mathbf{x}, \vec{\lambda}_0), \nabla_x F(\mathbf{x}, \vec{\lambda}_0), \mathbf{x}_0)$
 \triangleright with random restart
 - 2: **for** $i = 1$ to N **do**
 - 3: $\vec{\lambda}_i \sim \pi(\vec{\lambda}_{\min}, \vec{\lambda}_{\max})$
 - 4: $\mathbf{x}_i^* = OPT(F(\mathbf{x}, \vec{\lambda}_i), \nabla_x F(\mathbf{x}, \vec{\lambda}_i), \mathbf{x}^*)$
 - 5: **end for**
 - 6: Save $\{(\mathbf{x}_i, \vec{\lambda}_i)\}_{i=1}^N$
 - 7: Compute $\mathbf{z}_i^* = \frac{\mathbf{x}_i - \bar{\mathbf{x}}}{\sigma_{\mathbf{x}}}, i = 1, \dots, N$
 - 8: $\mathcal{L}(\mathbf{w}) = \sum_{i=1}^N \text{dist}(g(\mathbf{w}, \vec{\lambda}_i), \mathbf{z}_i^*)$ \triangleright with random restart
 - 9: $\mathbf{w}^* = OPT(\mathcal{L}, \nabla_w \mathcal{L}, \mathbf{w})$
-

Algorithm 1. SOMA SL.

Input $\mathbf{w}, \vec{\lambda}_0, \vec{\lambda}_{\min}, \vec{\lambda}_{\max}$, optimizer OPT

Output \mathbf{w}^*

```

1: for  $i = 1$  to  $L$  do
2:    $\vec{\lambda}_i \sim \pi(\vec{\lambda}_{\min}, \vec{\lambda}_{\max})$ 
3: end for
4:  $\mathcal{L}(\mathbf{w}) = 1 - \frac{1}{L} \sum_{i=1}^L F(g(\mathbf{w}, \vec{\lambda}_i), \vec{\lambda}_i)$ 
5:  $\mathbf{w}^* = OPT(\mathcal{L}, \nabla_{\mathbf{w}} \mathcal{L}, \mathbf{w})$   $\triangleright$  with random restart

```

Algorithm 2. SOMA BP.

For the control fields, we employ a Fourier ansatz:

$$u_j(t) = \sum_{k=1}^K x_{kj} \sin\left(\frac{k\pi t}{T}\right), j = 1, 2, \quad (31)$$

with K Fourier modes. For the QOC simulations, we set the central values $\delta_0 = 0$ GHz and $\alpha_0 = -0.34$ GHz [78]. The parameter vector of the QOC-problem class is

given by

$$\vec{\lambda} = (\delta, \alpha, \phi, \theta, T)^T. \quad (32)$$

For all the single-qubit gate simulations, we consider a multidimensional rectangle centered at $\vec{\lambda}_0 = (\delta_0, \alpha_0, \phi_0, \theta_0, T_0)^T$ and with upper bounds defined by $\pm\lambda_{\max}$. We consider four methods, which allow us to optimize multiple systems simultaneously: standard GRAPE; robust GRAPE, which uses the average GRAPE gradient over an ensemble of QOC problems in Eq. (15); a supervised training method, which we refer to as SOMA SL (Algorithm 1), using both linear and nonlinear models and where GRAPE solutions are first generated and then approximated via Eq. (18); and the unsupervised method, which trains a neural-network pulse directly on the fidelity of an ensemble of QOC problems using back propagation on Eq. (20). The latter is referred to as SOMA BP (Algorithm 2).

In the following section, we consider a QOC problem with $N = 500$ time steps of a Magnus-type time integrator, which approximates the unitary temporal evolution of the quantum system [72]. Our pulses are parametrized as in

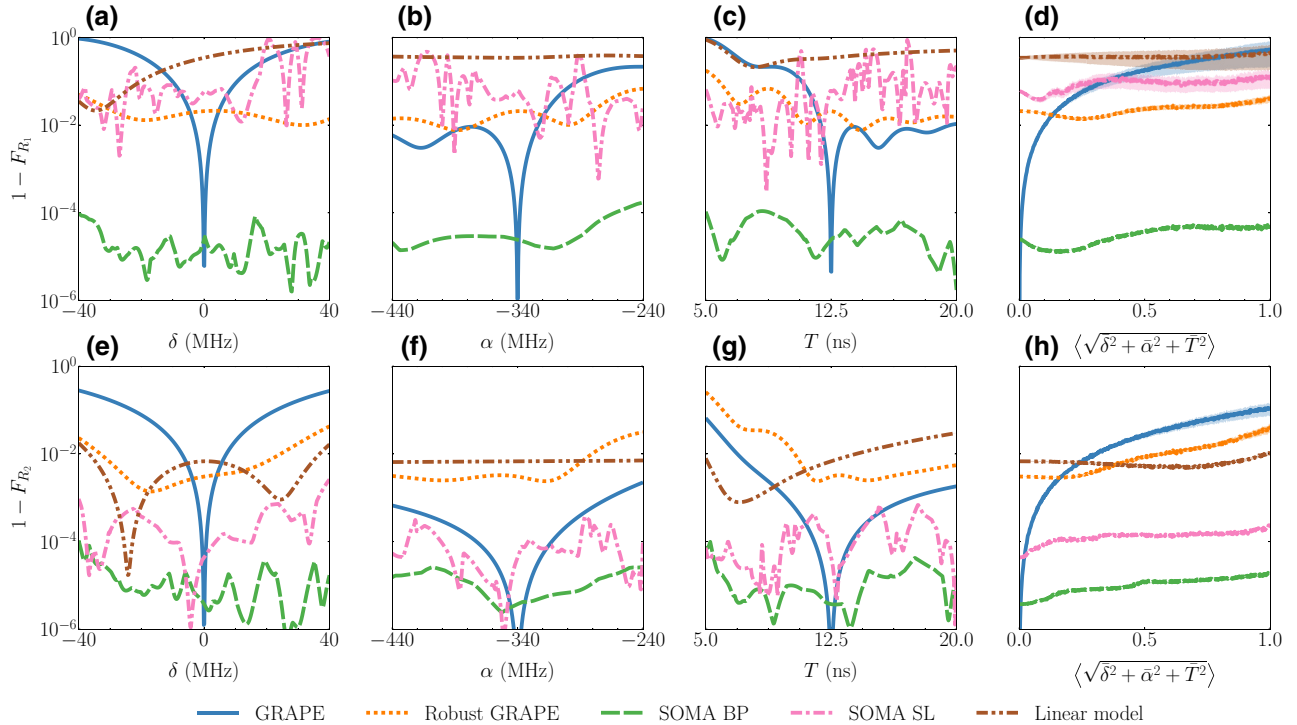


FIG. 4. The infidelity of pulses predicted by the different optimization methods for the $R_1(\theta = \pi/2)$ gate (first row) and the $R_2(\theta = \pi/2)$ gate (second row) as a function of three different quantum control problem parameters: (a),(e) the detuning δ between the qubit frequency and the driving frequency; (b),(f) the nonlinearity α of the qutrit; and (c),(g) the total gate duration T . (d),(h) The average performance (1000 test samples) of the algorithms as a function of the radial distance from the center $\vec{\lambda}_0$ of the parameter space. The terms δ, α, T indicate that the problem parameters are renormalized to the space $[0, 1]^D$ —see Appendix B. The shaded regions around each plot line show the standard deviation of the corresponding infidelities. A more detailed discussion about the standard deviation can be found in Appendix B. Both gates are optimized with four Fourier components for each one of the two control fields—see Eq. (28)—using the Hamiltonian in Eq. (29). The range of each parameter is given in Table I.

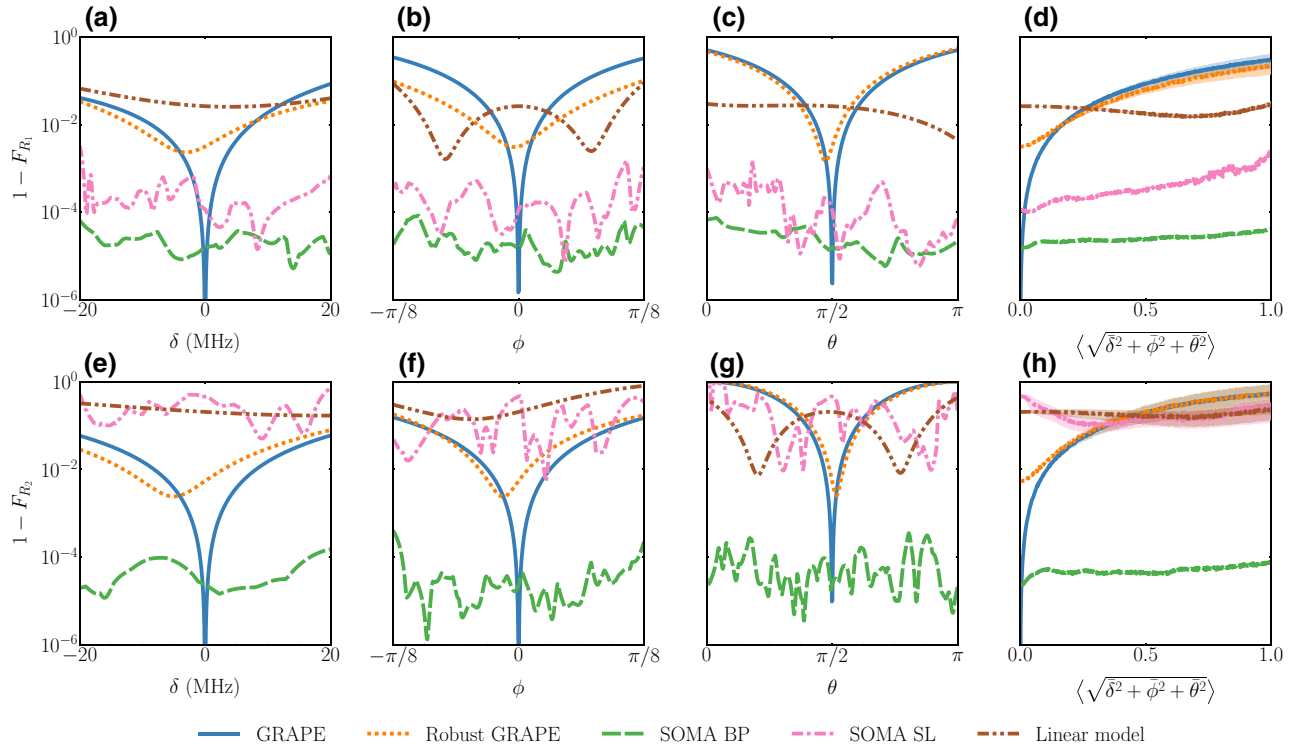


FIG. 5. The infidelity of pulses predicted by the different optimization methods for the $R_1(\theta)$ gate (first row) and the $R_2(\theta)$ gate (second row) as a function of three different quantum control problem parameters: (a),(e) the detuning δ between the qubit frequency and the driving frequency; (b),(f) the phase error ϕ between the σ_+ and σ_- terms; and (c),(g) the angle θ parametrizing the target gate. (d),(h) The average performance (1000 test samples) of the algorithms as a function of the radial distance from the center λ_0 of the parameter space. The terms $\bar{\delta}$, $\bar{\phi}$, and $\bar{\theta}$ indicate that the problem parameters are renormalized to the space $[0, 1]^D$ (see Appendix B). The shaded regions around each plot line show the standard deviation of the corresponding infidelities. A more detailed discussion about the standard deviation can be found in Appendix B. Both gates are optimized with four Fourier components for each one of the two control fields for $T = 10$ ns using the Hamiltonian in Eq. (29).

Eq. (28) by $K = 4$ Fourier components according to Eq. (31) for each one of the two control fields \hat{X} and \hat{Y} . The control fields are multiplied with a scaling factor equal to $1/A_0$ to ensure that the pulse amplitudes do not exceed the typical experimental maximal value of 1 GHz.

The results are divided into four blocks of four plots each (Figs. 4, 5), representing models trained on two different sets of problem parameters, the ranges of which are given in Table I and in Table II. Here, we limit the number of different parameters to three, although larger numbers are possible depending on the parameter range taken into account, the specifics of the physical system, the size of the neural network, and the number of systems sampled (see Appendix B). The first three plots in each row show the infidelity as a function of one varying parameter, while the other parameters are kept at their original values given by λ_0 . The fourth plot shows the average performance of the different methods as a function of the radial distance from the initial QOC problem λ_0 . This last plot ensures that the performance of the methods is stable when different parameters are changed simultaneously across the entire range of the parameter space.

In Fig. 4, the problem parameters δ (the qubit-drive detuning), α (the nonlinearity of the qutrit system), and T (the gate evolution time) and the target gates $R_1(\theta)$ defined by Eq. (23)—first row—and $R_2(\theta)$ defined by Eq. (24)—second row—are considered. The blue continuous line shows the GRAPE solution for the λ_0 parameter, the fidelity of which shows an exponential decay as a function of the distance from the center of the parameter space. The orange dashed line shows the performance of robust GRAPE and the brown dashed line the performance

TABLE I. The parameter range for the optimization of the single-qubit gates $R_1(\theta)$ and $R_2(\theta)$ as given in Fig. 4. For this simulation, we do not vary the angle parameters of the gate but, rather, the physical parameters of the qutrit, such as δ and α and the evolution time T .

	δ (MHz)	α (MHz)	ϕ	θ	$1/A_0$	T (ns)
Center	0	-340	0	$\frac{\pi}{2}$	0.01	10
Maximum	40	-240	0	$\frac{\pi}{2}$	0.01	20
Minimum	-40	-440	0	$\frac{\pi}{2}$	0.01	5

TABLE II. The parameter range for the optimization of the single-qubit gates $R_1(\theta)$ and $R_2(\theta)$ as given in Fig. 5. In this case, we not only vary the detuning but also the angles of the gates and the phase factor ϕ .

	δ (MHz)	α (MHz)	ϕ	θ	$1/A_0$	T (ns)
Center	0	-340	0	$\frac{\pi}{2}$	0.01	10
Maximum	20	-340	$\frac{\pi}{8}$	π	0.01	10
Minimum	-20	-340	$-\frac{\pi}{8}$	0	0.01	10

of SOMA SL with a linear model, both with infidelities around 10^{-2} . The pink dashed line shows the performance of SOMA SL with a neural network and the green dashed line the performance of SOMA BP. We observe that for the second gate $R_2(\theta)$, both methods are able to deliver an average infidelity below 10^{-4} , whereas for the first one $R_1(\theta)$, SOMA BP clearly outperforms SOMA SL.

In Fig. 5, the problem parameters are as follows: drive frequency detuning δ , the phase mismatch in the control fields ϕ , and the gate angle θ , together with the target gates $R_1(\theta)$ defined by Eq. (23)—first row—and $R_2(\theta)$ defined

by Eq. (24)—second row. We note the same general trends as in Fig. 4. In particular, we observe that although both SOMA methods are closer in terms of performance for $R_1(\theta)$, with a little advantage for SOMA BP, for gate $R_2(\theta)$ SOMA BP clearly outperforms SOMA SL. We can again note that SOMA BP mostly outperforms all the other methods, providing pulses that are stable over a large range of parameters.

For simulations with single-qubit gates, we use a two-layered network with 256 components per layer and eight neurons in output—the output space of the approximator has an output dimension $Q = 2K$.

B. CR gate with leakage

The cross-resonance (CR) gate is a two-qubit gate activated by microwave fields, which drive one of the qubits (target) at the frequency of the other (control). The gate is implemented in the context of quantum computing with superconducting qubits [79–82]. The gate gives rise to a ZX-type interaction [79], which can then be used to generate different two-qubit entangling unitaries. The CR gate

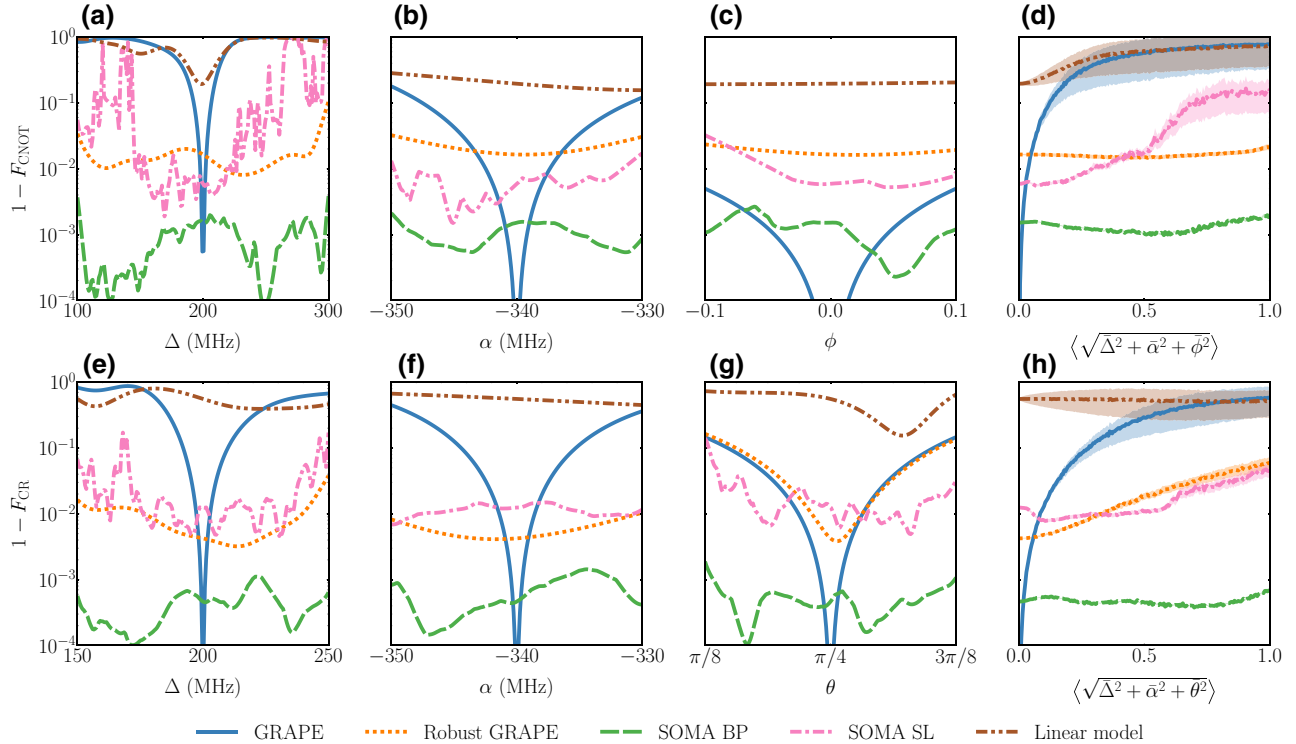


FIG. 6. The infidelity of pulses predicted by the different optimization methods for the direct controlled-NOT (CNOT) gate (first row) and the $CR(\theta)$ gate (second row) as a function of the different quantum control problem parameters: (a),(e) the frequency difference Δ between the two qubits; (b),(f) the nonlinearity α of the qutrits; (c) the phase error ϕ between control fields $\hat{\sigma}_+$ and $\hat{\sigma}_-$; and (g) the angle θ parametrizing the CR continuous family of gates given by Eq. (38). (d),(h) The average performance (1000 test samples) of the algorithms as a function of the radial distance from the center λ_0 of the parameter space. The terms Δ , $\bar{\alpha}$, $\bar{\phi}$, and $\bar{\theta}$ indicate that the problem parameters are renormalized to the space $[0, 1]^D$ —see Appendix B. The shaded regions around each plot line show the standard deviation of the corresponding infidelities. A more detailed discussion about the standard deviation can be found in Appendix B. Both gates are optimized with 30 Fourier components for each one of the four control fields and for total gate duration $T = 90$ ns.

can be embedded in a higher-dimensional system, e.g., where two qutrits are capacitively coupled together. In this case, the target gate is a two-qubit gate but the unitary generated by the Hamiltonian evolution is a two-qutrit gate; thus the fidelity is only computed with respect to the computational subspace—the CR gate is constructed using the $|0\rangle$ and $|1\rangle$ levels. The Hamiltonian of a two-transmon system reads [80]

$$\begin{aligned} H(t) &= H_d + H_c(t), \\ H_d &= \sum_{j=1}^2 \left(\omega_j \hat{b}_j^\dagger \hat{b}_j + \alpha_j \hat{b}_j \hat{b}_j^\dagger \hat{b}_j \hat{b}_j^\dagger \right) + J \left(\hat{b}_1 \hat{b}_2^\dagger + \hat{b}_1^\dagger \hat{b}_2 \right), \\ H_c(t) &= \sum_{j=1}^2 \Omega(t) e^{i\omega_j t + i\phi} \hat{b}_j^\dagger + \Omega(t)^* e^{-i\omega_j t - i\phi} \hat{b}_j, \end{aligned} \quad (33)$$

where J is the coupling strength of the transmon-transmon interaction, b_j is the lowering operator on the j th transmon, Ω is the driving field, and a Duffing-oscillator approximation is performed [76]. A further standard RWA allows us to simplify the problem, introducing at the same time the notation from Eq. (25), such that

$$H_d = \Delta \hat{n}_1 + \alpha (\hat{\Pi}_2^{(1)} + \hat{\Pi}_2^{(2)}) + J (\hat{b}_1 \hat{b}_2^\dagger + \hat{b}_1^\dagger \hat{b}_2), \quad (34)$$

$$\begin{aligned} H_c(t) &= \sum_{j=1}^2 \left(u_1^{(j)}(t) \left(e^{i\phi} \hat{b}_j^\dagger + e^{-i\phi} \hat{b}_j \right) \right. \\ &\quad \left. + u_2^{(j)}(t) \left(e^{i\phi} \hat{b}_j^\dagger - e^{-i\phi} \hat{b}_j \right) \right), \end{aligned} \quad (35)$$

where $\Delta = \omega_1 - \omega_2$. As control operators, we use the projectors $\hat{X}_j(\phi) = e^{-i\phi} \hat{b}_j + e^{i\phi} \hat{b}_j^\dagger$ and $i\hat{Y}_j(\phi) = e^{i\phi} \hat{b}_j^\dagger - e^{-i\phi} \hat{b}_j$, both on the control qutrit ($j = 1$) and the target qutrit ($j = 2$), as in Eq. (28), such that

$$H_c(t) = \sum_{j=1}^2 u_1^{(j)}(t) \hat{X}_j(\phi) + u_2^{(j)}(t) \hat{Y}_j(\phi), \quad (36)$$

where $u_j^{(1)}(t)$, $u_j^{(2)}(t)$ are the control fields on the j th qutrit (for the \hat{X}_j and the \hat{Y}_j operators, respectively). Note that all the fields operate at frequencies near-resonant to the target qubit. The Hamiltonian parameters are centered at values $\Delta_0 = 0.2$ GHz, $\alpha_0 = -0.34$ GHz, $J_0 = 0.01$ GHz. For each control field, we use a Fourier parametrization

[43,81]:

$$\forall j = 1, 2, \forall i = 1, 2 : u_i^{(j)}(t) = \sum_{k=1}^K x_k^{ij} \sin\left(\frac{k\pi t}{T}\right). \quad (37)$$

We further assume, as in the single-qubit case—see Eq. (28)—that the control fields are influenced by a phase factor ϕ , which then counts as a QOC-problem parameter in the two-qubit simulations.

For two-qubit gates with leakage, we present two different simulations, one for the controlled-NOT (CNOT) gate alone and one for a family of CR-like gates with the following parametrization:

$$\text{CR}(\theta) = \exp[i\theta(Z \otimes X)]. \quad (38)$$

For the two-qutrit gates, we use a two-layered neural network with 500 neurons per layer and 120 neurons in output. Here, the output of the approximator has a dimensionality $Q = 4K$, where $K = 30$ is the number of Fourier frequencies for each one of the four control fields—see Eq. (37). The time evolution is given by 5000 Magnus steps.

The results of the pulse-class learning are shown in Fig. 6. The evolution time is chosen as $T = 90$ ns, an improvement of about a factor of 2 over typical experimental durations. The first row of plots shows the results for the CNOT gate, where variations of the parameters Δ (the frequency difference between the control and target qubits), α (the nonlinearity of the qutrits), and ϕ (the phase term of the control fields) are considered. The second line shows the results for the CR(θ) gate as defined in Eq. (38). The maximal range for each parameter for each one of the two gates is described in Tables III and IV, respectively. We observe that, for both gates, the neural-network pulse trained with SOMA BP outperforms the other algorithms. For the CNOT gate, it can produce pulses that are robust over a very large detuning range ($|\Delta_{\max} - \Delta_{\min}| = 200$ MHz) by training on just 100 different systems sampled from a uniform distribution. For the CR gate, we choose a range of 100 MHz, while also taking into account the dependence on the angle θ . It is possible that a very

TABLE III. The parameter range for the optimization of the two-qubit CNOT gate. The range of values Δ , α , and ϕ , as well as the evolution time T , are determined by considering typical experimental settings of state-of-the-art superconducting quantum circuits [82–84]. For the value of ϕ , we assume a small phase error influencing the \hat{X} and \hat{Y} control fields both on the target and the control qubits.

	Δ (MHz)	α (MHz)	J (MHz)	ϕ	A_0	T (ns)
Center	200	−340	10	0	0.05	90
Maximum	300	−330	10	0.1	0.05	90
Minimum	100	−350	10	−0.1	0.05	90

TABLE IV. The parameter space for the optimization of the two-qubit CR gate as given in Eq. (38). Compared to the simulation of the CNOT gate, given in Table III, the range of the target-control frequency detuning Δ is reduced but the continuous parameter θ is also considered.

	Δ (MHz)	α (MHz)	J (MHz)	ϕ	θ	A_0	T (ns)
Center	200	-340	10	0	$\frac{\pi}{4}$	0.05	90
Maximum	250	-330	10	0	$\frac{\pi}{8}$	0.05	90
Minimum	150	-350	10	0	$\frac{3\pi}{8}$	0.05	90

large amount of samples or long training times could improve the performance of SOMA SL further; however, as the size of the Hilbert space increases, the optimization of large sample numbers on classical machines can become computationally too time consuming.

IV. PERFORMANCE ANALYSIS

In this section, we compare the different meta-optimization methods as the size of the QOC problem-parameter space increases. In particular, we vary the order of magnitude of the laser-frequency detuning δ , the qubit-frequency detuning Δ for the two-qubit gate, and the nonlinearity α for both the single-qubit gate and the CR gate. For the sake of this analysis, we only vary the maximal range of one single parameter at a time, while the other problem parameters have a fixed maximal range of 10^{-3} GHz. As hyperparameters for the different algorithms, we use the same values that have proved to be effective in the previous simulations. For SOMA SL, we use up to 10000 sample problems optimized with GRAPE, whereas for SOMA BP we optimize the average fidelity

with 500 system samples for the single-qubit gate and 100 systems for the two-qubit gate.

The results are shown in Fig. 7. We observe that although supervised training using the minima produced by robust GRAPE (blue line) shows lower infidelities for small variations of the detuning (up to 10 MHz), it fails when this value is increased to 100 MHz, whereas the neural network trained with back propagation of the fidelity still produces valid optima of the QOC problem. In particular, we observe a crossing between 10 and 100 MHz for both types of detuning (δ and Δ) and for the nonlinearity α in the two-qubit case, where the performance of the supervised method (SOMA SL green line) worsens dramatically, whereas SOMA BP (orange line) is able to keep the fidelity at high experimentally valid values. In general, we observe that SOMA SL significantly outperforms SOMA BP for small parameter variations, where the precision of the nonlinear regression is high enough to reproduce the pulse variations perfectly (see also Sec. 7). As a consequence, SOMA SL can be a useful tool to achieve adaptive robustness against small parameter variations, where it clearly surpasses SOMA BP, whereas the latter shines when the parameter variations and the number of Fourier components are comparatively larger. Furthermore, SOMA BP is capable of handling large variations of multiple parameters at the same time, as shown in Fig. 5. In all cases, at least one neural-network approach outperforms robust GRAPE and the linear regression model significantly. We argue that sampling large amounts of problem parameters could actually lead to a better performance for SOMA SL for *in situ* physical systems or numerical simulations where sampling proves fast and efficient. However, as we show in Fig. 8, it seems that for SOMA SL (green line), there exist systems where its

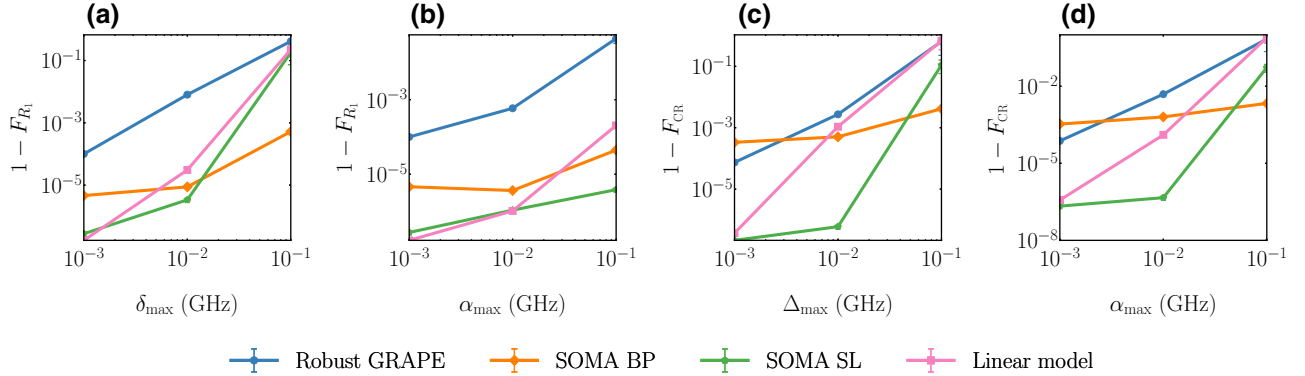


FIG. 7. A comparison of four different generalized optimization methods as a function of the parameter space size, when only one parameter range is varied, while all the others are kept constant at 10^{-3} GHz: (a),(b) the results for the single-qubit gate $R_1(\theta)$; (c),(d) the results for the CR(θ) gate. Both axes are on a logarithmic scale and represent the average infidelity over 1000 test samples. The values δ_{\max} , α_{\max} , and Δ_{\max} on the horizontal axes represent the maximum range of the corresponding problem parameter, δ , α , and Δ . The maximum range defines the parameter space over which the QOC problems are sampled. We observe how the neural network trained with the supervised algorithm outputs pulses with higher average fidelity than the other methods but it fails nonetheless when the detuning variation in the two different systems is increased up to the order of 100 MHz. In this case, the model trained with SOMA BP, however, still outputs high-fidelity pulses.




FIG. 8. An example of the average performance (1000 test samples) of robust GRAPE, SOMA SL, SOMA BP, and the linear model as a function of the number of (training) systems sampled for the single-qubit $R_2(\theta)$ gate with the same parameters as in Fig. 5. We observe here how in this case, the average infidelity in SOMA BP decreases as a function of the sample size, whereas the other methods show little to no improvement.

improvement as a function of the sample size is limited, which makes SOMA BP orange line a better option, if its implementation is possible. Likewise, robust GRAPE (blue line) and the linear model (pink line) gain limited benefit from increasing sample sizes. In particular for SOMA SL, we can often observe behaviors such as branching and outliers in the training data set, which probably contribute to a loss in the quality of the approximations. One may try to increase the precision of the model by using loss functions that are sensitive to outliers, such as the Huber loss [85], or to restrict the use of SOMA SL to systems where limited parameter drifting does not prevent the regressor from learning a high-quality representation of the solution space.

We also study whether varying the chosen samples during training (i.e., minibatching) can affect the performance of our algorithms; in particular, for robust GRAPE and SOMA BP. Computing the gradient over a batch of samples gives rise to a batched version of the algorithms. More specifically, a batched version of robust GRAPE, called bGRAPE, has been studied in Ref. [29], where impressive robustness is achieved by combining batches of problem parameters with momentum-based stochastic gradient descent. This is, of course, a different algorithm than L-BFGS-B, i.e. the optimization algorithm employed in all simulations discussed in this paper, and it does not similarly guarantee near-quadratic convergence [73]. For SOMA SL, computation of the gradient based on the MSE loss over a batch of samples leads to standard neural-network training with stochastic gradient descent. For the systems we consider, we do not observe an improvement of bGRAPE over robust GRAPE (called sGRAPE in Ref.

[29]), neither with L-BFGS-B nor with ADAM [86]. We believe that this is due to the use of Fourier components, which allow for more controllability of the quantum system [68], and the use of curvature information granted by both algorithms. Nonetheless, exploration of the effect of varying samples during training remains an interesting perspective worthy of further studies and commitment, both in the context of adaptive and robust control.

In the last part of the analysis, we discuss the scaling of the network approximations as we increase the output dimension, the number of qubits, and the input dimension. Since we aim at controlling single- and two-qubit gates, the number of expected controls only scales linearly with the number of qubits; and since the weak coupling between qubits drops off roughly exponentially with distance, we also do not expect the search complexity in state space to increase dramatically. In simulations, we expect both algorithms to behave similarly to GRAPE (or a different gradient-based control algorithm, if this is implemented) as the dimension of the Hilbert space increases. This is due to the underlying time evolution, which in both cases is given by the Trotterization or, as in our case, the Magnus expansion. The approximation of the quantum propagator affects the gradient-based optimization for single QOC problems, which generates the target data for SOMA SL but also acts as an activation function for the neural network [29] in SOMA BP. As a consequence, we do expect the state space and the equations of motion to scale exponentially with number of qubits, which is a general problem for all control and compilation tasks. Just as in the general case, we expect a combination of informed state-space parametrizations (such as tensor networks and sparse algebras) and of quantum-aided optimization (as in Sec. II) largely to address this important problem. Moreover, by considering increasing numbers of qubits, the number of pulse parameters and control fields increases consequently. However, control fields acting on different qubits usually commute, which can dramatically decrease correlations between the pulse components. Therefore, one does not need, in general, to have a single neural network outputting all pulse parameters at once but, rather, several different neural networks, one for each group of control fields, which commute between each other. For SOMA SL, we need to generate a large data set of QOC-problem solutions by means of a standard quantum control algorithm, which may be slow for many-qubit systems. However, this task can be easily parallelized, since the different optimizations are independent of each other. In this case, the main obstacle is not represented by the nonlinear regression over the data but, rather, by the quality of the data generated by the quantum control algorithm for each solution in the parameter space. Since the problem is high dimensional, the solution space will probably exhibit structures such as branching and outliers that are difficult to include in the nonlinear regression. A possible option here

TABLE V. A comparison of SOMA SL and SOMA BP in terms of their performance features. Here, n_{fev} and t_{GRAPE} are the maximal number of function evaluations and the computation time required by L-BFGS-B with GRAPE to solve a single quantum control problem. t_{REG} is the time required for the neural-network regression. t_{BPROP} is the time required by the neural-network back propagation for one sample.

Algorithm parameter	SOMA SL	SOMA BP
QOC cost-function evaluations	$n_{\text{fev}}N$	$n_{\text{fev}}LN_{\text{guesses}}$
Computation time	$t_{\text{GRAPE}}N + t_{\text{REG}}$	$t_{\text{BPROP}}L$
Number of samples	N	L

is to employ algorithms for data reduction and clustering, in order to obtain a high-quality representation of the solution space. As for SOMA BP, one may distinguish between simulation and experimental implementation. In the former case, the evaluation of the infidelity and its gradient as a function of the time evolution represents the main bottleneck (see Table V in Appendix B), together with vanishing gradients, which are also a well-known problem for other NISQ use cases and represent one of the main obstacles to any experimental application. Vanishing gradients could also be tackled with alternatives to back propagation (see, e.g., Ref. [87]). Finally, for large output spaces and very deep networks, GPU training and stochastic gradient-descent algorithms may provide useful speedup, as it is the standard in deep learning.

V. CONCLUSIONS

In this work, we show how to engineer solutions of problems in quantum optimal control that depend on problem parameters located outside the optimization routine. This includes physical system parameters, other external parameters such as the pulse time or bandwidth, and gate parameters such as rotation angles.

We therefore propose two methods to learn large classes of quantum gate-synthesis problems, SOMA SL and SOMA BP. We show through experimentally relevant examples that these methods prove able to learn adaptive solutions to generalized QOC problems. The output gates have fidelities that remain very high over the entire continuous parametrization of the gate sets, for typically large ranges as would be encountered experimentally. These continuous gate sets provide the opportunity to be used as computational primitives in compilation tasks, in NISQ variational algorithms, and for entire arrays of qubits rather than individually optimized ones.

ACKNOWLEDGMENTS

We thank Michael Schilling, Boxi Li, Martino Calzavara, and Markus Schmitt for useful discussions. We acknowledge support from the German Ministry for Education and Research, under the QSolid project with

Grant No. 13N16149, the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy—Cluster of Excellence Matter and Light for Quantum Computing (ML4Q) EXC 2004/1—390534769, the Jülich Supercomputing Center, and from the project AIDAS: AI, Data Analytics and Scalable Simulations, a joint virtual laboratory gathering the Forschungszentrum Jülich (FZJ) and the French Alternative Energies and Atomic Energy Commission (CEA). The simulations were realized using the JAX [69] and FLAX [88] libraries.

APPENDIX A: GRADIENTS OF THE FIDELITY

GRAPE [41,89] is the standard open-loop gradient-based method for QOC solutions. It can be implemented together with a Magnus-based propagator—see Ref. [72]. The gradient can then be used in combination with a gradient-based optimization algorithm, e.g., the L-BFGS-B algorithm [73,90]. Variations of GRAPE exist that exploit the advantages of parametrizing the pulse according to a specific set of basis functions (e.g., Fourier basis) [68] or compute the higher-order derivatives of the fidelity with respect to the pulse [51,89]. In this paper, GRAPE is always used in combination with L-BFGS-B, which uses fast Hessian estimation [91].

For a given set of parametrized functions $s_k : t \mapsto s_k(t)$, $k = 1, \dots, K$ describing the time dynamics of the control fields $u_j(t)$ with control parameters x_{kj} —see Eqs. (28) and (37)—and that can be time sliced in values $s_{ki} = s_k(t_i)$, $i = 0, \dots, N_{\text{evo}}$, the gradient of the cost function with respect to the control parameters can be computed using the chain rule [68]:

$$\frac{\partial F}{\partial x_{kj}} = \frac{\partial u_{ij}}{\partial x_{kj}} \frac{\partial F}{\partial u_{ij}}, \quad (\text{A1})$$

where the k index runs over the number of basis-function components, the i index over the time slice, and the j index over the different control operators. In a similar way, this can be applied to a neural-network parametrization $g : \mathbb{R}^D \mapsto \mathbb{R}^Q, \vec{\lambda} \mapsto g(\vec{\lambda})$ of the GRAPE pulse, mapping a given number of metaparameters to the pulse space, and the original formula can be rewritten to output the gradients of the fidelity with respect to the neural-network parameters (w_{ml} , b_l):

$$\frac{\partial F}{\partial w_{ml}} = \frac{\partial x_{kj}}{\partial w_{ml}} \frac{\partial u_{ij}}{\partial x_{kj}} \frac{\partial F}{\partial u_{ij}}, \quad (\text{A2})$$

$$\frac{\partial F}{\partial b_l} = \frac{\partial x_{kj}}{\partial b_l} \frac{\partial u_{ij}}{\partial x_{kj}} \frac{\partial F}{\partial u_{ij}}, \quad (\text{A3})$$

with the same indexing as in Eq. (A1). In this case, the neural-network output values, which give the coefficients

of the time-dependent basis functions, i.e., the terms $x_{ik} = g(\vec{\lambda})_{ik}$, depend on the QOC-problem parameters $\vec{\lambda}$.

APPENDIX B: IMPLEMENTATION DETAILS

For every physical system considered—the single-qubit gate with leakage and the two-qubit gate with leakage, each one with different target gates and system parameters—we first define an interval for each parameter. Values are sampled from a uniform distribution over the hypervolume defined by the intervals of parameters. The system together with the interval of parameters defines a family of QOC problems, which we analyze with one of three methods: robust control with GRAPE, which simply seeks for the best pulse for a set of different systems; SOMA SL, which first solves a sample of systems and then performs a regression over the sampled values—for the sake of completeness, we consider here both a linear and a nonlinear approach; and SOMA BP, where the network is trained directly on the average infidelity of an ensemble of systems with back propagation.

In the case of robust GRAPE, we sample L systems and run the optimization with random restart, i.e., we run the optimization $N_{\text{guesses}} = 5$ times with different conditions and then choose the pulse with the smallest average infidelity.

For SOMA SL, we sample up to $N = 10\,000$ points within the parameter space. For each one of these, we optimize the corresponding QOC problem with GRAPE—this can be performed with any proper optimal quantum control method—and then train the model to map the corresponding parameter to the optimal pulses.

As for SOMA BP, we sample L systems and run the neural-network training with random restart. The number of total samples required by the regression is usually larger than the one needed by the other two methods. The corresponding single-system optimization is nonetheless much faster and can be run in parallel. Therefore, we set $N = 10\,000$ for each simulation, whereas for the sake of comparison and due to the similarity of the other two methods, we always use the same number of samples (either $L = 500$ for the single-qubit case or $L = 100$ for the two-qubit case). Both SOMA methods employ two-layered neural networks with 256 neurons per layer for the single-qubit system and 500 neurons per layer for the two-qubit system. The linear model performs multilinear regression on the same data used by SOMA SL.

For the random restart of SOMA methods, we run the optimization $N_{\text{guesses}} = 5$ times with different initial conditions and then test the quality of the predictions on a test set. Afterward, we choose the model with the lowest average test infidelity. For both algorithms, we use the L-BFGS-B algorithm. Since the training of SOMA SL, which uses a MSE loss, is much faster than that for SOMA BP, we do not limit its maximal number of iterations. For SOMA

BP, however, this is limited to 6000 for the single-qubit gates and to 7000 for the two-qubit gates. Here, the input parameters for both SOMA SL and SOMA BP should be rescaled to lie, e.g., within the range $[0,1]$. For SOMA BP, this can be avoided in certain cases, as long as the size of the pulses remains large enough. For details about the performance of the algorithms, see Table V.

For each system, we evaluate the performance of the neural-network pulses on the entire parameter space. In order to visualize the performance of the algorithms considered in a one-dimensional plot, we consider the normalized parameter space, where all parameter axes are rescaled to the interval $[0,1]$:

$$f : \mathbb{R}^D \mapsto [0, 1]^D, \vec{\lambda} \longrightarrow \frac{\vec{\lambda} - \vec{\lambda}_{\min}}{|\vec{\lambda}_{\max} - \vec{\lambda}_{\min}|}. \quad (\text{B1})$$

By then considering D -dimensional spheres of radius $r \in [0, 1]^D$, we sample $N_s = 1000$ systems on the surfaces of such spheres and compute the average infidelity over these systems. By doing so, we can evaluate the performance over the parameter space, thereby ensuring that our methods are effective for every combination of problem-parameter values. The result is then averaged over N_s samples, i.e., we plot the mean and its standard. The latter is pictured as a shady region.

We also want to consider how the standard deviation of the average infidelity behaves when the different algorithms are tested against a batch of quantum systems sampled according to the parameter space. In particular, the standard deviation of the infidelity for N_s test systems is bounded by the average performance of the algorithms:

$$\sigma(\mathbf{w})^2 = \frac{1}{L} \sum_{i=1}^{N_s} F(\mathbf{w}, \vec{\lambda}_i)^2 - F_{\text{test}}^2 \leq 1 - F_{\text{test}}^2, \quad (\text{B2})$$

where F_{test} is the average fidelity of the trained pulses on the N_s test systems. For SOMA BP, assuming that overfitting is negligible, we have $\sigma(\mathbf{w})^2 \leq 1 - (1 - \mathcal{L}(\mathbf{w}))^2$, which means that the fidelity of the algorithms considered is guaranteed not to drop significantly below the average performance shown in Figs. 4, 5 and 6.


APPENDIX C: FIDELITY OF DISCRETE GATES AND THEIR PERFORMANCE

In this appendix, we briefly discuss how analog gates are expected to outperform equivalent-circuit decomposition with sequences of traditional discrete gates. In particular, we first consider the standard universal set of quantum gates $S = \{\text{CNOT}, T, S, H\}$ used in Fig. 1. We search for decompositions to determine the correct compilation sequence for the aforementioned gates. In particular, we consider the CR gate defined in Eq. (38) for an angle of

TABLE VI. The number of gates: the results of searching for CR-gate decompositions for four different angles— $\frac{\pi}{\sqrt{2}}$, $\frac{\pi}{\sqrt{3}}$, $\frac{\pi}{\sqrt{5}}$, and $\frac{\pi}{\sqrt{7}}$ —obtained with the help of exhaustive search (up to ten circuit layers) and stochastic descent (up to 20 circuit layers). The universal gate set is S . We observe here that while for $\frac{\pi}{\sqrt{7}}$ we can find a suitable decomposition with high fidelity, this is not the case for the remaining angle values. Moreover, two of the best decompositions contain two CNOT gates, which again leads us to a much longer gate-evolution time.


Angle θ	N_{CNOT}	N_T	N_S	N_H	Highest fidelity
$\frac{\pi}{\sqrt{2}}$	1	2	4	1	0.9819
$\frac{\pi}{\sqrt{3}}$	2	3	1	2	0.9778
$\frac{\pi}{\sqrt{5}}$	0	2	3	2	0.9728
$\frac{\pi}{\sqrt{7}}$	2	3	0	2	0.9999

$\theta = \pi/4$. The decomposition is given by the following circuit:



This decomposition consists of one entangling gate (the CNOT) and four local unitary operations and the representation is exact. Note that CNOT has to be implemented efficiently on the chosen quantum computing platform, which probably requires quantum control anyway. An entangling gate in the context of superconducting circuits usually requires $T \sim 150$ ns [92]. This is a little bit slower than the CR gate presented here but shows the equivalency of these gate sets. Nonetheless, SOMA also allows learning of both the CR and the CNOT as part of the same continuous gate set, thereby saving the additional cost of single-qubit gates (which typically take at least $T \sim 15$ ns each).

Where the advantage really shows up is for a gate angle that does not belong to the same entanglement class as CNOT. Here, it is clear how continuously parametrized gates and consequently SOMA can be beneficial to quantum compilation. A CR gate with an angle $\theta = \pi/8$ also requires five gates to be implemented, as shown by the following circuit:



but its circuit contains two CNOT gates. Since the CR ($\pi/8$) can take roughly half the time of a CR($\pi/4$), this implies that the CNOT-based circuit can be as much as 7 times slower than the continuously parametrized gate.

For other angles, it becomes increasingly difficult to find good circuits, with only partial approximations being

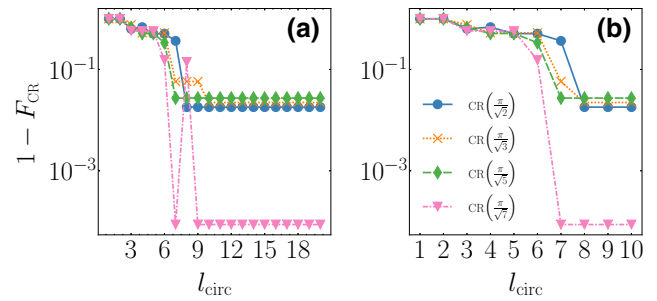


FIG. 9. Infidelities reached by applying (a) stochastic descent search and (b) brute-force (exhaustive) search on the CR gates defined by the angles in Table VI in order to decompose them into discrete circuits. As we can observe, both algorithms provide us with the same results. However, they cannot find a circuit representation with $F > 0.98$ for the first three gates, thus indicating that circuits decomposing these gates with a higher fidelity are longer, more error prone, and harder to discover.

possible at reasonably short depth. In order to test the quality of the approximation, we search for optimal discrete circuits representing a given circuit according to a given parametrization. We use exhaustive search [93] of all possible circuits (up to circuit depth 10) and stochastic descent, a special type of structured random search—see the reinforcement-learning (RL) approach in Ref. [26]—(up to circuit depth 20) to search for optimal decompositions of discrete gates and try to reproduce the chosen circuit with an increasing number of unitaries. The results are given in Table VI and Fig. 9 for the CR gate with angles $\pi/\sqrt{2}$, $\pi/\sqrt{3}$, $\pi/\sqrt{5}$, and $\pi/\sqrt{7}$. We see that although the fidelity of the discrete gates increases with the size of the quantum circuit, in three cases it cannot reach the value of $F = 0.99$ for circuits of depth smaller than 20. In the case of $\theta = \pi/\sqrt{7}$, a valid decomposition with fidelity $F = 0.9999$ is found. The search is performed exhaustively for $l_{\text{circ}} < 10$ and then using stochastic descent for $l_{\text{circ}} > 10$. Moreover, the decomposition of CR($\pi/\sqrt{7}$) contains two CNOT gates, which again implies a slowdown of the gate execution time.

Other examples for gates with superior analog performance can be seen in, e.g., Ref. [8], with superior performance especially expected for variational circuits.

- [1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, 2010).
- [2] C. M. Dawson and M. A. Nielsen, The Solovay-Kitaev algorithm, *Quantum Info. Comput.* **6**, 81 (2006).

- [3] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter, Elementary gates for quantum computation, *Phys. Rev. A* **52**, 3457 (1995).
- [4] A. Peres, Reversible logic and quantum computers, *Phys. Rev. A* **32**, 3266 (1985).
- [5] D. Gottesman, Theory of fault-tolerant quantum computation, *Phys. Rev. A* **57**, 127 (1998).
- [6] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [7] T. Proctor, M. Revelle, E. Nielsen, K. Rudinger, D. Lobser, P. Maunz, R. Blume-Kohout, and K. Young, Detecting and tracking drift in quantum information processors, *Nat. Commun.* **11**, 5396 (2020).
- [8] B. Foxen *et al.*, Demonstrating a Continuous Set of Two-Qubit Gates for Near-Term Quantum Algorithms, *Phys. Rev. Lett.* **125**, 120504 (2020).
- [9] J. Yao, L. Lin, and M. Bukov, Reinforcement Learning for Many-Body Ground-State Preparation Inspired by Counterdiabatic Driving, *Phys. Rev. X* **11**, 031070 (2021).
- [10] J. Preskill, Quantum computing in the NISQ era and beyond, *Quantum* **2**, 79 (2018).
- [11] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [12] R. Barends *et al.*, Digitized adiabatic quantum computing with a superconducting circuit, *Nature* **534**, 222 (2016).
- [13] B. Apolloni, C. Carvalho, and D. de Falco, Quantum stochastic optimization, *Stoch. Process. Their Appl.* **33**, 233 (1989).
- [14] T. Kadowaki and H. Nishimori, Quantum annealing in the transverse Ising model, *Phys. Rev. E* **58**, 5355 (1998).
- [15] 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).
- [16] E. Farhi, J. Goldstone, and S. Gutmann, A quantum approximate optimization algorithm, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028) (2014).
- [17] T. Hospedales, A. Antoniou, P. Micaelli, and A. Storkey, Meta-learning in neural networks: A survey, [arXiv:2004.05439](https://arxiv.org/abs/2004.05439) (2020).
- [18] D. Devikanniga, K. Vetrivel, and N. Badrinath, Review of meta-heuristic optimization based artificial neural networks and its applications, *J. Phys.: Conf. Ser.* **1362**, 012074 (2019).
- [19] S. Arimoto, S. Kawamura, and F. Miyazaki, Bettering operation of robots by learning, *J. Robot. Syst.* **1**, 123 (1984).
- [20] G. Tang, W. Sun, and K. Hauser, in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, Madrid, 2018), p. 3620.
- [21] A. D'Souza, S. Vijayakumar, and S. Schaal, in *Proceedings 2001 IEEE/RSJ International Conference on Intelligent Robots and Systems. Expanding the Societal Role of Robotics in the the Next Millennium (Cat. No. 01CH37180)* (IEEE, Maui, HI, USA, 2001), Vol. 1, p. 298.
- [22] N. Jetchev and M. Toussaint, in *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09* (Association for Computing Machinery, New York, 2009), p. 449.
- [23] S. Chen and J. T. Wen, Neural-learning trajectory tracking control of flexible-joint robot manipulators with unknown dynamics, [arXiv:1908.03269](https://arxiv.org/abs/1908.03269) (2019).
- [24] T. Ozaki, T. Suzuki, T. Furuhashi, S. Okuma, and Y. Uchikawa, Trajectory control of robotic manipulators using neural networks, *IEEE Trans. Ind. Electron.* **38**, 195 (1991).
- [25] L. Böttcher, N. Antulov-Fantulin, and T. Asikis, AI Pontryagin or how artificial neural networks learn to control dynamical systems, *Nat. Commun.* **13**, 333 (2022).
- [26] M. Bukov, A. G. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, Reinforcement Learning in Different Phases of Quantum Control, *Phys. Rev. X* **8**, 031086 (2018).
- [27] X.-C. Yang, M.-H. Yung, and X. Wang, Neural-network-designed pulse sequences for robust control of singlet-triplet qubits, *Phys. Rev. A* **97**, 042324 (2018).
- [28] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven, Universal quantum control through deep reinforcement learning, *npj Quantum Inf.* **5**, 33 (2019).
- [29] R.-B. Wu, H. Ding, D. Dong, and X. Wang, Learning robust and high-precision quantum controls, *Phys. Rev. A* **99**, 042327 (2019).
- [30] S. Lohani, B. T. Kirby, M. Brodsky, O. Danaci, and R. T. Glasser, Machine learning assisted quantum state estimation, *Mach. Learn.: Sci. Technol.* **1**, 035007 (2020).
- [31] J. R. Kuklinski, U. Gaubatz, F. T. Hioe, and K. Bergmann, Adiabatic population transfer in a three-level system driven by delayed laser pulses, *Phys. Rev. A* **40**, 6741 (1989).
- [32] N. V. Vitanov, A. A. Rangelov, B. W. Shore, and K. Bergmann, Stimulated Raman adiabatic passage in physics, chemistry, and beyond, *Rev. Mod. Phys.* **89**, 015006 (2017).
- [33] A. Sørensen and K. Mølmer, Quantum Computation with Ions in Thermal Motion, *Phys. Rev. Lett.* **82**, 1971 (1999).
- [34] A. Sørensen and K. Mølmer, Entanglement and quantum computation with ions in thermal motion, *Phys. Rev. A* **62**, 022311 (2000).
- [35] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [36] L. S. Theis, F. Motzoi, S. Machnes, and F. K. Wilhelm, Counteracting systems of diabaticities using DRAG controls: The status after 10 years, *EPL (Europhysics Letters)* **123**, 60001 (2018).
- [37] A. Bocharov, M. Roetteler, and K. M. Svore, Efficient Synthesis of Universal Repeat-until-Success Quantum Circuits, *Phys. Rev. Lett.* **114**, 080502 (2015).
- [38] F. Motzoi, M. P. Kaicher, and F. K. Wilhelm, Linear and Logarithmic Time Compositions of Quantum Many-Body Operators, *Phys. Rev. Lett.* **119**, 160503 (2017).
- [39] W. S. Warren, H. Rabitz, and M. Dahleh, Coherent control of quantum dynamics: The dream is alive, *Science* **259**, 1581 (1993).
- [40] R. Freeman, *Spin Choreography* (Oxford University Press, Oxford, 1998).
- [41] 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. Magn. Reson.* **172**, 296 (2005).

- [42] M. H. Goerz, D. Basilewitsch, F. Gago-Encinas, M. G. Krauss, K. P. Horn, D. M. Reich, and C. P. Koch, KROTOV: A Python implementation of Krotov's method for quantum optimal control, *SciPost Phys.* **7**, 80 (2019).
- [43] T. Caneva, T. Calarco, and S. Montangero, Chopped random-basis quantum optimization, *Phys. Rev. A* **84**, 022326 (2011).
- [44] 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).
- [45] J. J. Grefenstette, Optimization of control parameters for genetic algorithms, *IEEE Trans. Syst. Man Cybern.* **16**, 122 (1986).
- [46] S. Schaal and C. G. Atkeson, Learning control in robotics, *IEEE Robot. Autom. Mag.* **17**, 20 (2010).
- [47] C. Atkeson and J. McIntyre, in *Proceedings. 1986 IEEE International Conference on Robotics and Automation* (IEEE, San Francisco, CA, USA, 1986), Vol. 3, p. 1737.
- [48] B. Li, T. Calarco, and F. Motzoi, Non-perturbative analytical diagonalization of Hamiltonians with application to ZZ-coupling suppression and enhancement, [arXiv:2112.00039](https://arxiv.org/abs/2112.00039) (2021).
- [49] E. A. Martinez, T. Monz, D. Nigg, P. Schindler, and R. Blatt, Compiling quantum algorithms for architectures with multi-qubit gates, *New J. Phys.* **18**, 063029 (2016).
- [50] F. Frank, T. Unden, J. Zoller, R. S. Said, T. Calarco, S. Montangero, B. Naydenov, and F. Jelezko, Autonomous calibration of single spin qubit operations, *npj Quantum Inf.* **3**, 48 (2017).
- [51] M. Dalgaard, F. Motzoi, J. H. M. Jensen, and J. Sherson, Hessian-based optimization of constrained quantum control, *Phys. Rev. A* **102**, 042612 (2020).
- [52] C. R. de Oliveira, *Intermediate Spectral Theory and Quantum Dynamics (Progress in Mathematical Physics, 54)* (Birkhäuser, Basel, 2008), p. 423.
- [53] T. W. Borneman, M. D. Hürlimann, and D. G. Cory, Application of optimal control to CPMG refocusing pulse design, *J. Magn. Reson.* **207**, 220 (2010).
- [54] M. Möttönen, R. de Sousa, J. Zhang, and K. B. Whaley, High-fidelity one-qubit operations under random telegraph noise, *Phys. Rev. A* **73**, 022332 (2006).
- [55] B. Khani, S. T. Merkel, F. Motzoi, J. M. Gambetta, and F. K. Wilhelm, High-fidelity quantum gates in the presence of dispersion, *Phys. Rev. A* **85**, 022306 (2012).
- [56] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger, Optimized dynamical decoupling in a model quantum memory, *Nature* **458**, 996 (2009).
- [57] F. Motzoi, E. Halperin, X. Wang, K. B. Whaley, and S. Schirmer, Backaction-driven, robust, steady-state long-distance qubit entanglement over lossy channels, *Phys. Rev. A* **94**, 032313 (2016).
- [58] S. G. Schirmer, F. C. Langbein, C. A. Weidner, and E. Jonckheere, Robust control performance for open quantum systems, [arXiv:2008.13691](https://arxiv.org/abs/2008.13691) (2020).
- [59] M. M. Müller, S. Gherardini, and F. Caruso, Noise-robust quantum sensing via optimal multi-probe spectroscopy, *Sci. Rep.* **8**, 14278 (2018).
- [60] S. Wimperis, Broadband, narrowband, and passband composite pulses for use in advanced NMR experiments, *J. Magn. Reson., Ser. A* **109**, 221 (1994).
- [61] H. K. Cummins and J. A. Jones, Use of composite rotations to correct systematic errors in NMR quantum computation, *New J. Phys.* **2**, 006 (2000).
- [62] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (The MIT Press, Cambridge, Massachusetts, 2016), <http://www.deeplearningbook.org>.
- [63] R. S. Sutton, D. McAllester, S. Singh, and Y. Mansour, in *Advances in Neural Information Processing Systems*, edited by S. Solla, T. Leen, and K. Müller (The MIT Press, Cambridge, Massachusetts, 1999), Vol. 12.
- [64] M. A. Nielsen, A simple formula for the average gate fidelity of a quantum dynamical operation, *Phys. Lett. A* **303**, 249 (2002).
- [65] J. A. Nelder and R. Mead, A simplex method for function minimization, *Comput. J.* **7**, 308 (1965).
- [66] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhuber, Natural evolution strategies, *J. Mach. Learn. Res.* **15**, 949 (2014).
- [67] S. Mohamed, M. Rosca, M. Figurnov, and A. Mnih, Monte Carlo gradient estimation in machine learning, *J. Mach. Learn. Res.* **21**, 1 (2020).
- [68] F. Motzoi, J. M. Gambetta, S. T. Merkel, and F. K. Wilhelm, Optimal control methods for rapidly time-varying Hamiltonians, *Phys. Rev. A* **84**, 022307 (2011).
- [69] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, JAX: Composable transformations of PYTHON+NumPy programs, [http://github.com/google/jax](https://github.com/google/jax) (2018).
- [70] L. B. Rall, *Automatic Differentiation*, Lecture Notes in Computer Science (Springer, Berlin, 2006).
- [71] S. Blanes, F. Casas, J. Oteo, and J. Ros, The Magnus expansion and some of its applications, *Phys. Rep.* **470**, 151 (2009).
- [72] M. Dalgaard and F. Motzoi, Fast, high precision dynamics in quantum optimal control theory, *J. Phys. B: At. Mol. Opt. Phys.* **55**, 085501 (2022).
- [73] D. C. Liu and J. Nocedal, On the limited memory BFGS method for large scale optimization, *Math. Program.* **45**, 503 (1989).
- [74] IBM, IBM Quantum Composer guide, [http://quantum-computing.ibm.com/composer/docs/iqx/guide/advanced-single-qubit-gates/](https://quantum-computing.ibm.com/composer/docs/iqx/guide/advanced-single-qubit-gates/) (2022) [Online; accessed January 5, 2022].
- [75] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, *Annu. Rev. Condens. Matter Phys.* **11**, 369 (2020).
- [76] B. Khani, J. M. Gambetta, F. Motzoi, and F. K. Wilhelm, Optimal generation of Fock states in a weakly nonlinear oscillator, *Phys. Scr.* **T137**, 014021 (2009).
- [77] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom-Photon Interactions* (John Wiley & Sons, Nashville, Tennessee, 1998).
- [78] L. S. Theis, F. Motzoi, and F. K. Wilhelm, Simultaneous gates in frequency-crowded multilevel systems using fast,

- robust, analytic control shapes, *Phys. Rev. A* **93**, 012324 (2016).
- [79] G. S. Paraoanu, Microwave-induced coupling of superconducting qubits, *Phys. Rev. B* **74**, 140504 (2006).
- [80] C. Rigetti and M. Devoret, Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies, *Phys. Rev. B* **81**, 134507 (2010).
- [81] S. Kirchhoff, T. Keßler, P. J. Liebermann, E. Assémat, S. Machnes, F. Motzoi, and F. K. Wilhelm, Optimized cross-resonance gate for coupled transmon systems, *Phys. Rev. A* **97**, 042348 (2018).
- [82] M. Malekakhlagh and E. Magesan, Mitigating off-resonant error in the cross-resonance gate, *Phys. Rev. A* **105**, 012602 (2022).
- [83] M. Malekakhlagh, E. Magesan, and D. C. McKay, First-principles analysis of cross-resonance gate operation, *Phys. Rev. A* **102**, 042605 (2020).
- [84] K. Heya and N. Kanazawa, Cross-Cross Resonance Gate, *PRX Quantum* **2**, 040336 (2021).
- [85] P. J. Huber, Robust estimation of a location parameter, *Ann. Math. Stat.* **35**, 73 (1964).
- [86] D. P. Kingma and J. Ba, ADAM: A method for stochastic optimization, CoRR **abs/1412.6980** (2015).
- [87] D.-H. Lee, S. Zhang, A. Fischer, and Y. Bengio, Difference target propagation, [arXiv:1412.7525](https://arxiv.org/abs/1412.7525) (2014).
- [88] J. Heek, A. Levskaya, A. Oliver, M. Ritter, B. Rondepierre, A. Steiner, and M. van Zee, FLAX: A neural network library and ecosystem for JAX, <https://github.com/google/flax> (2020).
- [89] P. de Fouquieres, S. Schirmer, S. Glaser, and I. Kuprov, Second order gradient ascent pulse engineering, *J. Magn. Reson.* **212**, 412 (2011).
- [90] R. H. Byrd, J. Nocedal, and R. B. Schnabel, Representations of quasi-Newton matrices and their use in limited memory methods, *Math. Program.* **63**, 129 (1994).
- [91] P. Virtanen *et al.*, Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).
- [92] A. Kandala, K. X. Wei, S. Srinivasan, E. Magesan, S. Carnevale, G. A. Keefe, D. Klaus, O. Dial, and D. C. McKay, Demonstration of a High-Fidelity CNOT Gate For Fixed-Frequency Transmons With Engineered ZZ Suppression, *Phys. Rev. Lett.* **127**, 130501 (2021).
- [93] J. Nievergelt, in *Proceedings of the 27th Conference on Current Trends in Theory and Practice of Informatics, SOFSEM '00* (Springer-Verlag, Berlin, 2000), p. 18.