

This is the accepted manuscript made available via CHORUS. The article has been published as:

Optimal Control of Families of Quantum Gates

Frédéric Sauvage and Florian Mintert

Phys. Rev. Lett. **129**, 050507 — Published 29 July 2022

DOI: [10.1103/PhysRevLett.129.050507](https://doi.org/10.1103/PhysRevLett.129.050507)

Optimal control of families of quantum gates

Frédéric Sauvage^{1,2} and Florian Mintert¹

¹*Physics Department, Blackett Laboratory, Imperial College London, Prince Consort Road, SW7 2BW, United Kingdom*

²*Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

(Dated: June 30, 2022)

Quantum Optimal Control (QOC) enables the realization of accurate operations, such as quantum gates, and support the development of quantum technologies. To date, many QOC frameworks have been developed but those remain only naturally suited to optimize a single targeted operation at a time. We extend this concept to optimal control with a continuous family of targets, and demonstrate that an optimization based on neural networks can find families of time-dependent Hamiltonians realizing desired classes of quantum gates in minimal time.

After concerted efforts in the development of synthetic quantum systems we have access to a variety of systems with sufficiently long coherence time to perform a series of coherent operations. In the community's effort to turn such systems into technological applications, Quantum Optimal Control (QOC) [1, 2] helps to increase the precision and rate of desired operations. Common problems successfully addressed by means of QOC include the realization of quantum gates or entangled states in few-body or many-body systems [3–11] and the refinement of metrology protocols [12, 13].

Current tasks of optimal control are mostly focused on the realization of a single target operation, such as the preparation of one specific state or the implementation of one specific gate. Yet, as quantum technologies mature, it becomes important to enlarge the range of operations which can be accurately implemented on a device. For instance, in the context of Noisy-Intermediate Scale Quantum (NISQ) devices [14], augmenting the set of available elementary gates allows for more compact compilation of quantum circuits, *i.e.*, their decomposition into these elementary gates. Already the inclusion of continuous families of two-qubit gates to a typical gate-set, composed of single-qubit rotations and a two-qubit entangling gate, can lead to a significant reduction in gate count [15–17]. This, in turn, opens the possibility to run more expressive computations before the onset of decoherence, a key limitation in current technology. That is, the ability to implement a broader range of optimized operations has the potential to substantially increase the utility of current quantum hardware.

Despite the many flavours of QOC frameworks that have been proposed (e.g., [18–28]), it remains the case that current methodologies are only naturally suited to consider a single control task at a time. We thus aim at lifting the original scope of QOC from the control of a single target operation to the control of continuous families of targets. This is achieved with a neural-network (NN) modelling the dependency between Hamiltonians to be engineered and control tasks to be solved. Efficient training of the framework by means of gradient-descent is facilitated by recent advances in the field of automatic-differentiation [29]. Such framework, dubbed family-control, is sketched in Fig. 1 and is now explained in detail.

Typically, the central task in QOC is the identification of the time-dependent Hamiltonian $H(t)$ that induces a propagator

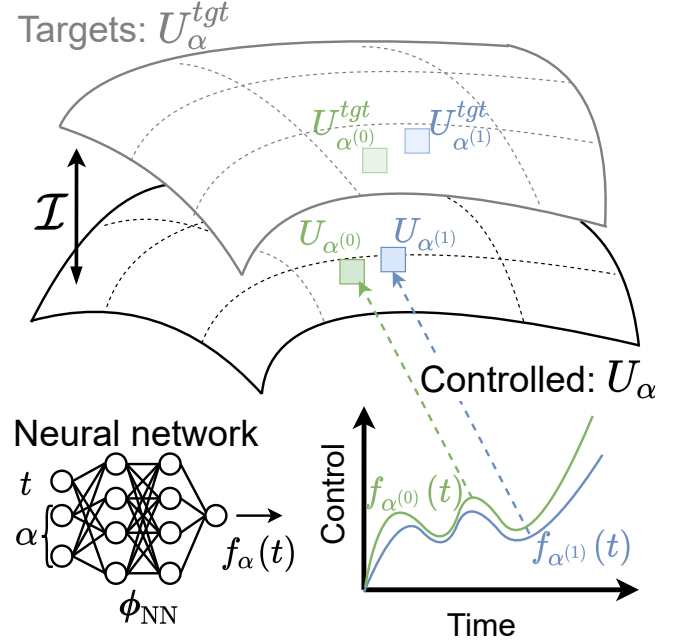


FIG. 1. Optimal control of a continuous family of target gates U_α^{tgt} indexed by the target parameter α which can be either a scalar or a vector. The time-dependent controls $f_\alpha(t)$ which now also depend on α are modeled by a neural network (NN). This NN effectively parameterizes a continuous family of controlled gates U_α where each point corresponds to the propagator obtained by evolving the system in time (dashed arrows) under the controls produced by the NN (illustrated for two set of target parameters $\alpha^{(0)}$ and $\alpha^{(1)}$). Training of the framework consists in optimizing the weights ϕ_{NN} of the NN such that the deviation \mathcal{I} between the controlled and target families of operations is minimized.

$U(t)$ with desired properties. This is formulated in terms of a cost functional $\mathcal{I}(H(t))$ to be minimized. A common example would be the task of realizing a target controlled–not gate $U^{tgt} = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes \sigma_x$ at a given time T , with the cost $\mathcal{I}(H(t)) = \|U(T) - U^{tgt}\|$ measuring the deviation between the controlled and target propagators. A corresponding task of family-control could be the realization of the family of target gates $U_\alpha^{tgt} = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes \exp(-i\alpha\sigma_x)$ with variable angle α . In this case, the overall task to be solved would be the identification of a continuum of Hamiltonians $H_\alpha(t)$, pa-

parameterized by the angle α , such that any of the propagators $U_\alpha = U_\alpha(t = T)$ induced by $H_\alpha(t)$ approximates the gate U_α^{tgt} as well as possible at $t = T$. The corresponding functional would thus become $\mathcal{I}(H_\alpha(t)) = \langle ||U_\alpha - U_\alpha^{tgt}|| \rangle_\alpha$ with an average over α .

The general formulation of such a family-control problem can be given in terms of the individual costs $\mathcal{I}_\alpha(H_\alpha(t))$, where the target parameter α can be a single scalar or a vector. The overall task to be solved is the identification of the continuum of time-dependent Hamiltonians $H_\alpha(t)$ that minimizes the averaged cost $\mathcal{I} = \langle \mathcal{I}_\alpha(H_\alpha(t)) \rangle_\alpha$.

In principle, this can be addressed as several control problems to be solved separately for a discretized set of target values $\{\alpha^{(i)}\}$, and, an additional step of interpolation for any new target with $\alpha \notin \{\alpha^{(i)}\}$. Hardly any control problem, however, has a unique solution, or at least a unique solution that can be found in practice. That is, there is no guarantee for two Hamiltonians $H_{\alpha^{(1)}}$ and $H_{\alpha^{(2)}}$ identified as optimal for similar values of $\alpha^{(1)}$ and $\alpha^{(2)}$, to be themselves similar. Any attempt to find an optimal Hamiltonian for a value of α between $\alpha^{(1)}$ and $\alpha^{(2)}$ in terms of an interpolation between $H_{\alpha^{(1)}}$ and $H_{\alpha^{(2)}}$ can thus result in a Hamiltonian that utterly fails to realize the desired task. To avoid this issue, it is desirable to require H_α to depend smoothly on α . Such requirement can be realized by means of an appropriate parameterization of the dependence of $H_\alpha(t)$ on both α and t . Given that neural networks (NNs) provide the flexible structure to approximate multivariate continuous functions up to arbitrary precision [30], these are deemed ideally suited for the task at hand.

Time-dependence in a Hamiltonian is typically realized in terms of temporally modulated electro-magnetic fields that appear as one or several control functions $f_\alpha(t)$ in the Hamiltonian. The scope of the NN is thus to model these functions. To this intent, the parameters α and the time t are taken to be the inputs of the NN, and the control values $f_\alpha(t)$ its outputs. An example of such a NN is illustrated in Fig. 1 for the case of two-dimensional parameters α and a single control function $f_\alpha(t)$; it is readily adapted to arbitrary dimensions of the parameters and number of controls by varying the sizes of the inputs and outputs accordingly.

The optimization – *i.e.*, training of the NN – can be achieved with a variety of techniques, but gradient-descent training has the advantage of simplicity and scalability to high-dimensional problems. Since the propagators U_α induced by the Hamiltonians H_α typically need to be constructed numerically, efficient means to take derivatives with respect to the control functions $f_\alpha(t)$ are essential. Recent advances in the field of automatic differentiation [29] give access to efficient differentiation over numerical solvers of differential equations. This allows to combine seamlessly gradients over the evolution of the system and over the weights of the NN. Finally, even though the evaluation of the averaged cost \mathcal{I} would always be based on a sum over discrete values of α rather than a proper integral, the output of the neural network is still continuous in α , and choosing different random sam-

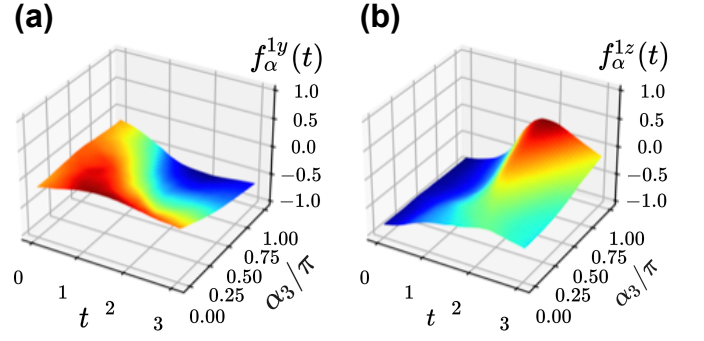


FIG. 2. Control of the family of arbitrary single-qubit rotations U_α^1 defined in Eq. (3) with the Hamiltonian in Eq. (1). The framework is successfully trained to implement any of the target rotations, resulting in an average infidelity of $\bar{\mathcal{I}} = 2 \times 10^{-4}$ (assessed on new targets not seen during training). To visualize the controls produced by the NN, α_1 and α_2 are kept fixed to a value of $3\pi/4$ and α_3 is varied in the range $[0, \pi]$. The amplitudes of the two control fields f_α^{1y} (panel a), and f_α^{1z} (panel b) produced by the NN are plotted as a function of both the time t and α_3/π .

pling points at each step in the training process avoids finding solutions with artefacts resulting from the sampling. Implementation details can be found in Sec. I of Supp. Matt. [31].

The following discussion exemplifies the framework sketched so far, with the realization of quantum gates induced by the n -qubit Hamiltonian

$$\mathcal{H}_\alpha(t) = \sum_{i < j=1}^n f_\alpha^{ij}(t) \sigma_x^{(i)} \sigma_x^{(j)} + \sum_{i=1}^n f_\alpha^{iy}(t) \sigma_y^{(i)} + f_\alpha^{iz}(t) \sigma_z^{(i)}, \quad (1)$$

with single-qubit Pauli $\sigma_y^{(i)}$ and $\sigma_z^{(i)}$ terms complemented with interactions $\sigma_x^{(i)} \sigma_x^{(j)}$ between pairs of qubits (i, j) . Overall, $C = 2n + n(n-1)/2$ time-dependent functions have to be learnt, including the single-qubit $f_\alpha^{iy}(t)$ and $f_\alpha^{iz}(t)$, and two-qubit controls $f_\alpha^{ij}(t)$. The Hamiltonian is sufficiently general so that any desired n -qubit unitary can be realized [32], but bounded control amplitudes $\{f_\alpha^{ij}, f_\alpha^{iy}, f_\alpha^{iz}\} \in [-1, 1]$ result in a finite minimal time required to realize a given unitary. Deviations between controlled and target gates are characterized by the gate infidelities

$$\mathcal{I}_\alpha(H_\alpha(t)) = 1 - \frac{1}{2^n} |\text{Tr}[U_\alpha^\dagger U_\alpha^{tgt}]|^2 \quad (2)$$

in the subsequent examples.

The basic workings of the framework can be illustrated with the task of realizing the manifold of single-qubit gates

$$U_\alpha^1 = \exp\left(-i\frac{\alpha_1}{2}\sigma_z\right) \exp\left(-i\frac{\alpha_2}{2}\sigma_y\right) \exp\left(-i\frac{\alpha_3}{2}\sigma_z\right) \quad (3)$$

for the three-dimensional target parameters α with components $\alpha_{j=1,2,3} \in [0, \pi]$, given the single-qubit version of $\mathcal{H}_\alpha(t)$ in Eq. (1) and a fixed gate time $T = \pi$. Here, and in all subsequent examples, the training stage is limited to 400 iterations, with each iteration corresponding to an average of




FIG. 3. Infidelities and times for the family U_α^1 of rotations when both the (target-dependent) control functions and times are learnt concurrently. For the sake of visualization, results are plotted for a two-dimensional subset of the three-dimensional family of targets, with fixed parameter $\alpha_3 = 3\pi/4$, and discretized $\alpha_1, \alpha_2 \in [0, \pi]$ over a grid of 75×75 regularly spaced points. For each of the corresponding target gates a heat-map indicates the values of the infidelities \mathcal{I}_α achieved (panel a, in logarithmic scale) and the control times T_α entailed by the framework (panel b).

the gate infidelity in Eq. (2) taken over 128 values of the parameters α uniformly sampled.

After training, the average gate infidelity $\bar{\mathcal{I}} = \langle \mathcal{I}_\alpha \rangle_\alpha$ resulting from the controls identified as optimal by the framework, is evaluated on an ensemble of 250 random values of α . Crucially, this average is taken with respect to new parameter values (*i.e.*, corresponding to targets not seen during training), and thus probes the ability of the framework to realize any gates belonging to the targeted family. In this example, the average infidelity is as low as $\bar{\mathcal{I}} = 2[3] \times 10^{-4}$, where the number in brackets indicates the standard deviation of the distribution. Fig. 2 depicts the two control functions $f_\alpha^{1y}(t)$ and $f_\alpha^{1z}(t)$ (panels (a) and (b) respectively) as function of α_3/π and t for $\alpha_1 = \alpha_2 = 3\pi/4$, substantiating that the solutions produced by the NN are indeed well-behaved, continuous functions of both the parameters α and time t .

A comparison between family-control and current QOC approaches is provided in the Sec. II. of the Supp. Mat. [31]. As discussed, the latter involves individual optimizations of a discrete set of gates, and a subsequent step of interpolation. Similar training complexity for the individual approach is found for a discretization of each α_i in Eq. (3) over less than $N_d = 10$ distinct values. Refining such coarse discretization results in substantially increased training effort – scaling cubically with N_d given the 3d nature of the gates family. More crucially, interpolating between individual control solutions is found to often yield close-to-vanishing fidelities.

In addition to an optimization of control functions, the present framework is also well suited to identify a minimal gate time T , or even minimal target-dependent gate times T_α (Sec. I B of Supp. Mat. [31]). The latter is achieved by introducing a second neural network with the gate time T_α as an output. Given the new cost $\mathcal{I}_\alpha + \mu \times T_\alpha$ comprised of the gate infidelity and the gate time as a penalty weighted with a scalar factor $\mu > 0$, this second neural network can be trained similarly to the case discussed above.

In Fig. 3 are reported the results from such an optimization with a small value $\mu = 10^{-2}$ of the weight suitable to find high-fidelity gates close to the minimally required time. Panel (a) depicts the infidelity \mathcal{I}_α of the resulting gates for $\alpha_3 = 3\pi/4$ as function of α_1 and α_2 . Typical values are smaller than 10^{-3} and the average infidelity with the average taken over all three components of α is $\bar{\mathcal{I}} = 4[5] \times 10^{-4}$.

Panel (b) depicts the target-dependent minimized gate times, with again a fixed value of $\alpha_3 = 3\pi/4$ but varied α_1 and α_2 . The shortest gate time is obtained for $\alpha_1 = \alpha_2 = 0$ (*i.e.*, for the target parameters $\alpha^{(0)} = [0, 0, 3\pi/4]$) in which case the constant control amplitudes $f_{\alpha^{(0)}}^{1z}(t) = 1$ and $f_{\alpha^{(0)}}^{1y}(t) = 0$ induce the desired gate $U_{\alpha^{(0)}}^1 = \exp(-i\frac{3\pi}{8}\sigma_z)$ after a time $T_{\alpha^{(0)}} = 3\pi/8$. The obtained gate times grow with increasing values of α_1 and α_2 , but always remain below the value of π used in the above example.

While the ability to realize single-qubit gates is of substantial practical value, it is certainly not the challenging control problem that helps to demonstrate the actual strength of the framework. This is better achieved in terms of two-qubit and three-qubit gates that are building blocks of quantum algorithms or digital quantum simulations.

TABLE I. Control under the Hamiltonian of Eq. (1) of families of $n = 2$ and 3 qubit gates, corresponding to $C = 5$ and 9 control functions to be learnt respectively. For each problem the family of target U_α^{tgt} considered and the domain χ of the target parameters α are reported. Results are provided in terms of the average (and standard deviations in bracket) infidelities $\bar{\mathcal{I}}$, and of the ratio R between the average gate times resulting from a decomposition in terms of elementary gates and the times necessitated by the framework. Additional elements of training the families (vi) and (vii) are provided in the main text.

U_α^{tgt}	χ	$\bar{\mathcal{I}}[10^{-4}]$	R
(i) $ 0\rangle\langle 0 \otimes I + 1\rangle\langle 1 \otimes \exp(-i\alpha_1\sigma_z^{(2)})$	$[0, \pi]$	1[1]	2.0
(ii) $\exp(-i\alpha_1\sigma_z^{(1)}\sigma_z^{(2)})$	$[0, \frac{\pi}{2}]$	0[0]	2.0
(iii) $ 0\rangle\langle 0 \otimes I + 1\rangle\langle 1 \otimes U_\alpha^1$	$[0, \pi]^3$	3[4]	2.1
(iv) $\exp(-i\sum_{j \in \{x,y,z\}} \alpha_j \sigma_j^{(1)} \sigma_j^{(2)})$	$[0, \frac{\pi}{2}]^3$	4[4]	2.6
(v) $\exp(-i\alpha_1\sigma_z^{(1)}\sigma_z^{(2)}\sigma_z^{(3)})$	$[0, \frac{\pi}{2}]$	1[0]	4.1
(vi) $\exp(-i\sum_{j \in \{x,y,z\}} \alpha_j \sigma_j^{(1)} \sigma_j^{(2)} \sigma_j^{(3)})$	$[0, \frac{\pi}{2}]^3$	9[8]	> 10
(vii) $(I - 11\rangle\langle 11) \otimes I + 11\rangle\langle 11 \otimes U_\alpha^1$	$[0, \pi]^3$	6[5]	> 10

Table I summarizes the results for a few selected families of two- and three-qubit gates with the domain χ of the parameters α depicted in column 3 and the obtained average infidelities (in multiples of 10^{-4}) in column 4 (further details of the NNs used are reported in Sec. I C of Supp. Mat. [31]). The two-qubit gates (i) to (iv) involve optimizations over $C = 5$ control functions, and the three-qubit gates (v) to (vii) involve $C = 9$ control functions.

Consistently with the previous findings, low infidelities $\bar{\mathcal{I}} < 5 \times 10^{-4}$ are achieved for any of the families of two-qubit gates (i-iv) and for the one-dimensional family of three-qubit gates (v).

A straightforward application of the above framework to

the problems (vi) and (vii) – corresponding to 9 time-dependent controls to be learnt and 3-dimensional families to be realized – however results in higher infidelities ($3[2.5] \times 10^{-3}$ for (vi) and $1.7[1.5] \times 10^{-3}$ for (vii)) than in the other cases. Yet, the results of the optimizations contain clear indications towards steps to reach higher fidelities that are now further discussed.

First, the lowest fidelities are systematically obtained for values of α close to the boundary of its admissible domain χ (as can also be seen in Fig. 3(a)). Enlarging the range of values used for training by 20% resolves this effect. Second, the control functions identified as optimal have general properties that can be exploited to reduce the number of independent functions that need to be learnt. In case (vi), the control solutions discovered by the framework satisfy the relation $f_{\alpha}^{13} = f_{\alpha}^{1z} = f_{\alpha}^{3z} = 0$ and in case (vii) they satisfy $f_{\alpha}^{1z} = f_{\alpha}^{2z} = 0$, $f_{\alpha}^{1y} = f_{\alpha}^{2y}$ and $f_{\alpha}^{13} = f_{\alpha}^{23}$. This indicates that only 6 and 5 independent control functions, out of the 9 possible, are needed for the cases (vi) and (vii) respectively. The infidelities listed in Table I, for families (vi) and (vii), result from an optimization with enlarged domain χ and reduced number of control functions, and their magnitude is comparable to those of the other cases.

While generally the non-uniqueness of solutions of optimal control problems makes it difficult to understand why a solution returned by a specific algorithm does achieve the goal that it is meant to achieve, it seems that the requirement of smooth dependence on the parameters α helps the NN to identify common features of all control pulses within the family, and to avoid unnecessary terms in the Hamiltonian that would obscure its working principle.

Beyond this conceptual benefit and the low infidelities achieved, the gain in gate time is also of high practical relevance. Since state-of-the-art implementation of unitaries on quantum devices rely on their decompositions in terms of elementary gates, the times T^{dec} entailed by such decompositions provide well-defined baselines. Given the freedom offered by the control Hamiltonian in Eq. (1) these decompositions are performed in terms of the gate-set of rotations generated by the single-qubit σ_y and σ_z and two-qubit $\sigma_x\sigma_x$ operators, for which qiskit’s [33] compiling routine is employed with the highest level of optimization available (Sec. II of Supp. Matt. [31]).

Column 5 of Table I depicts the ratio R between the averaged durations $\langle T_{\alpha}^{dec} \rangle_{\alpha}$ obtained with compiled gate circuits and the durations obtained with the present techniques. In all cases there is an improvement of at least a factor of 2, but, in cases (vi) and (vii), the improvement is substantially larger. This suggests that compilation techniques (*i.e.*, discrete optimizations) struggle with these complex three-qubit gates, whereas the continuous optimization realized in terms of NN does not suffer from these limitations.

The ability to accurately control entire families of gates in reduced time, especially for complex gates, highlights the benefits of family-control. Given that the automatic differentiation techniques [29], that ensures the efficient training of

the framework, can be applied to any system of ordinary differential equations (ODE), family-control can find direct application to a broad range of quantum systems, such as superconducting qubits. Since those are non-linear oscillators with a ladder of excited states, further studies would include suppression of leakage to such states. Similarly, trapped ions or opto-mechanical systems with several interacting degrees of freedom pose control problems that can be addressed with the present techniques.

Applying QOC to open systems [34–37] allows to take into account (and minimize) the detrimental effects of an external environment, and is within the direct reach of family-control. The dynamics of open quantum systems is mostly described by means of a master equation, effectively an ODE, and is thus amendable to the methodology described. While simulating such dynamics is inherently more demanding than for closed systems, this overhead can be mitigated by evolving only a constant number of carefully selected initial states [35]. Going further, automatic differentiation has now been extended to the treatment of stochastic differential equations [38, 39] such that family-control can also be applied to open systems simulated with quantum trajectories [37] and could even generalize to problems of control with active feedback [40, 41].

While optimal control is traditionally realized in terms of control pulses designed in numerical experiments, fundamental limitations in modelling and simulating the dynamics of composite quantum systems resulted in a shift towards designing control pulses in laboratory experiments [42–45]. Just like many techniques for individual control targets could be generalized to this setting, also family-control could be trained based exclusively on experimental data, either in situations where gradients can be experimentally estimated [46], or by resorting to gradient-free optimization strategies [47].

Essentially, the methodology that was presented here enables the control of a quantum system in different contexts. In the examples investigated, this context was in one to one correspondence with the target gate to be realized, that is, the overall details of the system under control were kept fixed and only the targets were varied. More generally, the scheme based on NNs allows to tailor controls to be applied to any relevant context variable. For instance, the inputs of the NN could also include intrinsic details of the controlled system (such as varied energy detunings [48] or sizes [49]) or extrinsic (such as environmental heating rates [50] or nearby operations inducing cross-talk [51]). Provided that the effects of these context variables can be simulated and that the corresponding optimal controls are expected to vary continuously with these variables, one would learn to accurately operate a quantum device in very broad situations.

We are indebted to stimulating discussions with Selwyn Simsek. This work is supported through a studentship in the Quantum Systems Engineering Skills and Training Hub at Imperial College London funded by EPSRC (Grant No. EP/P510257/1), and through funding from the QuantERA ER-ANET Cofund in Quantum Technologies implemented within the European Union’s Horizon 2020 Programme under the

project Theory-Blind Quantum Control TheBlinQC and from EPSRC under the grant EP/R044082/1.

-
- [1] J Werschnik and E K U Gross, “Quantum optimal control theory,” *J. Phys. B* **40**, R175–R211 (2007).
- [2] Steffen J. Glaser, Ugo Boscain, Tommaso Calarco, Christiane P. Koch, Walter Köckenberger, Ronnie Kosloff, Ilya Kuprov, Burkhard Luy, Sophie Schirmer, Thomas Schulte-Herbrüggen, Dominique Sugny, and Frank K. Wilhelm, “Training schrödinger’s cat: quantum optimal control,” *Eur. Phys. J. D* **69**, 279 (2015).
- [3] N. Timoney, V. Elman, S. Glaser, C. Weiss, M. Johanning, W. Neuhauser, and Chr. Wunderlich, “Error-resistant single-qubit gates with trapped ions,” *Phys. Rev. A* **77**, 052334 (2008).
- [4] Florian Dolde, Ville Bergholm, Ya Wang, Ingmar Jakobi, Boris Naydenov, Sébastien Pezzagna, Jan Meijer, Fedor Jelezko, Philipp Neumann, Thomas Schulte-Herbrüggen, Jacob Biamonte, and Jörg Wrachtrup, “High-fidelity spin entanglement using optimal control,” *Nature Communications* **5**, 3371 (2014).
- [5] T. Choi, S. Debnath, T. A. Manning, C. Figgatt, Z.-X. Gong, L.-M. Duan, and C. Monroe, “Optimal quantum control of multimode couplings between trapped ion qubits for scalable entanglement,” *Phys. Rev. Lett.* **112**, 190502 (2014).
- [6] Reinier W. Heeres, Philip Reinhold, Nissim Ofek, Luigi Frunzio, Liang Jiang, Michel H. Devoret, and Robert J. Schoelkopf, “Implementing a universal gate set on a logical qubit encoded in an oscillator,” *Nat. Commun.* **8**, 94 (2017).
- [7] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas, “Fast quantum logic gates with trapped-ion qubits,” *Nature* **555**, 75–78 (2018).
- [8] Ming Gong, Ming-Cheng Chen, Yarui Zheng, Shiyu Wang, Chen Zha, Hui Deng, Zhiguang Yan, Hao Rong, Yulin Wu, Shaowei Li, Fusheng Chen, Youwei Zhao, Futian Liang, Jin Lin, Yu Xu, Cheng Guo, Lihua Sun, Anthony D. Castellano, Haohua Wang, Chengzhi Peng, Chao-Yang Lu, Xiaobo Zhu, and Jian-Wei Pan, “Genuine 12-qubit entanglement on a superconducting quantum processor,” *Phys. Rev. Lett.* **122**, 110501 (2019).
- [9] Harry Levine, Alexander Keesling, Giulia Semeghini, Ahmed Omran, Tout T. Wang, Sepheer Ebadi, Hannes Bernien, Markus Greiner, Vladan Vuletić, Hannes Pichler, and Mikhail D. Lukin, “Parallel implementation of high-fidelity multiqubit gates with neutral atoms,” *Phys. Rev. Lett.* **123**, 170503 (2019).
- [10] A. Omran, H. Levine, A. Keesling, G. Semeghini, T. T. Wang, S. Ebadi, H. Bernien, A. S. Zibrov, H. Pichler, S. Choi, J. Cui, M. Rossignolo, P. Rembold, S. Montangero, T. Calarco, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, “Generation and manipulation of schrödinger cat states in rydberg atom arrays,” *Science* **365**, 570–574 (2019).
- [11] Arthur Larrouy, Sabrina Patsch, Rémi Richaud, Jean-Michel Raimond, Michel Brune, Christiane P. Koch, and Sébastien Gleyzes, “Fast navigation in a large hilbert space using quantum optimal control,” *Phys. Rev. X* **10**, 021058 (2020).
- [12] F. Poggiali, P. Cappellaro, and N. Fabbri, “Optimal control for one-qubit quantum sensing,” *Phys. Rev. X* **8**, 021059 (2018).
- [13] Phila Rembold, Nimba Oshnik, Matthias M. Müller, Simone Montangero, Tommaso Calarco, and Elke Neu, “Introduction to quantum optimal control for quantum sensing with nitrogen-vacancy centers in diamond,” *AVS Quantum Sci.* **2**, 024701 (2020).
- [14] John Preskill, “Quantum Computing in the NISQ era and beyond,” *Quantum* **2**, 79 (2018).
- [15] B. Foxen, C. Neill, A. Dunsworth, P. Roushan, B. Chiaro, A. Megrant, J. Kelly, Zijun Chen, K. Satzinger, R. Barends, F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, S. Boixo, D. Buell, B. Burkett, Yu Chen, R. Collins, E. Farhi, A. Fowler, C. Gidney, M. Giustina, R. Graff, M. Harrigan, T. Huang, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, P. Klimov, A. Korotkov, F. Kostritsa, D. Landhuis, E. Lucero, J. McClean, M. McEwen, X. Mi, M. Mohseni, J. Y. Mutus, O. Naaman, M. Neeley, M. Niu, A. Petukhov, C. Quintana, N. Rubin, D. Sank, V. Smelyanskiy, A. Vainsencher, T. C. White, Z. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis (Google AI Quantum), “Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms,” *Phys. Rev. Lett.* **125**, 120504 (2020).
- [16] Nathan Lacroix, Christoph Hellings, Christian Kraglund Andersen, Agustin Di Paolo, Ants Remm, Stefania Lazar, Sebastian Krinner, Graham J. Norris, Mihai Gabureac, Johannes Heinsoo, Alexandre Blais, Christopher Eichler, and Andreas Wallraff, “Improving the performance of deep quantum optimization algorithms with continuous gate sets,” *PRX Quantum* **1**, 110304 (2020).
- [17] Deanna M. Abrams, Nicolas Didier, Blake R. Johnson, Marcus P. da Silva, and Colm A. Ryan, “Implementation of xy entangling gates with a single calibrated pulse,” *Nat. Electron.* **3**, 744–750 (2020).
- [18] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen, and Steffen J. Glaser, “Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms,” *J. Magn. Reson.* **172**, 296–305 (2005).
- [19] Tommaso Caneva, Tommaso Calarco, and Simone Montangero, “Chopped random-basis quantum optimization,” *Phys. Rev. A* **84**, 022326 (2011).
- [20] Daniel M. Reich, Mamadou Ndong, and Christiane P. Koch, “Monotonically convergent optimization in quantum control using krotov’s method,” *J. Chem. Phys.* **136**, 104103 (2012).
- [21] Ehsan Zahedinejad, Joydip Ghosh, and Barry C. Sanders, “High-fidelity single-shot toffoli gate via quantum control,” *Phys. Rev. Lett.* **114**, 200502 (2015).
- [22] Jun Li, Xiaodong Yang, Xinhua Peng, and Chang-Pu Sun, “Hybrid quantum-classical approach to quantum optimal control,” *Phys. Rev. Lett.* **118**, 150503 (2017).
- [23] Nelson Leung, Mohamed Abdelhafez, Jens Koch, and David Schuster, “Speedup for quantum optimal control from automatic differentiation based on graphics processing units,” *Phys. Rev. A* **95**, 042318 (2017).
- [24] Shai Machnes, Elie Assémat, David Tannor, and Frank K. Wilhelm, “Tunable, flexible, and efficient optimization of control pulses for practical qubits,” *Phys. Rev. Lett.* **120**, 150401 (2018).
- [25] Marin Bukov, Alexandre G. R. Day, Dries Sels, Phillip Weinberg, Anatoli Polkovnikov, and Pankaj Mehta, “Reinforcement learning in different phases of quantum control,” *Phys. Rev. X* **8**, 031086 (2018).
- [26] Murphy Yuezhen Niu, Sergio Boixo, Vadim N. Smelyanskiy, and Hartmut Neven, “Universal quantum control through deep reinforcement learning,” *npj Quantum Inf.* **5**, 33 (2019).
- [27] Frédéric Sauvage and Florian Mintert, “Optimal quantum control with poor statistics,” *PRX Quantum* **1**, 020322 (2020).
- [28] Luuk Coopmans, Di Luo, Graham Kells, Bryan K. Clark, and Juan Carrasquilla, “Protocol discovery for the quantum control of majoranas by differentiable programming and natural evolu-

- tion strategies,” *PRX Quantum* **2**, 020332 (2021).
- [29] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud, “Neural ordinary differential equations,” in *Adv. Neural Inf. Process. Syst.*, Vol. 31 (2018).
 - [30] Kurt Hornik, “Approximation capabilities of multilayer feed-forward networks,” *Neural Networks* **4**, 251–257 (1991).
 - [31] See Supplemental Material at [url] for implementation details, a comparison of family-control to current QOC approaches, and description of the methodology used for the compilation of quantum gates in terms of an elementary gate-set. This includes Refs. [52–64].
 - [32] Sonia G Schirmer, H Fu, and Allan I Solomon, “Complete controllability of quantum systems,” *Phys. Rev. A* **63**, 063410 (2001).
 - [33] IBM collaboration, “Qiskit: An open-source framework for quantum computing,” (2021).
 - [34] T Schulte-Herbrüggen, A Spörl, N Khaneja, and S J Glaser, “Optimal control for generating quantum gates in open dissipative systems,” *Journal of Physics B: Atomic, Molecular and Optical Physics* **44**, 154013 (2011).
 - [35] Michael H Goerz, Daniel M Reich, and Christiane P Koch, “Optimal control theory for a unitary operation under dissipative evolution,” *New Journal of Physics* **16**, 055012 (2014).
 - [36] Christiane P Koch, “Controlling open quantum systems: tools, achievements, and limitations,” *Journal of Physics: Condensed Matter* **28**, 213001 (2016).
 - [37] Mohamed Abdelhafez, David I. Schuster, and Jens Koch, “Gradient-based optimal control of open quantum systems using quantum trajectories and automatic differentiation,” *Phys. Rev. A* **99**, 052327 (2019).
 - [38] Junteng Jia and Austin R Benson, “Neural jump stochastic differential equations,” in *Adv. Neural Inf. Process. Syst.*, Vol. 32 (2019).
 - [39] Christopher Rackauckas, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar, Dominic Skinner, Ali Ramadhan, and Alan Edelman, “Universal differential equations for scientific machine learning,” arXiv preprint (2020).
 - [40] Frank Schäfer, Pavel Sekatski, Martin Koppenhöfer, Christoph Bruder, and Michal Kloc, “Control of stochastic quantum dynamics by differentiable programming,” *Mach. Learn.: Sci. and Technol.* **2**, 035004 (2021).
 - [41] Riccardo Porotti, Antoine Essig, Benjamin Huard, and Florian Marquardt, “Deep reinforcement learning for quantum state preparation with weak nonlinear measurements,” arXiv preprint (2021).
 - [42] J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, P. J. J. O’Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Cleland, and John M. Martinis, “Optimal quantum control using randomized benchmarking,” *Phys. Rev. Lett.* **112**, 240504 (2014).
 - [43] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, “Leakage reduction in fast superconducting qubit gates via optimal control,” *npj Quantum Inf.* **7**, 14 (2021).
 - [44] Sean Greenaway, Frédéric Sauvage, Kiran E. Khosla, and Florian Mintert, “Efficient assessment of process fidelity,” *Phys. Rev. Research* **3**, 033031 (2021).
 - [45] Yuval Baum, Mirko Amico, Sean Howell, Michael Hush, Maggie Liuzzi, Pranav Mundada, Thomas Merkh, Andre RR Carvalho, and Michael J Biercuk, “Experimental deep reinforcement learning for error-robust gateset design on a superconducting quantum computer,” arXiv preprint (2021).
 - [46] Leonardo Banchi and Gavin E. Crooks, “Measuring Analytic Gradients of General Quantum Evolution with the Stochastic Parameter Shift Rule,” *Quantum* **5**, 386 (2021).
 - [47] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever, “Evolution strategies as a scalable alternative to reinforcement learning,” arXiv preprint (2017).
 - [48] Riccardo Porotti, Dario Tamascelli, Marcello Restelli, and Enrico Prati, “Coherent transport of quantum states by deep reinforcement learning,” *Commun. Phys.* **2**, 61 (2019).
 - [49] S. van Frank, M. Bonneau, J. Schmiedmayer, S. Hild, C. Gross, M. Cheneau, I. Bloch, T. Pichler, A. Negretti, T. Calarco, and S. Montangero, “Optimal control of complex atomic quantum systems,” *Sci. Rep.* **6**, 34187 (2016).
 - [50] Karl P Horn, Florentin Reiter, Yiheng Lin, Dietrich Leibfried, and Christiane P Koch, “Quantum optimal control of the dissipative production of a maximally entangled state,” *New J. Phys.* **20**, 123010 (2018).
 - [51] Adam Winick, Joel J. Wallman, and Joseph Emerson, “Simulating and mitigating crosstalk,” *Phys. Rev. Lett.* **126**, 230502 (2021).
 - [52] Michael Bartholomew-Biggs, Steven Brown, Bruce Christianson, and Laurence Dixon, “Automatic differentiation of algorithms,” *J. Comput. Appl. Math.* **124**, 171–190 (2000).
 - [53] Atilim Gunes Baydin, Barak A Pearlmutter, Alexey Andreyevich Radul, and Jeffrey Mark Siskind, “Automatic differentiation in machine learning: a survey,” *J. Mach. Learn. Research* **18** (2018).
 - [54] Ricky T. Q. Chen, “Pytorch implementation of differentiable ode solvers,” (2018).
 - [55] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala, “Pytorch: An imperative style, high-performance deep learning library,” in *Adv. Neural Inf. Process. Syst.*, Vol. 32 (2019).
 - [56] Diederik P Kingma and Jimmy Ba, “Adam: A method for stochastic optimization,” arXiv preprint (2014).
 - [57] Marc Claesen and Bart De Moor, “Hyperparameter search in machine learning,” in *Proc. of the 11th Metaheuristics Int. Conf.* (2015).
 - [58] James Bergstra and Yoshua Bengio, “Random search for hyperparameter optimization,” *Journal of machine learning research* **13** (2012).
 - [59] Xavier Glorot, Antoine Bordes, and Yoshua Bengio, “Deep sparse rectifier neural networks,” in *Proc. Fourteenth Int. Conf. Artif. Intel. Stat.*, Proceedings of Machine Learning Research, Vol. 15 (2011).
 - [60] Xiao-Ming Zhang, Zezhu Wei, Raza Asad, Xu-Chen Yang, and Xin Wang, “When does reinforcement learning stand out in quantum control? a comparative study on state preparation,” *npj Quantum Information* **5**, 85 (2019).
 - [61] Richard S Sutton and Andrew G Barto, *Reinforcement learning: An introduction* (MIT press, 2018).
 - [62] Jesper Hassleris Mohr Jensen, Jens Jakob Sørensen, Klaus Mølmer, and Jacob Friis Sherson, “Time-optimal control of collisional $\sqrt{\text{swap}}$ gates in ultracold atomic systems,” *Phys. Rev. A* **100**, 052314 (2019).
 - [63] Martín Larocca, Esteban Calzetta, and Diego A. Wisniacki, “Exploiting landscape geometry to enhance quantum optimal control,” *Phys. Rev. A* **101**, 023410 (2020).
 - [64] Martín Larocca, Esteban Calzetta, and Diego Wisniacki, “Fourier compression: A customization method for quantum

control protocols,” Phys. Rev. A **102**, 033108 (2020).