

# Taming Quantum Systems: A Tutorial for Using Shortcuts-To-Adiabaticity, Quantum Optimal Control, and Reinforcement Learning

Callum W. Duncan<sup>1,2</sup>, Pablo M. Poggi<sup>2,3,\*</sup>, Marin Bukov<sup>4</sup>, Nikolaj Thomas Zinner<sup>5,6</sup> and Steve Campbell<sup>7,8,9</sup>

<sup>1</sup>*Aegiq Ltd., Cooper Buildings, Arundel Street, Sheffield S1 2NS, United Kingdom*

<sup>2</sup>*Department of Physics, SUPA and University of Strathclyde, Glasgow G4 0NG, United Kingdom*

<sup>3</sup>*Center for Quantum Information and Control, Department of Physics and Astronomy, University of New Mexico, Albuquerque, New Mexico 87131, USA*

<sup>4</sup>*Max Planck Institute for the Physics of Complex Systems, Nöthnitzer Str. 38, 01187 Dresden, Germany*

<sup>5</sup>*Department of Physics and Astronomy, Aarhus University, DK-8000 Aarhus C, Denmark*

<sup>6</sup>*Kvantify Aps, DK-2100 Copenhagen, Denmark*

<sup>7</sup>*School of Physics, University College Dublin, Belfield, Dublin 4, Ireland*

<sup>8</sup>*Centre for Quantum Engineering, Science, and Technology, University College Dublin, Belfield, Dublin 4, Ireland*

<sup>9</sup>*Dahlem Center for Complex Quantum Systems, Freie Universität Berlin, Arnimallee 14, 14195 Berlin, Germany*

 (Received 13 May 2025; revised 7 August 2025; published 31 October 2025)

Precise manipulation of quantum effects at the atomic and nanoscale has become an essential task in ongoing scientific and technological endeavors. Quantum control methods are thus routinely exploited for research in areas such as quantum materials, quantum chemistry, and atomic and molecular physics, as well as in the development of quantum technologies like computing, simulation, and sensing. Here, we present a pedagogical introduction to the basics of quantum control methods in tutorial form, with the aim of providing newcomers to the field with the core concepts and practical tools to use these methods in their research. We focus on three areas: shortcuts to adiabaticity, quantum optimal control, and machine-learning-based control. We lay out the basic theoretical elements of each area in a pedagogical way and describe their application to a series of example cases. For these, we include detailed analytical derivations as well as extensive numerical results. As an outlook, we discuss quantum control methods in the broader context of quantum technologies development and complex quantum systems research, outlining potential connections and synergies between them.

DOI: [10.1103/j8c7-v2hd](https://doi.org/10.1103/j8c7-v2hd)

## CONTENTS

I. INTRODUCTION	2	E. Variational adiabatic gauge potential	11
A. Scope of this tutorial	3	F. Examples	12
B. Methodologies of quantum control	4	1. Example: The Landau-Zener model	12
II. SHORTCUTS TO ADIABATICITY	4	2. Example: The Ising model in one dimension	13
A. The adiabatic theorem	5	3. Example: The Ising model with all-to-all couplings (Lipkin-Meshkov-Glick model)	18
B. Exact counterdiabatic term	6	4. Example: Variational counterdiabatic driving	22
C. Adiabatic gauge potential	8	G. Experimental implementations of counterdiabatic driving	28
D. Counterdiabatic driving and eigenstate phase accumulation	9	III. QUANTUM OPTIMAL CONTROL	28
		A. Basic statement of an optimal control problem	29
		B. Gradient-based optimal control using GRAPE	31
		C. Examples	33

\*Contact author: [pablo.poggi@strath.ac.uk](mailto:pablo.poggi@strath.ac.uk)

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

1. Example: optimal control in the Landau-Zener model	33
2. Example: optimal control of a single-qubit gate	34
3. Example: generation of many-body entanglement	36
D. Experimental implementations of quantum optimal control	37
IV. REINFORCEMENT LEARNING FOR OPTIMAL QUANTUM CONTROL	38
A. Optimal control via policy gradients	40
B. Examples	41
1. Example: Universal single-qubit state preparation	42
2. Example: RL vs. counterdiabatic driving in the presence of trotterization errors	44
3. Example: Continuous single-qubit feedback control using quantum data	46
C. Implementations of reinforcement learning in experiment	50
V. ADVANTAGES, LIMITATIONS, AND HYBRID ALGORITHMS	51
VI. FUTURE CHALLENGES AND PROSPECTS FOR QUANTUM CONTROL	54
A. Quantum control to push fundamental limits in quantum dynamics	54
B. Connections between counterdiabatic driving, the adiabatic gauge potential, and chaos	55
C. Adiabatic gauge potential and quantum geometry	55
D. Shortcuts-to-adiabaticity and periodically driven (Floquet) systems	55
E. Quantum control to characterize non-equilibrium dynamics	56
F. Numerical approaches to the calculation of counterdiabatic driving and the adiabatic gauge potential	56
G. Hybrid control techniques	57
H. Quantum control of open systems	57
I. Quantum annealing and shortcuts-to-adiabaticity	57
J. Variational quantum algorithms and quantum optimal control	58
K. Quantum control for quantum computing	58
JUPYTER NOTEBOOKS	60
ACKNOWLEDGMENTS	60
DATA AVAILABILITY	60
REFERENCES	60

## I. INTRODUCTION

Ever since experiments gained access to the quantum regime, there has been a continuing need to develop and

perfect efficient techniques to manipulate quantum systems. These manipulations occur in real time, and the corresponding time-dependent control parameter can be described by a function of time, called the *control protocol*. The recent development of quantum-enhanced devices has led to a significant surge in activity which has witnessed some remarkable advances, both theoretical and experimental, in the use of these control techniques; for example quantum control has found fertile application in assessing the thermodynamics of quantum systems, providing a means to boost the performance of nanoscale quantum heat engines, to develop ultraprecise quantum sensors, and to robustly implement quantum gates [1,2].

Unlike classical mechanics, where the most common way of achieving control is to apply mechanical forces, in quantum control electromagnetic fields prevail, both in the dc and the ac regime. This can be traced back to the necessity to couple to Angstrom-sized atoms, and even smaller objects, where quantum effects become prevalent. Moving beyond the control of a single system to few and multi-particle settings introduces other control knobs besides external potentials, in particular, varying the interaction strength. For instance, in ultracold atoms this can be achieved using so-called Feshbach resonances [3], while in systems of optical tweezer arrays, one can directly control the positions of the trapped atoms which effectively fine-tunes the strength of the position-dependent interaction potential [4]. The ability to manipulate an arbitrary quantum system then becomes a multifaceted problem, one which must also be cognisant of other physically relevant constraints, for example, a finite control pulse energy and/or bandwidth, smoothness of the control protocol, or the intrinsic decoherence time which will bound the protocol time duration that can vary greatly between different physical platforms [1,2,5,6]. These constraints can feed directly into characterizing the efficacy of a control protocol with the most common figure of merit being the fidelity (quantum overlap), between the protocol-evolved final state,  $|\psi(t)\rangle$  and the desired target state,  $|\psi_*\rangle$ ,

$$\mathcal{F} = |\langle \psi_* | \psi(t) \rangle|^2 \quad (1)$$

where we have assumed the states in Eq. (1) are pure (as will be the case for all examples in this tutorial); however, we remark that these states could also be eigenstates or mixed ensembles. While we will focus on the use of fidelity as our key figure of merit throughout this tutorial, it is important to note that when it comes to many-body states this should be done with care since it often decreases exponentially with the number of particles. Indeed, if we consider a product state with single particle fidelity of 0.99, the many-body fidelity of a system with 100 particles is  $0.99^{100} \approx 0.36 \ll 0.99$ . In general there are many possible cost functions that can be used when optimizing the dynamics of a quantum system, for example, the energy

(density) is typically minimized in the process of preparing the ground state of a quantum system. More generally, one might be interested in maximizing and/or minimizing the expectation values of a specific observable, e.g., the magnetization, the momentum distribution, the strength of a correlation function or less frequently accessible quantities, such as quantum entanglement [7–10]. Regardless of the specific choice, the techniques presented below can be readily adapted for a given situation.

In large, spatially extended systems, it is often useful to consider two kinds of control protocols: global and local. Global protocols usually act (almost) homogeneously on the entire system, while local protocols can either be applied to a smaller part of the system or to individual constituents, e.g., the atoms of a chain using single-site addressability. Whatever the means of controlling a quantum system, it is important to keep in mind that the control protocols that can be emulated in experiments typically correspond to local (or sums of local) operators in the Hamiltonian description of the system; this is particularly relevant for many-body control. Thus, locality imposes a formidable, yet very physical, constraint on the accessibility of control fields [11]. The various control techniques that are known to be effective for single, isolated quantum systems must be carefully analyzed in view of this constraint when the system size is scaled up. Providing an introduction to the tools and techniques necessary to understand the interplay between implementability and scalability of control protocols for quantum systems is the focus of this tutorial.

### A. Scope of this tutorial

We aim to provide a succinct entry point to some of the state-of-the-art approaches to controlling quantum systems. There are many excellent reviews available on these and related topics, see for example Refs. [1,2,5,12–18], which the interested reader can delve into. Here, we aim to achieve two goals: Firstly, to provide a highly pedagogical introduction to the most prevalent techniques used to control many-body systems; thus, in addition to the technical details of the control protocols themselves, we

also provide extensive details on other important tools for solving paradigmatic many-body systems. Secondly, by introducing a suite of techniques arising from the various subcommunities that work on quantum control, our aim is to highlight the benefits that can arise from employing hybridized control techniques, i.e., those that take advantage of aspects of several control strategies.

With these aims in mind, in what follows we present the basic ingredients for three of the leading approaches to achieving non-adiabatic control (Fig. 1): shortcuts-to-adiabaticity, quantum optimal control, and reinforcement learning. They are presented in isolation (and in no particular order) to ensure that the core theoretical underpinnings of each approach are presented in a suitably pedagogical manner. As such, the interested reader can dive straight into a given technique without the need to work through any preceding sections. For a given task, the choice of which approach to employ depends on a range of competing factors. The experimental architecture clearly imposes certain constraints, e.g., the available set of controls and/or measurements. Beyond such hardware limitations, however, other requirements necessarily requested from the dynamics will also impact the choice of control technique to use. For instance, if the system must follow a specific dynamical path in Hilbert space, e.g., remaining in the ground state, then the suite of tools from shortcuts-to-adiabaticity are likely to be the most suitable choice. However, in many situations the intermediate dynamics is not of any particular relevance and we are only interested in, e.g., state transfer or state preparation. In these cases, often the goal is simply to find a protocol that works under the given constraints and, therefore, the task becomes principally about pulse shape optimization where optimal control provides several efficient approaches. If, on the other hand, one is interested in automating the procedure of finding optimal protocols, then adopting a deep learning approach is advisable, as it allows trained RL agents to generalize and produce control protocols even for parameter values not encountered during optimization; moreover, RL algorithms are particularly useful for on-line feedback control based on (partial) measurements. In Sec. V we provide further discussion along these lines,




FIG. 1. Schematic of the main three approaches to quantum control that are discussed in this tutorial article: shortcuts to adiabaticity (Sec. II), optimal control (Sec. III), and reinforcement learning (Sec. IV).

aiming to highlight some of the strengths and weaknesses of the techniques that we discuss in this tutorial and highlight the benefit that comes from combining techniques to tackle a problem.

Another motivation for this tutorial, that we believe sets this work apart from other reviews and pedagogical resources on the topic, is to provide an “introductory one-stop-shop” for the various control subcommunities. To facilitate this, we explicitly demonstrate the application of all protocols for the manipulation of a single two-level quantum system, in addition to several other relevant examples. These additional examples focus on some of the most common areas to which a given technique is applied. Therefore in Sec. II we consider the exact and approximate control of critical many-body systems, Sec. III provides examples relevant to quantum computation focussing on the implementation of gates and the generation of multipartite entanglement, while Sec. IV explores universal single-qubit state preparation and active qubit feedback control using an ancilla. In the hope of bringing some consistency to the nomenclature, we also provide tables summarizing the notation employed when introducing a given technique in Table I for shortcuts-to-adiabaticity, Table III for optimal control, and Table IV for reinforcement learning. Finally, the reader can refer to Table VI for a list of abbreviations commonly used throughout this tutorial. We provide Jupyter notebooks for interested readers to explore further the control techniques introduced on GitHub [19].

## B. Methodologies of quantum control

We can broadly define three classes of quantum systems of increasing control complexity: Few (including single) particle systems are often used to demonstrate the basics of control methods and can also benchmark numerical quantum control algorithms against exactly solvable problems. However, they also bear relevance to real systems such as qubits (two-level systems) and qutrits or  $\Lambda$  systems (three-level systems). In these settings common control problems include state population transfer or finding optimal single- or two-qubit gates. The relative simplicity of the models typically can allow for the inclusion of dissipation effects [20]. The next class of systems that we will be primarily interested in are fermionic band models and weakly interacting bosons described by quadratic Hamiltonians. Examples in this category include topological insulators, weakly-interacting bosons, the transverse-field Ising model, and the Kitaev honeycomb model. The translation invariant counterparts of such models are exactly solvable in momentum space, which offers the possibility to explore some control techniques analytically. In particular, two-band models can be viewed as a collection of independent two-level systems labeled by (quasi)momentum. As

we will demonstrate, control over the individual momentum modes allows for optimal control of the entire system, however, the momentum-dependence ultimately compromises the locality of the applied control fields in real space. Nevertheless, useful insight into the most relevant operators for achieving control can be learned which can prove valuable in moving to the final class of systems. The most complex control class comprises many-body systems with no known exact solutions, e.g., non-integrable models with Hubbard- and Heisenberg-type Hamiltonians. In general we can expect here optimal control techniques to work whenever the state of interest is gapped in energy; for gapless models much less is known regarding their controllability, and there are many open questions at present.

Broadly speaking, we distinguish between adiabatic and non-adiabatic techniques when classifying quantum control protocols. Adiabatic protocols (which we will rigorously define for quantum systems in the proceeding sections) make use of the discrete structure of quantized spectra and high-fidelity control timescales are inversely proportional to the energy gaps of the controlled level to its neighboring levels. A famous example is the Landau-Zener model, where an exponentially small fraction is left in the excited state after the completion of the adiabatic process [21]. Stimulated rapid adiabatic passage (STIRAP) [22] is another technique which exploits spectral characteristics to achieve robust state transfer. In this tutorial, we will focus explicitly on non-adiabatic techniques. This class encompasses counterdiabatic (CD) protocols that generate a transitionless time evolution with respect to the instantaneous basis of the Hamiltonian describing the system of interest, which is typically far from the adiabatic regime. It also captures the broader set of Hamiltonians that are tailored to bring the controlled system into a target state within a fixed protocol duration (but allows for the creation of excitations during the evolution so long as they are removed before the protocol comes to an end) as is typical in many quantum optimal control approaches.

## II. SHORTCUTS TO ADIABATICITY

Many practical tasks involve quantum adiabatic processes and, therefore, a significant body of work has amassed in developing techniques that either exploit the inherent adiabatic dynamics characteristic of slowly driven systems or, as will be the focus of this section, achieve an effective adiabatic dynamics in a finite time. We will first carefully define the concept of adiabaticity before focusing on counterdiabatic driving, which is arguably the conceptually simplest control method for achieving the desired evolution and works by directly suppressing any non-adiabatic transitions. We will review some modern reformulations and extensions of this idea, in particular discussing how to construct the control fields through

TABLE I. Summary of notation for the most common quantities in shortcuts-to-adiabaticity. Note that the symbol  $\mathcal{A}$  (and any of parameter dependent variations) always refers to the same object, i.e., the exact counterdiabatic term.

Object	Proposed Symbol
Original system Hamiltonian	$H(t)$
Berry correction (no diagonal elements, also known as Kato potential)	$\mathcal{A}, \mathcal{A}_\lambda, \mathcal{A}(t)$
Gauge potential (table point iv) [ $U$ 's]	$\dot{\lambda}(i\partial_\lambda U)U^\dagger$
Total generator (original system Hamiltonian + controls) for Kato/Berry	$H_{\text{CD}}(t) = H(t) + \mathcal{A}_\lambda$
Total generator (original system Hamiltonian + controls) for AGP	$H_{\text{CD}}(t) = H(t) + \dot{\lambda}(i\partial_\lambda U)U^\dagger$
Approximate AGP	$A_\lambda^{(x)}$ or $A_\lambda^{[x]}$

the adiabatic gauge potential and how to determine local approximations to achieving the control using variational techniques. We will explicitly consider the control of a simple two-level system described by the Landau-Zener model in Sec. II F 1, which will serve as a paradigmatic example that will also be considered when discussing the other control techniques that are a focus of this tutorial, namely optimal control in Sec. III C 1 and reinforcement learning in Secs. IV B 1 and IV B 2. We will also consider the application of these techniques to simple many-body settings embodied by three variations of the Ising model: (i) nearest neighbor, (ii) infinite range, and (iii) non-integrable.

### A. The adiabatic theorem

It is important from the outset to give a precise definition of an adiabatic process to avoid some common confusion: we recall that, in thermodynamics, adiabatic processes are those that do not transfer heat between the system and its environment, in the sense that they occur on timescales faster than the scales required for heat exchange; they may be slow or fast. By contrast, in quantum mechanics, as we shall see below, adiabatic evolution is defined by the lack of population transfer between an eigenstate of interest undergoing unitary evolution and the rest of the states in the instantaneous spectrum. Note that the energy of the state of interest can change in the process of adiabatic evolution, as the drive does work on the system, but this does not give rise to any heat; heat in quantum mechanics is defined microscopically as arising from any transitions between the instantaneous eigenstates of the Hamiltonian, caused by a non-adiabatic evolution. Adiabaticity in quantum mechanics is thus different from the notion of adiabaticity in classical thermodynamics; in fact, quantum adiabaticity is closer to what is called a quasistatic process in thermodynamics where the system remains in equilibrium throughout the evolution [23]. In this tutorial, the concept of adiabaticity refers to quantum mechanics.

To understand control protocols that achieve an effectively adiabatic evolution, let us now formalize the concept of adiabaticity for quantum dynamics. The suite of techniques termed “shortcuts-to-adiabaticity” arise precisely

from exploring approaches that mimic an adiabatic dynamics on, in principle, arbitrary timescales. Thus, we begin with a recapitulation of the adiabatic theorem, which already provides a wealth of knowledge regarding the types of generators that allow for the control of a quantum system. There are several rigorous derivations of the adiabatic theorem, starting from the earliest incarnations arising from Erhenfest’s analysis and the proof by Born and Fock [24] which have later been refined and extended by Kato to the case of degenerate eigenstates [25]. Herein, we follow a similar treatment found in many standard texts and reviews, e.g., Ref. [26].

Consider a quantum system described by the time-dependent non-degenerate Hamiltonian  $H(t)$  which evolves according to the time-dependent Schrödinger equation

$$i\hbar\partial_t|\psi(t)\rangle = H(t)|\psi(t)\rangle. \quad (2)$$

Furthermore, it is clear that at every instant in time the system also satisfies the time-independent Schrödinger equation

$$H(t)|n(t)\rangle = E_n(t)|n(t)\rangle \quad (3)$$

where  $\{|n(t)\rangle\}$  and  $\{E_n(t)\}$  are the *instantaneous* eigenstates and eigenenergies of the system, i.e., this is simply the eigenvalue equation for the matrix  $H(t)$  treating  $t$  as a parameter. Since the states  $\{|n(t)\rangle\}$  form a valid basis for the Hilbert space of the system at every instant of time, we can use them to express the solution to Eq. (2), i.e., we can assume an ansatz solution of the form

$$|\psi(t)\rangle = \sum_n c_n(t)|n(t)\rangle \quad (4)$$

with  $c_n(t)$  some complex coefficients. Putting Eq. (4) into Eq. (2) and using the spectral decomposition of  $H(t)$  we have

$$\begin{aligned} i\hbar \left[ \partial_t \left( \sum_n c_n(t) |n(t)\rangle \right) \right] &= H(t) \left[ \sum_n c_n(t) |n(t)\rangle \right] \\ &= \sum_k E_k(t) |k(t)\rangle \langle k(t)| \left[ \sum_n c_n(t) |n(t)\rangle \right]. \end{aligned} \quad (5)$$

Dropping the explicit time-dependence for brevity [i.e.,  $H(t) \equiv H$ ,  $|n(t)\rangle \equiv |n\rangle$ , and  $E_n(t) \equiv E_n$ ] and denoting the time derivative of any quantity  $x(t)$  by  $\partial_t x(t) \equiv \dot{x}$

$$\begin{aligned} i\hbar \left[ \sum_n (\dot{c}_n |n\rangle + c_n |\dot{n}\rangle) \right] &= \sum_{n,k} c_n E_k |k\rangle \langle k|n\rangle \\ &= \sum_n c_n E_n |n\rangle. \end{aligned} \quad (6)$$

To get a set of differential equations for the coefficients  $c_n$  we project onto the  $m$ th eigenstate,  $\langle m|$ , of  $H$ ,

$$i\hbar \left[ \sum_n (\dot{c}_n \langle m|n\rangle + c_n \langle m|\dot{n}\rangle) \right] = \sum_n c_n E_n \langle m|n\rangle, \quad (7)$$

which again due to orthonormality allows us to simplify most of the sums. Separating terms with  $m = n$  from  $m \neq n$  we find

$$\begin{aligned} i\hbar \left( \dot{c}_m + c_m \langle m|\dot{m}\rangle + \sum_{n \neq m} c_n \langle m|\dot{n}\rangle \right) &= c_m E_m \quad (8) \\ \implies i\hbar \dot{c}_m = c_m E_m - i\hbar c_m \langle m|\dot{m}\rangle - i\hbar \sum_{n \neq m} c_n \langle m|\dot{n}\rangle \\ &= (E_m - i\hbar \langle m|\dot{m}\rangle) c_m - i\hbar \sum_{n \neq m} c_n \langle m|\dot{n}\rangle. \end{aligned} \quad (9)$$

The key to the adiabatic theorem lies in the last term of Eq. (9). If this term is negligible, the evolution of the coefficients,  $c_m$ , obeys a simple first-order differential equation

$$\begin{aligned} \frac{\dot{c}_m}{c_m} &= -\frac{i}{\hbar} (E_m - i\hbar \langle m|\dot{m}\rangle) \\ \implies c_m &= \exp \left[ -\frac{i}{\hbar} \int_0^t E_m - i\hbar \langle m|\dot{m}\rangle ds \right]. \end{aligned} \quad (10)$$

This can be conveniently cast as

$$c_m(t) = \exp[-i(\phi_d(t) + \phi_g(t))], \quad (11)$$

with the dynamic phase  $\phi_d(t) = (1/\hbar) \int_0^t E_m(s) ds$  and the geometric phase  $\phi_g(t) = -\int_0^t i \langle m(s)|\partial_s m(s)\rangle ds$ , where we have briefly reintroduced the explicit time-dependence. Clearly, if the system begins in the  $n$ th eigenstate of  $H(0)$ , then the only non-trivial equation for the coefficients will correspond to the initial condition  $c_n(0) = 1$ , and it will remain in this state only accumulating the dynamical and geometric phases [27].

Thus, to understand the adiabatic theorem we must examine the last term in Eq. (9) more carefully. Notice

that this term effectively captures the instantaneous coupling between eigenstates due to the rate of change of the Hamiltonian. In order for this term to be negligible, we require  $\langle m|\dot{n}\rangle \rightarrow 0$ . Let us analyze this term more closely by taking the time derivative of Eq. (3)

$$\begin{aligned} \partial_t (H |n\rangle) &= \partial_t (E_n |n\rangle) \\ \dot{H} |n\rangle + H |\dot{n}\rangle &= \dot{E}_n |n\rangle + E_n |\dot{n}\rangle. \end{aligned} \quad (12)$$

Projecting again onto the  $m$ th state and recalling that the sum in Eq. (9) involves only  $m \neq n$

$$\begin{aligned} \langle m|\dot{H} |n\rangle + \langle m|H |\dot{n}\rangle &= \dot{E}_n \langle m|n\rangle + E_n \langle m|\dot{n}\rangle \\ \langle m|\dot{H} |n\rangle + E_m \langle m|\dot{n}\rangle &= E_n \langle m|\dot{n}\rangle \\ \implies \langle m|\dot{n}\rangle &= \frac{\langle m|\dot{H} |n\rangle}{E_n - E_m}. \end{aligned} \quad (13)$$

This equation captures the adiabatic condition by precisely expressing the error accumulated due to a non-adiabatic drive. While we shall focus on the insight into the controllability of quantum systems, it is worth stressing that the adiabatic theorem provides a versatile starting point for a range of physical settings, including models of quantum computation [28], topological states of matter [26], and for understanding the dynamics of critical systems according to the Kibble-Zurek mechanism [29–35]. While the preceding derivation applies to a very broad range of settings, we note that it requires modification when, e.g., the system exhibits degeneracies in its energy spectrum [25] or for periodically driven systems [36–38].

## B. Exact counterdiabatic term

While Eq. (13) has been well known for almost a century, only recently has it been shown that this term can be used to determine precisely the correction that needs to be added to the generator of the dynamics in order to suppress any non-adiabatic transitions for an arbitrary evolution governed by some time-dependent process,  $\lambda(t)$ . Demiralp and Rice [39,40], and independently Berry [41], showed that augmenting the system Hamiltonian with the so-called counterdiabatic (CD) term ensures that the dynamics precisely tracks an instantaneous adiabatic evolution. This CD term follows directly from Eq. (13) and is given by

$$\mathcal{A} \equiv \mathcal{A}_\lambda \equiv \mathcal{A}(t) = i\hbar \sum_n \sum_m \frac{|m\rangle \langle m|\dot{H} |n\rangle \langle n|}{E_n - E_m}, \quad n \neq m. \quad (14)$$

An important note: in order to bring some consistency to notation, here and throughout the symbol  $\mathcal{A}$  (including any parameter dependent variations) will be reserved for the

*exact* counterdiabatic term, while (variations) of the symbol  $A$  will be used to denote approximations. We refer to Tables I, III, and IV, which should serve as a reference point to clarify the conventions used throughout this tutorial.

We assume that there are no degeneracies in the spectrum, so the denominator in Eq. (14) is finite for all times,  $t$ . If we consider evolving a specific initial eigenstate of  $H(0)$ , the unitary dynamics governed by the Hamiltonian

$$H_{\text{CD}} = H + \mathcal{A}, \quad (15)$$

ensures that at every instant of time, the state of the system is guaranteed to be found (up to an overall phase) in the corresponding instantaneous eigenstate of the original Hamiltonian,  $H$ . Since this holds for any choice of eigenstate, it follows that this Hamiltonian will also realize an adiabatic dynamics for an arbitrary state of the system.

While we have arrived at Eq. (14) by carefully deriving the adiabatic condition, Eq. (13), this term can be directly determined starting from a more “traditional” approach of inverse engineering the dynamics such that a given target state is exactly achieved [41]. Assuming the system starts in the eigenstate  $|n(0)\rangle$  (where for clarity we have explicitly included the time-dependence) an adiabatic evolution implies that the state at some final time,  $T$ , will have accumulated the dynamical and geometric phases only, i.e., it will be of the form

$$\begin{aligned} |\psi_n(T)\rangle &= \exp[-i(\phi_d(T) + \phi_g(T))] |n(T)\rangle \\ &= \exp\left[-\frac{i}{\hbar} \int_0^T E_n(t) - i\hbar \langle n(t)|\partial_t n(t)\rangle dt\right] |n(T)\rangle. \end{aligned} \quad (16)$$

The goal of inverse engineering is to determine the Hamiltonian that will evolve the initial state to precisely this adiabatically connected target state, i.e.,

$$|\psi(T)\rangle = U(T) |n(0)\rangle \equiv |\psi_n(T)\rangle. \quad (17)$$

To determine what Hamiltonian generates such dynamics, we use the fact that any unitary,  $U(t)$ , satisfies

$$i\hbar \partial_t U(t) = H' U(t) \implies H' = i\hbar (\partial_t U(t)) U^\dagger(t), \quad (18)$$

where we have omitted the explicit time-dependence of  $H'$  for clarity. Given that our target is  $|\psi_n(T)\rangle$  our required unitary is therefore

$$\begin{aligned} U(T) &= \sum_n U_n(T) \\ &= \sum_n \exp\left[-\frac{i}{\hbar} \int_0^T E_n(t) - i\hbar \langle n(t)|\partial_t n(t)\rangle dt\right] \\ &\quad \times |n(T)\rangle \langle n(0)|. \end{aligned} \quad (19)$$

We can now use Eq. (18) to determine the corresponding Hamiltonian. Recalling that  $\partial_t e^{f(t)} = f'(t) e^{f(t)}$  and for brevity defining  $f(T) = \exp\left[-(i/\hbar) \int_0^T E_n(t) - i\hbar \langle n(t)|\partial_t n(t)\rangle dt\right]$ , we can determine the ingredients necessary for evaluating Eq. (18) by focusing on a single value of  $n$ ,

$$\begin{aligned} \partial_t U_n(t) &= \left(-\frac{i}{\hbar} E_n - \langle n(t)|\partial_t n(t)\rangle e^{f(t)}\right) |n(t)\rangle \langle n(0)| + e^{f(t)} |\partial_t n(t)\rangle \langle n(0)| \\ U_n^\dagger(t) &= e^{f(t)*} |n(0)\rangle \langle n(t)| \\ \implies i\hbar (\partial_t U_n) U_n^\dagger &= i\hbar \left(\left(-\frac{i}{\hbar} E_n(t) - \langle n(t)|\partial_t n(t)\rangle e^{f(t)}\right) |n(t)\rangle \langle n(0)| + e^{f(t)} |\partial_t n(t)\rangle \langle n(0)|\right) \times e^{f(t)*} |n(0)\rangle \langle n(t)| \\ &= E_n |n(t)\rangle \langle n(t)| - i\hbar \langle n(t)|\partial_t n(t)\rangle e^{f(t)} e^{f(t)*} |n(t)\rangle \langle n(t)| + i\hbar e^{f(t)} e^{f(t)*} |\partial_t n(t)\rangle \langle n(t)|. \end{aligned} \quad (20)$$

While not immediately obvious, the term  $\langle n(t)|\partial_t n(t)\rangle$  is purely imaginary [26], and therefore the function defined by  $f(t)$  is imaginary; to see this, one can take the derivative of the normalization condition  $\langle n(t)|n(t)\rangle = 1$  and rearrange terms. We can simplify the above expression since  $e^{f(t)} e^{f(t)*} = 1$ , and therefore the Hamiltonian that gives rise to a perfectly adiabatic dynamics is given by

$$\begin{aligned} H' &= H_{\text{CD}} = \sum_n E_n(t) |n(t)\rangle \langle n(t)| + i\hbar \sum_n (|\partial_t n(t)\rangle \langle n(t)| - \langle n(t)|\partial_t n(t)\rangle |n(t)\rangle \langle n(t)|) \\ &= H + \mathcal{A}. \end{aligned} \quad (21)$$

where the second term on the right-hand side is identical to Eq. (14) [41].

This formulation is the starting point for a wealth of studies regarding the control of quantum systems [2]. As we shall see, when directly applied to many-body systems, the CD term generally gives rise to a highly non-local interaction term thus limiting its direct experimental applicability to single or few-body systems [42,43]. Nevertheless, the determination of Eq. (14) still provides useful insight into the types of interactions that render a complex system controllable and this insight can then be used to restrict the control Hamiltonian to experimentally relevant interactions. As we will now discuss, the adiabatic gauge potential provides an alternative approach to deriving the CD term that allows us to achieve this aim.

### C. Adiabatic gauge potential

An alternative method to think about the non-adiabatic terms in Eq. (13) is through the adiabatic gauge potential (AGP) [44,45]. We begin by considering a state  $|\psi(\lambda(t))\rangle$  evolving under the time-dependent Schrödinger equation of Eq. (2) where we change slightly notation such that we describe the time-dependence by the parameter  $\lambda(t)$ . From here, we will drop the explicit time-dependence and retain only the  $\lambda$ -dependence. In principle, we can always construct the unitary transformation  $U(\lambda)$  to a “co-moving” reference frame which diagonalizes the Hamiltonian for all times. Given that a diagonal Hamiltonian cannot induce any diabatic transitions, states evolving under it simply acquire a phase with frequencies corresponding to the eigenenergies. However, we know that physics must be consistent between reference frames. Therefore, performing the unitary transformation on the equation of motion in this case must produce additional non-diagonal terms and these are given by the adiabatic gauge potential. This is indeed the case, as the transformation of a Hamiltonian by a general unitary for which  $\partial_t U \neq 0$

gives

$$H(t) \rightarrow U^\dagger(t)H(t)U(t) - i\hbar U^\dagger(t)\partial_t U(t). \quad (22)$$

Converting this to the specific problem at hand, using the unitary which diagonalizes the Hamiltonian we get

$$\begin{aligned} H(\lambda) &\rightarrow U^\dagger(\lambda)H(\lambda)U(\lambda) - i\hbar \dot{\lambda} U^\dagger(\lambda)\partial_\lambda U(\lambda) \\ &= \tilde{H}(\lambda) - \dot{\lambda}\tilde{\mathcal{A}}_\lambda, \end{aligned} \quad (23)$$

where we have used the fact that  $\partial_t = \dot{\lambda}\partial_\lambda$ , with tildes denoting operators that are in the co-moving frame and we have defined the AGP in the co-moving frame as

$$\tilde{\mathcal{A}}_\lambda = i\hbar U^\dagger(\lambda)\partial_\lambda U(\lambda). \quad (24)$$

We can find the adiabatic gauge potential in the original reference frame by inspecting its matrix elements with respect to a fixed parameter independent basis,  $\{|n\rangle_0\}$ , with  $|m(\lambda)\rangle = \sum_n U_{nm}(\lambda) |n\rangle_0$ , giving [45]

$$\begin{aligned} \langle m|_0 \tilde{\mathcal{A}}_\lambda |n\rangle_0 &= i\hbar \langle m|_0 U^\dagger(\lambda)\partial_\lambda U(\lambda) |n\rangle_0 \\ &= i\hbar \langle m(\lambda)| \partial_\lambda |n(\lambda)\rangle = \langle m(\lambda)| \mathcal{A}_\lambda |n(\lambda)\rangle, \end{aligned} \quad (25)$$

meaning that the gauge potential in the lab frame can be thought of as proportional to the derivative along the parameter path,

$$\mathcal{A}_\lambda = i\hbar \partial_\lambda = i\hbar [\partial_\lambda U(\lambda)]U^\dagger(\lambda). \quad (26)$$

If we modify the original Hamiltonian by adding the AGP then we obtain

$$H(\lambda) + \dot{\lambda}\mathcal{A}_\lambda \rightarrow \tilde{H}(\lambda) + \dot{\lambda}\tilde{\mathcal{A}}_\lambda - \dot{\lambda}\tilde{\mathcal{A}}_\lambda = \tilde{H}(\lambda), \quad (27)$$

and since  $\tilde{H}(\lambda)$  is diagonal, it cannot cause diabatic transitions between the eigenstates. Thus, addition of the AGP

TABLE II. The table summarizes the acquired dynamical and geometric phases following different types of adiabatic evolution  $i\partial_t|\psi(t)\rangle = \mathcal{H}(t)|\psi(t)\rangle$  with  $|\psi(t)\rangle = |n(0)\rangle$ , and explains the relation between the generators of counterdiabatic evolution introduced in the text. The geometric phase is a property of the eigenstate manifold and remains the same irrespective of the Hamiltonian  $\mathcal{H}$  used to generate the evolution; however, the choice of  $\mathcal{H}$  modifies the dynamics and hence the total phase accumulated by the wavefunction. The dynamical and geometric phases are defined by  $\phi_d^{(n)}(t) = \int_0^t E_n(\lambda(s))ds$  and  $\phi_g^{(n)}(t) = -\int_0^t \langle n(\lambda(s))|i\partial_s n(\lambda(s))\dot{\lambda}(s)\rangle ds$ , respectively; the gauge  $\chi(\lambda(t)) = \int_0^t \partial_s \chi(\lambda(s))\dot{\lambda}(s)ds$  is defined in Eq. (32).

	$\mathcal{H} = H(t)$ in the adiabatic limit: $T \rightarrow \infty$ , $\dot{\lambda} \rightarrow 0$ , $\dot{\lambda}T \rightarrow const.$	$\mathcal{H} = H(t) + \dot{\lambda}\mathcal{A}_\lambda = H_{CD}$ (Kato gauge potential) parallel-transport gauge	$\mathcal{H} = H(t) + (i\partial_t U)U^\dagger = H_{CD}$ , $\chi(\lambda) = 0$ gauge	$\mathcal{H} = \mathcal{A}_\lambda$ (Kato gauge potential), parallel-transport gauge	$\mathcal{H} = (i\partial_t U)U^\dagger = -iU\partial_t U^\dagger$ , $\chi(\lambda) = 0$ gauge	gauge potential in generic gauge $\chi(\lambda)$
dynamical phase $\phi_d$	$\phi_d(t)$	$\phi_d(t)$	$\phi_d(t) - \phi_g(t)$	0	$-\phi_g(t)$	$-\phi_g(t) + \chi(\lambda(t))$
geometric phase $\phi_g$	$\phi_g(t)$	$\phi_g(t)$	$\phi_g(t)$	$\phi_g(t)$	$\phi_g(t)$	$\phi_g(t)$
total phase $\phi_d + \phi_g$	$\phi_d(t) + \phi_g(t)$	$\phi_d(t) + \phi_g(t)$	$\phi_d(t)$	$\phi_g(t)$	0	$\chi(\lambda(t))$

is equivalent to CD driving up to a dynamical global phase (see Table II). This can also be seen by considering the differentiation with respect to  $\lambda$  of the condition

$$\langle m(\lambda) | H(\lambda) | n(\lambda) \rangle = 0 \quad \text{for } m \neq n, \quad (28)$$

$$\begin{aligned} \partial_\lambda \langle m(\lambda) | H(\lambda) | n(\lambda) \rangle &= (\partial_\lambda \langle m(\lambda) |) H(\lambda) | n(\lambda) \rangle + \langle m(\lambda) | (\partial_\lambda H(\lambda)) | n(\lambda) \rangle + \langle m(\lambda) | H(\lambda) (\partial_\lambda | n(\lambda) \rangle) \\ &= (E_m(\lambda) - E_n(\lambda)) \langle m(\lambda) | \partial_\lambda n(\lambda) \rangle + \langle m(\lambda) | (\partial_\lambda H(\lambda)) | n(\lambda) \rangle = 0, \end{aligned} \quad (29)$$

where we have used the fact that  $\langle \partial_\lambda m(\lambda) | n(\lambda) \rangle = -\langle m(\lambda) | \partial_\lambda n(\lambda) \rangle$  which follows from differentiating the condition  $\langle n(\lambda) | m(\lambda) \rangle = 0$ . By rearranging above we can find the matrix elements for the adiabatic gauge potential as

$$\begin{aligned} \langle m(\lambda) | \mathcal{A}_\lambda | n(\lambda) \rangle &= i\hbar \langle m(\lambda) | \partial_\lambda n(\lambda) \rangle \\ &= i\hbar \frac{\langle m(\lambda) | (\partial_\lambda H(\lambda)) | n(\lambda) \rangle}{E_m(\lambda) - E_n(\lambda)}, \end{aligned} \quad (30)$$

which clearly is closely related to Eq. (13). It therefore follows that we can write the off-diagonal part of the AGP in the instantaneous eigenbasis of the Hamiltonian as exactly the counterdiabatic driving term, i.e.,

$$\begin{aligned} \mathcal{A}(t) &= \dot{\lambda} \sum_{m \neq n} |m(\lambda)\rangle \langle m(\lambda) | \mathcal{A}_\lambda | n(\lambda) \rangle \langle n(\lambda) | \\ &= i\hbar \dot{\lambda} \sum_{m \neq n} \frac{|m(\lambda)\rangle \langle m(\lambda) | (\partial_\lambda H(\lambda)) | n(\lambda) \rangle \langle n(\lambda) |}{E_m(\lambda) - E_n(\lambda)}. \end{aligned} \quad (31)$$

While Eq. (31) itself does not provide any additional information regarding the controllability of complex quantum systems, it provides the starting point for taking a variational approach to constructing the control Hamiltonian. This provides a means to enforce the locality constraints in an insightful manner as well as allowing to assess the relevance of non-local interaction terms [46].

The diagonal elements of the gauge potential,  $\langle n(\lambda) | \mathcal{A}_\lambda | n(\lambda) \rangle = i \langle n(\lambda) | \partial_\lambda n(\lambda) \rangle$ , are known as the Berry connection. Importantly, they are not physical, i.e., they cannot be observed in experiments. To see this, consider the gauge transformation  $|n(\lambda)\rangle \rightarrow \exp(i\chi^{(n)}(\lambda)) |n(\lambda)\rangle$  which depends on the parameter  $\lambda$  via the phase  $\chi^{(n)}(\lambda)$ : since it attaches an overall global phase  $\chi^{(n)}(\lambda)$ , it cannot change the values of physical observables; on the other hand, it leads to a shift in the Berry connection:

$$i \langle n(\lambda) | \partial_\lambda n(\lambda) \rangle \rightarrow i \langle n(\lambda) | \partial_\lambda n(\lambda) \rangle - \partial_\lambda \chi^{(n)}(\lambda). \quad (32)$$

A physically observable quantity is the Berry phase: for an adiabatic trajectory along a *closed* curve  $\mathcal{C}$ , the Berry phase

where  $|n(\lambda)\rangle$  and  $|m(\lambda)\rangle$  are two eigenstates, i.e., the above is nothing more than the orthogonality condition. The differentiation of this orthogonality condition gives us

can be written as  $\gamma = \phi_g(T) = \oint_{\mathcal{C}} i \langle n(\lambda) | \partial_\lambda n(\lambda) \rangle d\lambda$ , which is gauge-invariant.

#### D. Counterdiabatic driving and eigenstate phase accumulation

Counterdiabatic time evolution generated by  $H_{\text{CD}}$  [Eq. (21)] or  $\mathcal{A}_\lambda$  [Eq. (26)] and starting from an eigenstate of the Hamiltonian leaves the population in the adiabatically connected instantaneous eigenstate. However, as we mentioned above, such evolution often results in a global phase; we then find the system in the instantaneous eigenstate at time  $t$ , *up to a phase*. In specific cases, we can decompose this global phase into a dynamical phase  $\phi_d$  and a geometric phase  $\phi_g$ . Table II summarizes this behavior and helps us distinguish between the zoo of different Hamiltonians  $\mathcal{H}$  we introduced above, which all solve the Schrödinger initial value problem

$$i\partial_t |\psi(t)\rangle = \mathcal{H}(t) |\psi(t)\rangle, \quad |\psi(0)\rangle = |n(0)\rangle. \quad (33)$$

In the following, it will be useful to define the co-moving reference frame via the unitary

$$\begin{aligned} V(t_2, t_1) &= \mathcal{W}(t_2, t_1) S_{\lambda(t_1)}, \\ \mathcal{W}(t_2, t_1) &= \mathcal{T}_\lambda \exp \left( -i \int_{\lambda(t_1)}^{\lambda(t_2)} d\lambda \mathcal{A}_\lambda \right) \end{aligned} \quad (34)$$

where  $\mathcal{T}_\lambda$  is the path-ordered exponential, and  $S_\lambda$  diagonalizes the Hamiltonian  $\mathcal{H}(\lambda)$ , i.e.,  $S_\lambda^\dagger \mathcal{H}(\lambda) S_\lambda = \sum_n \mathcal{E}_n(\lambda) |n_0\rangle \langle n_0| = \tilde{\mathcal{D}}(\lambda)$ . The operator  $\mathcal{W}(t_2, t_1)$  is known as the Wilson line operator, and it evolves the eigenstates at  $\lambda(t_1)$  to the eigenstates at  $\lambda(t_2)$ . If we define  $V^\dagger(t, 0)$  to transform from the lab frame to the co-moving frame, we have

$$\begin{aligned} \tilde{\mathcal{H}}(t) &= V^\dagger(t, 0) \mathcal{H}(t) V(t, 0) - iV^\dagger(t, 0) \partial_t V(t, 0) \\ &= \tilde{\mathcal{D}}(\lambda(t)) - \dot{\lambda} \tilde{\mathcal{A}}_\lambda, \end{aligned} \quad (35)$$

where in the first term we used that the combination  $\mathcal{W}(t, 0) S_{\lambda(0)}$  diagonalizes the Hamiltonian  $\mathcal{H}(\lambda(t))$  [ $S_{\lambda(0)}$

diagonalizes the Hamiltonian at  $t = 0$ , and  $\mathcal{W}(t, 0)$  propagates the eigenstates to time  $t$ ]. In the second (Galilean) term, note that the derivative  $\partial_t$  leaves  $S_{\lambda(0)}$  intact, and the expression follows using the unitarity of  $S$  and the definition of  $\mathcal{W}$  above. Finally, recall that the evolution operator  $U(t_2, t_1)$  in the lab frame can be obtained from the evolution operator in the co-moving frame  $\tilde{U}(t_2, t_1)$ , using the relation

$$\begin{aligned} U(t_2, t_1) &= V(t_2, t_1) \tilde{U}(t_2, t_1) V^\dagger(t_1, 0) \\ &= \mathcal{W}(t_2, t_1) S_{\lambda(t_1)} \tilde{U}(t_2, t_1) S_{\lambda(t_1)}^\dagger \mathcal{W}^\dagger(t_1, 0). \end{aligned} \quad (36)$$

Let us now consider six different cases for  $\mathcal{H}(t)$ , and discuss the phases accumulated by the eigenstates:

- (1) As we have shown in Sec. II B above, when  $\mathcal{H}(t) = H(t)$  and we consider evolution in the perfect adiabatic limit ( $T \rightarrow \infty, \dot{\lambda} \rightarrow 0, \dot{\lambda}T \rightarrow \text{const.}$ , and we are slower compared to the inverse energy gap to adjacent eigenstates), at time  $t$  we find the system in the state  $|\psi(t)\rangle = \exp[-i(\phi_d(t) + \phi_g(t))] |n(t)\rangle$ . This is nothing else but the content of the adiabatic theorem we proved above;
- (2) If we evolve using the counterdiabatic Hamiltonian  $\mathcal{H}(t) = H(t) + \dot{\lambda} \mathcal{A}_\lambda = H_{\text{CD}}(t)$ , the evolved state accumulates both a dynamical and a geometric phase. The difference to (1) is that we have relaxed the condition of working in the adiabatic limit. This statement is what counterdiabatic driving is designed for. This is easiest to see in the co-moving frame, where

$$\begin{aligned} \tilde{\mathcal{H}}(t) &= \tilde{H}(t) + \dot{\lambda} \tilde{\mathcal{A}}_\lambda(t) - V^\dagger(t, 0) i \partial_\lambda V(t, 0) \\ &= \tilde{H}(t) + \dot{\lambda} \tilde{\mathcal{A}}_\lambda(t) - \dot{\lambda} \tilde{\mathcal{A}}_\lambda(t) \\ &= \sum_n E_n(t) |n_0\rangle \langle n_0|, \end{aligned} \quad (37)$$

see Eq. (27), which is diagonal since  $\tilde{\mathcal{A}}_\lambda$  is canceled by the Galilean term  $V^\dagger(t, 0) i \partial_\lambda V(t, 0)$ . Time evolution in the co-moving frame then amounts to integrating the diagonal elements, and recovers precisely the dynamical phase  $\phi_{d,n}(t) = \hbar^{-1} \int_0^t E_n(s) ds$ :

$$\tilde{U}(t, 0) = \exp\left(-i \sum_n \phi_{d,n}(t) |n_0\rangle \langle n_0|\right). \quad (38)$$

Going back to the lab frame, cf. Eq. (36), we have

$$\begin{aligned} U(t, 0) &= \mathcal{W}(t, 0) S_{\lambda(0)} \exp\left(-i \sum_n \phi_{d,n}(t) |n_0\rangle \langle n_0|\right) \\ &\quad \times S_{\lambda(0)}^\dagger \mathcal{W}(0, 0). \end{aligned} \quad (39)$$

We can simplify this expression by noticing that (a)  $\mathcal{W}(0, 0) = 1$ , and (b) recalling that  $S_{\lambda(0)} |n_0\rangle = |n[\lambda(0)]\rangle$  by definition. This gives

$$U(t, 0) = \mathcal{W}(t, 0) \sum_n e^{-i\phi_{d,n}(t,0)} |n[\lambda(0)]\rangle \langle n[\lambda(0)]|. \quad (40)$$

If we now consider a closed loop,  $\lambda(0) = \lambda(T)$ , the Wilson line operator contains the Berry phases [25,47]:

$$\begin{aligned} \mathcal{W}(T, 0) &= \mathcal{T}_\lambda \exp\left(-i \oint d\lambda \mathcal{A}_\lambda\right) \\ &= \sum_n e^{-i\phi_{g,n}} |n[\lambda(0)]\rangle \langle n[\lambda(0)]|, \end{aligned} \quad (41)$$

where the geometric phase is  $\phi_{g,n} = -i\hbar \int_0^T \langle n(t) | \partial_t n(t) \rangle \dot{\lambda}(t) dt$ . Equation (40) then coincides with Eq. (19) we derived before, since for a closed loop  $|n(0)\rangle = |n(T)\rangle$ . Thus, we find that the total accumulated phase by each eigenstate is  $\phi_d(t) + \phi_g(t)$ —precisely as in the adiabatic limit, but without the restriction of working in the adiabatic regime. This is counterdiabatic driving.

- (3) If instead we evolve using another counterdiabatic Hamiltonian  $\mathcal{H}(t) = H(t) + \dot{\lambda}(i\partial_\lambda U(\lambda)) U^\dagger(\lambda) = H_{\text{CD}}(t)$ , the evolved state accumulates only a dynamical phase, i.e.,  $|\psi(t)\rangle = \exp(-i\phi_d(t)) |n(t)\rangle$ , because the geometric phase cancels out. To see this, consider the co-moving frame, where

$$\begin{aligned} \tilde{\mathcal{H}} &= \tilde{H}(t) + \dot{\lambda} U^\dagger(\lambda) (i\partial_\lambda U(\lambda)) - \dot{\lambda} \tilde{\mathcal{A}}_\lambda \\ &= \sum_n (E_n(t) + i \langle n(s) | \partial_s n(s) \rangle \dot{\lambda}(s)) |n_0\rangle \langle n_0|, \end{aligned} \quad (42)$$

which is again diagonal since the off-diagonal parts in  $U^\dagger(\lambda) (i\partial_\lambda U(\lambda))$  and  $\tilde{\mathcal{A}}_\lambda$  cancel leaving only the diagonal piece  $i \langle n(s) | \partial_s n(s) \rangle \dot{\lambda}(s)$ . Time evolution in the co-moving frame then amounts to integrating the eigenenergies of  $\tilde{\mathcal{H}}$ , which is precisely the difference between the dynamical and the geometric phases,  $\phi_d(t) - \phi_g(t)$ . Going back to the lab frame and considering a closed loop as in example (ii) above, the Wilson line operator adds an extra geometric phase  $+\phi_g$ . Summing all contributions, we find that the geometric phases cancel out, and we are left with the dynamical phase only.

- (4) If, on the other hand, we use only the gauge potential  $\mathcal{H}(t) = \dot{\lambda} \mathcal{A}(t)$  to evolve the eigenstates (also known as the Kato gauge potential [25]), the system does not accumulate a dynamical phase, but

it does accumulate a geometric phase, so the state at time  $t$  is  $|\psi(t)\rangle = \exp(-i\phi_g(t))|n(t)\rangle$ . To see the absence of a dynamical phase, note that in the co-moving frame, the Hamiltonian  $\tilde{\mathcal{H}}(t) = \dot{\lambda}\tilde{\mathcal{A}}_\lambda - \dot{\lambda}\tilde{\mathcal{A}}_\lambda \equiv 0$  vanishes identically. The gauge choice that corresponds to the Kato potential is the so-called parallel-transport gauge. It is a proper gauge along the trajectory  $\lambda(t)$ .

- (5) Whenever we evolve with the gauge potential  $\mathcal{H}(t) = \dot{\lambda}(i\partial_\lambda U(\lambda))U^\dagger(\lambda)$  only, we do not accumulate any global phases, and hence  $|\psi(t)\rangle = |n(t)\rangle$ . The best way to think about this is that the effective dynamical phase is equal in magnitude and opposite in sign to the geometric phase generated by  $\dot{\lambda}(i\partial_\lambda U(\lambda))U^\dagger(\lambda)$ , so their sum vanishes. This is seen most easily in the co-moving frame from above, where the Hamiltonian  $\tilde{\mathcal{H}}(t) = \tilde{\mathcal{H}}(t) = -\sum_n i\langle n(t)|\partial_t n(t)\rangle|n_0\rangle\langle n_0|$  produces a dynamical phase which is the negative of the geometric phase from the Wilson line operator in the lab frame. Note that, in this case, the gauge choice  $\chi(\lambda)$  is fixed such that the Berry connection is given by  $i\langle n(\lambda)|\partial_\lambda n(\lambda)\rangle$ .
- (6) And finally, if we consider evolution with the adiabatic gauge potential in a generic gauge  $\chi^{(m)}(\lambda)$ , we will find the state  $|\psi(t)\rangle = \exp(-i\chi(t))|n(t)\rangle$ , where  $\chi(t)$  is neither the dynamical nor the geometric phase.

### E. Variational adiabatic gauge potential

We will now outline a variational approach which allows us to calculate approximate adiabatic gauge potentials for complex many-body problems, even when the eigenstates are not known. This approach starts from noting that the adiabatic gauge potential defined by Eq. (31) can be alternatively given by the condition [44,45,48]

$$[A_\lambda, H_0] = i\hbar(\partial_\lambda H_0(\lambda) + F_{\text{ad}}) \quad (43)$$

with the generalized adiabatic force operator

$$F_{\text{ad}} = -\sum_n \partial_\lambda E_n(\lambda) |n(\lambda)\rangle\langle n(\lambda)|. \quad (44)$$

Note, that  $F_{\text{ad}}$  is equivalent to  $\partial_\lambda H(\lambda)$  in the energy eigenbasis. We can therefore define an operator

$$G_\lambda(A_\lambda) = \partial_\lambda H_0(\lambda) + \frac{i}{\hbar}[A_\lambda, H_0], \quad (45)$$

where now  $A_\lambda$  is an *approximation* of the adiabatic gauge potential; clearly if  $A_\lambda = \mathcal{A}_\lambda$  then  $G_\lambda(\mathcal{A}_\lambda) = -F_{\text{ad}}$ . This allows us to take a variational approach to finding the adiabatic gauge potential by minimizing the distance between

$G_\lambda(A_\lambda)$  and  $-F_{\text{ad}}$ . Choosing the Frobenius norm this distance is given by

$$\begin{aligned} D^2(A_\lambda) &= \text{Tr}[(G_\lambda(A_\lambda) + F_{\text{ad}})^2] \\ &= \text{Tr}[G_\lambda^2(A_\lambda)] + \text{Tr}[F_{\text{ad}}^2] + 2\text{Tr}[G_\lambda(A_\lambda)F_{\text{ad}}]. \end{aligned} \quad (46)$$

We can use the cyclic and linear mapping properties of the trace in the eigenbasis to simplify the last term

$$\begin{aligned} \text{Tr}[G_\lambda(A_\lambda)F_{\text{ad}}] &= \text{Tr}[F_{\text{ad}}\partial_\lambda H_0(\lambda)] + \frac{i}{\hbar}\text{Tr}[F_{\text{ad}}[A_\lambda, H_0]] \\ &= -\text{Tr}[F_{\text{ad}}^2] - \frac{i}{\hbar}\text{Tr}[[F_{\text{ad}}, A_\lambda]H_0] \\ &= -\text{Tr}[F_{\text{ad}}^2], \end{aligned} \quad (47)$$

where we have used the fact that  $F_{\text{ad}}$  will commute with the Hamiltonian. Therefore, in order to construct an approximation to the AGP we must minimize the distance given by

$$D^2(A_\lambda) = \text{Tr}[G_\lambda^2(A_\lambda)] - \text{Tr}[F_{\text{ad}}^2]. \quad (48)$$

Since the adiabatic force does not explicitly depend on  $A_\lambda$ , the minimization procedure is equivalent to simply minimizing the norm of  $G_\lambda(A_\lambda)$ . We can simplify the procedure by minimizing the associated action, i.e., define

$$S(A_\lambda) = \text{Tr}[G_\lambda^2(A_\lambda)] \quad \text{and minimize} \quad \frac{\partial S(A_\lambda)}{\partial A_\lambda} = 0. \quad (49)$$

This will give an approximate adiabatic gauge potential, or equivalently, an approximate counterdiabatic term. The power of this approach stems from the freedom in choosing what operators appear in  $A_\lambda$ . By restricting to, for example, only local operators we can construct a control term that will perform as well as possible for a many-body system under such a constraint and we can readily expand the operator set to more complex terms, introducing higher order interactions. As a corollary, this approach therefore also provides valuable insight into the relevance of the non-local interaction terms that the exact CD term typically requires [46,49].

Using the approach of variational CD for an arbitrary many-body quantum system, we may have simply moved the complexity from knowing the instantaneous eigenstates to finding a suitable ansatz for  $A_\lambda$ . Note, in general, there will be exponentially many operators on which the adiabatic gauge potential could potentially have support, and this could still be unwieldy even when truncating the number of bodies involved or the range of the operators. However, a physically motivated option does exist [50],

first, we can express the adiabatic gauge potential as

$$\mathcal{A}_\lambda = \lim_{\epsilon \rightarrow 0^+} \int_0^\infty dt e^{-\epsilon t} (e^{-iH(\lambda)t} \partial_\lambda H(\lambda) e^{iH(\lambda)t} - \mathcal{M}_\lambda), \quad (50)$$

with  $\mathcal{M}_\lambda$  representing the diagonal terms, which are not relevant. The first term of the integral can be expanded via the Baker-Campbell-Hausdorff formula, taking a representative exponent  $X$  and operator  $Y$ ,

$$e^X Y e^{-X} = Y + [X, Y] + \frac{1}{2}[X, [X, Y]] + \dots + \frac{1}{n!}[X, \underbrace{[X, \dots [X, Y] \dots]}_n] + \dots, \quad (51)$$

with the commutator  $[X, \cdot]$  repeated  $n$  times in the general term. Using this we obtain

$$\mathcal{A}_\lambda = i \sum_{n=1}^{\infty} \alpha_n \underbrace{[H(\lambda), [H(\lambda), \dots [H(\lambda), \partial_\lambda H(\lambda)]]]}_n, \quad (52)$$

where the  $\alpha_n$  coefficients still need to be solved for using the variational approach. While not immediately obvious, we find that  $\mathcal{A}_\lambda$  is off-diagonal and that only odd terms in  $n$  actually contribute to the off-diagonal [50]. Therefore, we can trivially set all  $\alpha_n = 0$  for all  $n = 2m$  with  $m \in \mathbb{Z}$  and write the adiabatic gauge potential as

$$\mathcal{A}_\lambda = i \sum_{l=1}^{\infty} \alpha_l \underbrace{[H(\lambda), [H(\lambda), \dots [H(\lambda), \partial_\lambda H(\lambda)]]]}_{2l-1}. \quad (53)$$

Now, to take an informed ansatz we can consider a truncation of the commutation relation. Note, that while this commutator ansatz will restrict to those operators that  $\mathcal{A}_\lambda$  can physically have support on, there still may exist exponentially many terms for an arbitrary many-body Hamiltonian [51]. It should also be noted that a finite truncation of  $l = 1, \dots, p$  does not mean that you have all operators

for up to  $p$ -body interactions or that it does not contain operators of more than  $p$  bodies. In addition, each order of the expansion of the commutator can contain repeating terms, and they are not orthogonal to each other, though often in practice one orthogonalizes each term to avoid having an overly expressive ansatz [51–53]. The use of the commutator ansatz at finite truncation and its connection to Krylov space is an active area of study which we will briefly discuss in Sec. VI.

## F. Examples

With the necessary groundwork established we now move to the main purpose of this tutorial: applying these techniques to several systems including those with direct relevance for many-body systems. We take a deliberately pedagogical approach providing explicit calculations of the various control methods. In doing so we therefore also detail some well-known techniques for solving paradigmatic many-body models, in particular, Jordan-Wigner, Holstein-Primakoff, and Bogoliubov transformations.

### 1. Example: The Landau-Zener model

We begin by considering the two-level Landau-Zener (LZ) model with Hamiltonian given by

$$H \equiv H(t) = \Delta \sigma^x + v(t) \sigma^z, \quad (54)$$

where  $\sigma^a$ ,  $a \in x, y, z$  are the Pauli matrices

$$\sigma^x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \quad \sigma^y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \sigma^z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad (55)$$

and we assume units such that  $\hbar = 1$ . Its simplicity notwithstanding, the LZ model captures a remarkably diverse range of physical phenomena [21]. Relevant to many-body physics, it captures the central aspects of critical systems and in particular two-band models [33]. The corresponding eigenstates and eigenenergies can be written

$$\begin{aligned} |\phi_g\rangle &= \cos(\theta/2) |0\rangle + \sin(\theta/2) |1\rangle, & E_g &= -\sqrt{\Delta^2 + v(t)^2} \\ |\phi_e\rangle &= \sin(\theta/2) |0\rangle - \cos(\theta/2) |1\rangle, & E_e &= \sqrt{\Delta^2 + v(t)^2} \end{aligned} \quad (56)$$

with  $\tan \theta = \Delta/v(t)$ . Employing Eq. (14) we see that the counterdiabatic Hamiltonian is given by

$$\begin{aligned} \mathcal{A} &= i \left( \frac{|\phi_g\rangle\langle\phi_g| \partial_t v(t) \sigma^z |\phi_e\rangle\langle\phi_e|}{E_g - E_e} + \frac{|\phi_e\rangle\langle\phi_e| \partial_t v(t) \sigma^z |\phi_g\rangle\langle\phi_g|}{E_e - E_g} \right) \\ &= i \frac{\partial_t v(t)}{2\sqrt{\Delta^2 + v(t)^2}} \left( |\phi_e\rangle\langle\phi_e| \sigma^z |\phi_g\rangle\langle\phi_g| - |\phi_g\rangle\langle\phi_g| \sigma^z |\phi_e\rangle\langle\phi_e| \right) = -\frac{\partial_t v(t) \sin \theta}{2\sqrt{\Delta^2 + v(t)^2}} \sigma^y = -\frac{\Delta \partial_t v(t)}{2(\Delta^2 + v(t)^2)} \sigma^y. \end{aligned} \quad (57)$$

We can arrive at the same expression for  $\mathcal{A}$  using Eq. (21)

$$\mathcal{A} = i \left( |\dot{\phi}_g\rangle\langle\phi_g| - \langle\phi_g|\dot{\phi}_g\rangle |\phi_g\rangle\langle\phi_g| + |\dot{\phi}_e\rangle\langle\phi_e| - \langle\phi_e|\dot{\phi}_e\rangle |\phi_e\rangle\langle\phi_e| \right) = i(i\dot{\theta}\sigma^y) = -\frac{\Delta \partial_t \nu(t)}{2(\Delta^2 + \nu(t)^2)} \sigma^y. \quad (58)$$

For any choice of ramp,  $\nu(t)$ , evolving the system according to the Hamiltonian

$$H_{\text{CD}} = H + \mathcal{A} = \Delta\sigma^x + \nu(t)\sigma^z - \frac{\Delta \partial_t \nu(t)}{2(\Delta^2 + \nu(t)^2)} \sigma^y, \quad (59)$$

will ensure a perfectly adiabatic dynamics. As an example, if the system is initialized in the ground state for some initial value of the field,  $\nu_0$ , and we linearly ramp the system according to  $\nu(t) = \nu_0 + \nu_d(t/T)$ , we can understand why this Hamiltonian achieves a perfectly adiabatic evolution by comparing the energy spectra of the original and the controlled Hamiltonians [54]. In Fig. 2, the solid curves show the behavior of the eigenenergies of the bare LZ model for a linear ramp and we clearly see the avoided crossing occurring for  $t/T = 0.5$ . In contrast, the energy spectrum of the control Hamiltonian shows that adding the counterdiabatic term modifies the energy gap in such a way that, rather than exhibiting an avoided crossing, the energy levels are pushed apart in precisely the right manner to ensure that the rate at which the system is being driven is still slow enough, with respect to this new energy gap, that it continues to evolve adiabatically. Clearly the counterdiabatic term fundamentally changes the spectral properties of the system in order to achieve adiabaticity. Indeed, using the fact that  $\mathcal{A}$  is a Hamiltonian term, several




FIG. 2. Time-dependence of the energy eigenvalues for the Landau-Zener model (solid, black) and for the controlled evolution, Eq. (59) (dashed) for two different protocol durations assuming a linear ramp  $\nu(t) = -5 + 10(t/T)$  and  $\Delta = 1$ .

works use this as the basis for examining the thermodynamics of quantum control arguing that  $\mathcal{A}$  can be thought of as an energetic toll that must be paid in order to manipulate the system [55–61]. For simple models such as Eq. (54) it is intuitive to expect that  $\mathcal{A}$  would modify the spectrum by simply opening an energy gap; however, this is not necessarily true, as has been shown by closing energy gaps in examples including non-interacting tight-binding models [62] and many-body spin-1/2 lattice gauge models [63]. In fact, the study of what constitutes the important characteristics that the counterdiabatic term induces in the spectrum of a complex many-body system is an open question. While interesting and insightful in its own right, understanding the control requirements for the LZ model is vital when moving to the genuine many-body Ising model.

## 2. Example: The Ising model in one dimension

The transverse-field nearest-neighbor Ising model describes a collection of interacting spin- $\frac{1}{2}$  degrees of freedom on a lattice under the action of a perpendicular magnetic field. This model is often taken as an example to first consider new control protocols due to its well-known solutions and relation to other higher-complexity cases that are studied [32,64,65]. We will consider the one-dimensional (1D) Ising model for simplicity, which is given by

$$H = -\sum_{j=1}^N \left( \sigma_j^x \sigma_{j+1}^x + g \sigma_j^z \right), \quad (60)$$

$\sigma_j^a$ ,  $a \in x, y, z$  representing the Pauli matrix acting on the  $j$ th site and  $g$  the relative strength of an applied external field. This model exhibits a quantum phase transition at  $g = g_c = \pm 1$  between a paramagnetic ( $|g| > 1$ ) and ferromagnetic ( $|g| < 1$ ) phase. As it is the only parameter, we will consider the scenario where  $g = g(t)$ , i.e., we drive  $g$  and are interested in implementing CD terms such that we negate any diabatic transitions so that, e.g., we can remain in the ground state of the system in a finite driving time.

*The Jordan-Wigner transformation.* It is not always possible to derive the exact CD for arbitrary Hamiltonians, but luckily the Ising model is exactly solvable through the Jordan-Wigner transformation. This allows for the exact CD to be obtained, as was done in Refs. [42,66,67]. First, we will assume periodic boundary conditions, i.e.,  $\sigma_{N+1} = \sigma_1$ , for an  $N$  site system with an even number of spins and  $N \geq 2$ . It is simpler to consider the term across the periodic

boundary separately and to write the Hamiltonian as

$$H = - \sum_{j=1}^{N-1} \sigma_j^x \sigma_{j+1}^x - \sigma_N^x \sigma_1^x - g \sum_{j=1}^N \sigma_j^z. \quad (61)$$

The Jordan-Wigner transformation is used to transform the Hamiltonian from the interacting spin basis to a basis of non-interacting fermions, and a useful summary of its implementation can be found in Refs. [66,68]. We can define a set of creation and annihilation fermionic operators from the spin basis as

$$c_j^\dagger = K_j^\dagger \sigma_j^-, \quad c_j = K_j \sigma_j^+, \quad (62)$$

with

$$\sigma_j^\pm = \frac{1}{2} \left( \sigma_j^x \pm i \sigma_j^y \right) \quad K_j = e^{i\pi \sum_{k<j} \sigma_k^+ \sigma_k^-} = \prod_{k<j} (1 - 2n_k), \quad (63)$$

where  $n_j = c_j^\dagger c_j$  is the number operator and  $K_j$ , with eigenvalues  $\pm 1$ , is the non-local string operator which counts the parity of the number of fermions on the lattice before site  $j$ . The fermionic operators obey the anti-commutation relations

$$\{c_i, c_j^\dagger\} = \delta_{ij} \quad \{c_i, c_j\} = \{c_i^\dagger, c_j^\dagger\} = 0. \quad (64)$$

The one-body term is then relatively straightforward as we can write the  $z$ -Pauli matrix on site  $j$  as

$$\sigma_j^z = 1 - 2c_j^\dagger c_j, \quad (65)$$

which can be seen from

$$\begin{aligned} \sigma_j^z &= 1 - 2\sigma_j^- \sigma_j^+ = 1 - \frac{1}{2} \left( \sigma_j^x - i \sigma_j^y \right) \left( \sigma_j^x + i \sigma_j^y \right) = 1 - \frac{1}{2} \left( \sigma_j^x \sigma_j^x - i^2 \sigma_j^y \sigma_j^y - i \sigma_j^y \sigma_j^x + i \sigma_j^x \sigma_j^y \right) \\ &= 1 - \frac{1}{2} \left( 1 + 1 - i[\sigma_j^y, \sigma_j^x] \right) = 1 - \left( 1 - \sigma_j^z \right) = \sigma_j^z, \end{aligned} \quad (66)$$

note that the ordering of the Pauli matrices is important as they do not commute and we have used the fact that  $[\sigma_j^y, \sigma_j^x] = -2i\sigma_j^z$ . We can also write out the  $x$ -Pauli matrix as,

$$\sigma_j^x = K_j \left( c_j^\dagger + c_j \right) = K_j \left( K_j^\dagger \sigma_j^- + K_j \sigma_j^+ \right) = \sigma_j^- + K_j K_j \sigma_j^+ = \sigma_j^- + \sigma_j^+ = \sigma_j^x, \quad (67)$$

where we used the fact that  $K_j K_j$  acting on any state is equivalent to  $\mathbb{1}_j$ , i.e., the identity operator. With this we can write the Ising Hamiltonian given by Eq. (61) as

$$\begin{aligned} H &= - \sum_{j=1}^{N-1} K_j \left( c_j^\dagger + c_j \right) K_{j+1} \left( c_{j+1}^\dagger + c_{j+1} \right) - K_N \left( c_N^\dagger + c_N \right) \left( c_1^\dagger + c_1 \right) - g \sum_{j=1}^N \left( 1 - 2c_j^\dagger c_j \right) \\ &= - \sum_{j=1}^{N-1} K_j K_j \left( c_j^\dagger + c_j \right) \left( 1 - 2n_j \right) \left( c_{j+1}^\dagger + c_{j+1} \right) - K_N \left( c_N^\dagger + c_N \right) \left( c_1^\dagger + c_1 \right) - g \sum_{j=1}^N \left( 1 - 2c_j^\dagger c_j \right). \end{aligned} \quad (68)$$

We can simplify this further by considering each term in turn. The first term can be written as

$$\begin{aligned} & - \sum_{j=1}^{N-1} K_j K_j \left( c_j^\dagger + c_j \right) \left( 1 - 2n_j \right) \left( c_{j+1}^\dagger + c_{j+1} \right) = - \sum_{j=1}^{N-1} \left( c_j^\dagger + c_j \right) \left( 1 - 2c_j^\dagger c_j \right) \left( c_{j+1}^\dagger + c_{j+1} \right) \\ &= - \sum_{j=1}^{N-1} \left( c_j^\dagger + c_j - 2c_j^\dagger c_j^\dagger c_j - 2c_j c_j^\dagger c_j \right) \left( c_{j+1}^\dagger + c_{j+1} \right) = - \sum_{j=1}^{N-1} \left( c_j^\dagger + c_j - 2c_j \right) \left( c_{j+1}^\dagger + c_{j+1} \right) \\ &= \sum_{j=1}^{N-1} \left( c_j - c_j^\dagger \right) \left( c_{j+1}^\dagger + c_{j+1} \right), \end{aligned} \quad (69)$$

the second term—as

$$\begin{aligned}
& -K_N (c_N^\dagger + c_N) (c_1^\dagger + c_1) = -K_{N+1} (1 - 2n_N) (c_N^\dagger + c_N) (c_1^\dagger + c_1) \\
& = -P (c_N^\dagger + c_N - 2c_N^\dagger c_N c_N^\dagger - 2c_1^\dagger c_1 c_1) (c_1^\dagger + c_1) = -P (c_N^\dagger + c_N - 2c_N^\dagger) (c_1^\dagger + c_1) \\
& = -P (c_N - c_N^\dagger) (c_1^\dagger + c_1),
\end{aligned} \tag{70}$$

where  $P = K_{N+1}$  is the parity operator for the number of fermions in the full  $N$ -site system. The third term cannot be simplified further, though it is common to see it written in a different form given by

$$\sum_{j=1}^N (1 - 2c_j^\dagger c_j) = \sum_{j=1}^N (c_j^\dagger c_j + c_j c_j^\dagger - 2c_j^\dagger c_j) = \sum_{j=1}^N (c_j c_j^\dagger - c_j^\dagger c_j), \tag{71}$$

where we have used the anti-commutation relations given in Eq. (64). All this together means we can write Eq. (68) as

$$H = \sum_{j=1}^{N-1} (c_j - c_j^\dagger) (c_{j+1}^\dagger + c_{j+1}) - P (c_N - c_N^\dagger) (c_1^\dagger + c_1) - g \sum_{j=1}^N (c_j c_j^\dagger - c_j^\dagger c_j). \tag{72}$$

This Hamiltonian is invariant under the parity operator and we can therefore split it into the even parity ( $P = 1$ ) and odd parity ( $P = -1$ ) subspaces, and write the full Hamiltonian as

$$H = H^+ P^+ + H^- P^-, \tag{73}$$

with

$$P^\pm = \frac{1}{2} (1 \pm P), \quad \text{and} \quad H^\pm = \sum_{j=1}^N (c_j - c_j^\dagger) (c_{j+1}^\dagger + c_{j+1}) - g \sum_{j=1}^N (c_j c_j^\dagger - c_j^\dagger c_j), \tag{74}$$

where  $H^+$  has anti-periodic boundary conditions, i.e.,  $c_{N+1} \equiv -c_1$ , and  $H^-$  has periodic boundary conditions, i.e.,  $c_{N+1} \equiv c_1$ .

*Momentum space.* We will now focus our attention on the positive parity subspace, with the consideration of the negative parity subspace being similar to the discussion below. The dynamics of the Ising spin chain given by Hamiltonian Eq. (60) is then described by the Schrödinger equation and the fermionic Hamiltonian

$$H^+ = \sum_{j=1}^N (c_j - c_j^\dagger) (c_{j+1}^\dagger + c_{j+1}) - g \sum_{j=1}^N (c_j c_j^\dagger - c_j^\dagger c_j), \tag{75}$$

with anti-periodic boundary conditions,  $c_{N+1} = -c_1$ .

It will now be convenient to write the model in (quasi)momentum  $k$ -space and to do this we can first write the fermion operators in  $k$ -space as

$$c_k = \frac{e^{-i\phi}}{\sqrt{N}} \sum_{j=1}^N e^{-ikj} c_j, \quad c_j = \frac{e^{i\phi}}{\sqrt{N}} \sum_k e^{ikj} c_k, \tag{76}$$

with the overall phase given by  $\phi$  changing the phase of the pairs that will be created in  $k$ -space as we will see below. Note, the choice of parity subspace informs the correct values of the quasimomentum  $k$ . For anti-periodic boundary conditions we must satisfy  $c_{N+1} \equiv -c_1$ , which after transforming to  $k$ -space will mean  $e^{ikN} = -1$ , which has solutions

$$k = \pm \frac{(2m-1)\pi}{N} \quad \text{with } m = 1, 2, \dots, \frac{N}{2}. \tag{77}$$

First we can Fourier transform the local term

$$\begin{aligned} \sum_{j=1}^N (c_j c_j^\dagger - c_j^\dagger c_j) &= \frac{1}{N} \sum_j \sum_{k,k'} (e^{-i(k'-k)j} c_k c_{k'}^\dagger - e^{-i(k'-k)j} c_{k'}^\dagger c_k) \\ &= \sum_{k,k'} (\delta_{k,k'} c_k c_{k'}^\dagger - \delta_{k,k'} c_{k'}^\dagger c_k) = \sum_k (c_k c_k^\dagger - c_k^\dagger c_k), \end{aligned} \quad (78)$$

where we have used the fact that  $\delta_{m,n} = \sum_j e^{-i(n-m)j} / N$ . We can follow similar steps with the nearest-neighbor term to obtain

$$\begin{aligned} \sum_{j=1}^N (c_j - c_j^\dagger) (c_{j+1}^\dagger + c_{j+1}) &= \sum_{k,k'} (e^{2i\phi} e^{ik'} \delta_{-k,k'} c_k c_{k'} - e^{-2i\phi} e^{-ik'} \delta_{-k,k'} c_k^\dagger c_{k'}^\dagger - \delta_{k,k'} c_k^\dagger c_k + \delta_{k,k'} c_k c_k^\dagger) \\ &= \sum_k (2 \cos k c_k^\dagger c_k + e^{2i\phi} e^{ik} c_k c_{-k} - e^{-2i\phi} e^{-ik} c_k^\dagger c_{-k}^\dagger), \end{aligned} \quad (79)$$

where in the final step we have neglected a constant energy offset. Combining these, and using the anti-commutation relations as well as the fact that  $\sum_k \cos k = 0$ , we can obtain the  $k$ -space Hamiltonian to be

$$H = \sum_k [(\cos k - g) (c_k c_k^\dagger - c_k^\dagger c_k) + e^{2i\phi} e^{ik} c_k c_{-k} - e^{-2i\phi} e^{-ik} c_k^\dagger c_{-k}^\dagger]. \quad (80)$$

All terms in the Hamiltonian now come in momentum pairs,  $(k, -k)$ , and we can define the full Hamiltonian by only considering the positive  $k$  values. Let us now define a fermionic two-component spinor

$$\Psi_k = \begin{pmatrix} c_k \\ c_{-k}^\dagger \end{pmatrix} \quad \Psi_k^\dagger = \begin{pmatrix} c_k^\dagger c_{-k} \end{pmatrix}, \quad (81)$$

which we can use to rewrite the Hamiltonian as

$$\begin{aligned} H &= \sum_{k>0} H_k \quad \text{with} \\ H_k &= 2\Psi_k^\dagger \begin{pmatrix} g - \cos k & -ie^{-2i\phi} \sin k \\ ie^{2i\phi} \sin k & -g + \cos k \end{pmatrix} \Psi_k. \end{aligned} \quad (82)$$

We can write out the above  $2 \times 2$  matrix in terms of a psuedo-spin as is shown in Eq. (54) of Ref. [68], the choice of  $\phi$  then comes down to the direction of the effective magnetic field for this psuedo-spin, for the Ising model we have written the effective magnetic field of the psuedo-spin is in the  $x - z$  plane which corresponds to  $\phi = -\pi/4$ . Taking this value of  $\phi$  we can write the Hamiltonian in momentum space as

$$H_k = 2\Psi_k^\dagger [(g - \cos k) \sigma^z + \sin k \sigma^x] \Psi_k. \quad (83)$$

*Back to Landau-Zener.* We have now managed to decompose the 1D Ising model into a collection of independent two-level systems. Note that, under Schrödinger

dynamics, each two-level system will undergo a Landau-Zener transition. We can now utilize the results of Sec. II F 1, identifying  $\Delta = 2 \sin k$  and  $v(t) = 2g - 2 \cos k$ , and the CD term can be written as

$$A = - \sum_{k>0} \frac{\dot{g} \sin k}{2(1 + g^2 - 2g \cos k)} \Psi_k^\dagger \sigma^y \Psi_k. \quad (84)$$

If one is only interested in being able to numerically calculate the outcome of a protocol with CD this is where we could stop. However, while we have the CD term in a nice form, expressing this in  $k$ -space may not be ideal for physical implementation. In addition, as the CD term is local in  $k$ -space, we should expect that it will be highly non-local in physical space; but this physical space alone is the space of the non-interacting fermions, and we must remember that the native basis for our initial Hamiltonian is in fact spin-1/2's.

Having Fourier converted the problem of spin-1/2's to non-interacting fermions and then transformed to  $k$ -space, we now find ourselves with the CD term but with the need to trek back along the path we have followed to this point. First we can write out the CD term in the fermion operators as

$$A = -i \sum_{k>0} \frac{\dot{g} \sin k}{2(1 + g^2 - 2g \cos k)} (c_k c_{-k} - c_k^\dagger c_{-k}^\dagger), \quad (85)$$

then utilizing the inverse Fourier transform defined in Eq. (76) we can write this in position space as

$$\mathcal{A} = -\frac{i}{2N}\dot{g} \sum_{j,j'}^N \sum_{k>0} \frac{\sin k}{1+g^2-2g\cos k} \sin((j-j')k) \times (c_j c_{j'} + c_j^\dagger c_{j'}^\dagger), \quad (86)$$

$$\begin{aligned} \mathcal{A} &= -\frac{i}{2}\dot{g} \sum_{j,j'}^N f(j-j') (c_j c_{j'} + c_j^\dagger c_{j'}^\dagger) \\ &= -\frac{i}{2}\dot{g} \sum_{j,j'}^N f(j-j') \left[ K_j K_{j'} \sigma_j^+ \left( \prod_{j \leq m < j'} \sigma_m^z \right) \sigma_{j'}^+ + K_j^\dagger K_{j'}^\dagger \sigma_j^- \left( \prod_{j \leq m < j'} \sigma_m^z \right) \sigma_{j'}^- \right] \\ &= -\dot{g} \sum_{j,j'}^N f(j-j') \left[ \sigma_j^y \left( \prod_{j < m < j'} \sigma_m^z \right) \sigma_{j'}^x + \sigma_j^x \left( \prod_{j < m < j'} \sigma_m^z \right) \sigma_{j'}^y \right], \end{aligned} \quad (87)$$

where in the last step we have removed a  $\sigma_j^z$  from the product and used the fact that  $\sigma_j^x \sigma_j^z = -i\sigma_j^y$  and  $\sigma_j^y \sigma_j^z = i\sigma_j^x$ .

Achieving perfect control still requires knowledge of the coefficients,  $f(j-j')$ , the precise determination of which is highly non-trivial [42,67], as we will discuss explicitly shortly. Nevertheless, Eq. (87) is the exact CD term for the 1D Ising model in the native spin-1/2 basis. As anticipated, it is a highly non-local operator which couples spins across the full system and, in this regard, it is instructive to examine how the complexity of the CD term emerges as we scale up the size of the system. For  $N=2$  we can directly compute the exact  $\mathcal{A}$  using Eq. (14) and find

$$\mathcal{A} \propto (\sigma_1^y \sigma_2^x + \sigma_1^x \sigma_2^y) \quad (88)$$

similarly for  $N=3$

$$\mathcal{A} \propto (\sigma_1^y \sigma_2^x + \sigma_1^x \sigma_2^y + \sigma_2^y \sigma_3^x + \sigma_2^x \sigma_3^y + \sigma_3^y \sigma_1^x + \sigma_3^x \sigma_1^y), \quad (89)$$

where we see the operators entering the control terms are precisely those appearing in Eq. (87). Thus, while Eq. (87) provides the overall form of the counterdiabatic term, it still remains to determine the correct form of the coefficients,  $f(j-j')$ . For  $N=2$  it is easy to find the explicit form of the coefficient by direct calculation, resulting in the CD term taking the compact form

$$\mathcal{A} = -\frac{\dot{g}}{4(1+g^2)} (\sigma_1^y \sigma_2^x + \sigma_1^x \sigma_2^y). \quad (90)$$

remembering that we only need to consider positive  $k$ . We see that, as expected, the CD term is highly non-local in physical space as it is a sum over all pairs of sites in the 1D lattice. We can now write the CD term in terms of Pauli matrices, using Eq. (62),

Already determining the precise coefficients even for  $N=3$  becomes significantly more involved; however, if we restrict to only controlling the ground state we find by direct calculation

$$\begin{aligned} A_{\text{GS}} &= -\dot{g} \frac{1+g}{8(1+g^3)} (\sigma_1^y \sigma_2^x + \sigma_1^x \sigma_2^y + \sigma_2^y \sigma_3^x + \sigma_2^x \sigma_3^y \\ &\quad + \sigma_3^y \sigma_1^x + \sigma_3^x \sigma_1^y), \end{aligned} \quad (91)$$

where the choice of symbol reflects the fact that, despite the operators appearing in  $A_{\text{GS}}$  corresponding precisely to those that make up  $\mathcal{A}$ , the choice of coefficients means that this counterdiabatic term will not suppress excitations for all eigenstates of the Hamiltonian. In general, one can take a similar approach to Ref. [42] and assume an ansatz form for the coefficients. If we restrict to control of the ground state for an even size chain with periodic boundary conditions, then we can use the results of Ref. [67], which provides an analytic form for the coefficients. For consistency with Refs. [42,67] we slightly rearrange Eq. (87) into the form

$$A_{\text{GS}} = -\dot{g} \left[ \sum_{m=1}^{N/2-1} h_m(g) A_{\text{GS}}^{[m]} + \frac{1}{2} h_{N/2}(g) A_{\text{GS}}^{[N/2]} \right] \quad (92)$$

where the coefficients are now given precisely by

$$h_m(g) = \frac{g^{2m} + g^N}{8g^{m+1}(1+g^N)}. \quad (93)$$

It should be clear that each  $m$  term in Eq. (92) corresponds to a different “range” of interactions given by

$$A_{\text{GS}}^{[m]} = \sum_{n=1}^N \left[ \sigma_n^y \left( \prod_{j=n+1}^{n+m-1} \sigma_j^z \right) \sigma_{n+m}^x + \sigma_n^x \left( \prod_{j=n+1}^{n+m-1} \sigma_j^z \right) \sigma_{n+m}^y \right], \quad (94)$$

with  $m = 1$  corresponding to nearest-neighbor terms, i.e., no  $\sigma_j^z$ 's appearing as, e.g., seen in the  $N = 2$  and 3 cases, and we extend up to a maximal range of  $N/2$  due to the assumed periodic boundary conditions. Note that the factor of  $1/2$  appearing in front of  $A_{\text{GS}}^{[N/2]}$  is due to the even system size which results in there being only a single spin at the “maximal” distance of  $N/2$  from site 1. This allows us to write out explicitly the control term that will maintain adiabatic dynamics for the ground state for  $N = 4$ , which is the smallest system size such that there are contributions arising from both terms in Eq. (92) with

$$A_{\text{GS}} = A_{\text{GS}}^{(2)} = -\dot{g} \left[ h_1(g) A_{\text{GS}}^{[1]} + \frac{1}{2} h_2(g) A_{\text{GS}}^{[2]} \right] \quad \text{where} \quad (95)$$

$$h_1(g) A_{\text{GS}}^{[1]} = \frac{1+g^2}{8(1+g^4)} (\sigma_1^y \sigma_2^x + \sigma_1^x \sigma_2^y + \sigma_2^y \sigma_3^x + \sigma_2^x \sigma_3^y + \sigma_3^y \sigma_4^x + \sigma_3^x \sigma_4^y + \sigma_4^y \sigma_1^x + \sigma_4^x \sigma_1^y) \quad (96)$$

$$h_2(g) A_{\text{GS}}^{[2]} = \frac{g}{4(1+g^4)} (\sigma_1^y \sigma_2^z \sigma_3^x + \sigma_1^x \sigma_2^z \sigma_3^y + \sigma_2^y \sigma_3^z \sigma_4^x + \sigma_2^x \sigma_3^z \sigma_4^y + \sigma_3^y \sigma_4^z \sigma_1^x + \sigma_3^x \sigma_4^z \sigma_1^y + \sigma_4^y \sigma_1^z \sigma_2^x + \sigma_4^x \sigma_1^z \sigma_2^y) \quad (97)$$

where for consistency of notation with Sec. II F 4 we introduce  $A_{\text{GS}}^{(m)}$  to denote the cumulative control term utilizing all interactions with ranges up to and including  $m$ -body terms, i.e., the sum of the individual terms with a specific range  $A_{\text{GS}}^{[m]}$ . In Fig. 3 we show the relevance of the different contributions in achieving control of the full system. We see that if we include only the two-body control terms,

i.e., evolve according to  $H - \dot{g} h_1(g) A_{\text{GS}}^{[1]}$ , these terms provide the most significant improvement in the ability control of the system when compared to the performance if we include only the three body control terms, i.e., evolve according to  $H - \dot{g} h_2(g) A_{\text{GS}}^{[2]}$ . Clearly, given that we are considering only  $N = 4$ , perfect control is achieved when both two- and three-body terms are present and the system is evolved under  $H + A_{\text{GS}}^{(2)}$ .




FIG. 3. Ising model for  $N = 4$  with  $g(t) = 0.1 + 1.9t/T$  fixing  $T = 1$  and starting in the ground state. We show the fidelity with the instantaneous ground state for the full counterdiabatic term (black) involving all two and three body interactions, i.e., Eq. (95), implementing only the two-body terms, i.e., Eq. (96) (dashed, red), implementing only the three-body terms, i.e., Eq. (97) (dot-dashed, orange). We also show the evolution for the bare Hamiltonian (blue).

### 3. Example: The Ising model with all-to-all couplings (Lipkin-Meshkov-Glick model)

While the previous section considered the solvable Ising model with nearest-neighbor interactions, for which the exact CD term for any  $N$  can be determined, we can also look at a complementary interaction setting, where all constituents mutually interact with one another. This all-to-all coupling, or infinite range interaction, is captured by the Lipkin-Meshkov-Glick (LMG) Hamiltonian

$$H = -\frac{1}{N} \sum_{j < k}^N \sigma_j^x \sigma_k^x - \sum_j^N g(t) \sigma_j^z. \quad (98)$$

Originally proposed as a model to study shape phase transitions in atomic nuclei, Eq. (98) has since come to encompass a wide range of physical phenomena including as a model of magnetism [69,70] and as a paradigmatic model for studying quantum phase transitions [71]. Notice that unless warranted for clarity, we use  $g(t) \equiv g$  for brevity.

From the quantum many-body perspective, this model is interesting due to its very rich phase diagram, host to

ground and excited state quantum phase transitions thanks to the high degree of symmetry the all-to-all interaction provides. In the thermodynamic limit,  $N \rightarrow \infty$ , the ground state phase transition is between symmetric ( $0 < g < 1$ ) and symmetry broken ( $g > 1$ ) phases. In the symmetric phase, even and odd parity subspaces are degenerate (and quasi-degenerate for large but finite  $N$  with the gap between even and odd sub-sector energies inversely scaling with system size), while in the symmetry broken phase this degeneracy is lifted [69,72]. Despite the similarities with the Ising model, while Eq. (60) can be solved for arbitrary values of  $N$  the same is not true for Eq. (98). However, due to the high degree of symmetry that this Hamiltonian exhibits, it is more amenable to a mean-field type approach, allowing for the model to be exactly solvable in the thermodynamic limit by mapping to a quantum harmonic oscillator. This raises interesting consequences when aiming to achieve control of the model: tools and techniques have been developed to control a harmonic oscillator [73]; that said, clearly for any finite  $N$  the mapping will only be an approximation and therefore leaves questions as to the efficacy of “naively” applying such control approaches. To examine this, however, we must first solve the model. In what follows we will focus on the symmetry broken phase only, i.e.,  $g > 1$ , carefully detailing the steps necessary to find the solution. An almost identical calculation can be performed for the symmetric phase, with some caveats that we will discuss later and details of which can be found in, e.g., Ref. [74].

The first step is to exploit the symmetry induced by the homogeneous all-to-all interactions. This means that it is more convenient to write Eq. (98) using collective spin operators,  $S_\alpha = \frac{1}{2} \sum_n \sigma_n^\alpha$  giving

$$H = -\frac{2}{N} S_x^2 - 2g S_z + \frac{1}{2} \mathbb{1}. \quad (99)$$

Notice that the last term, proportional to  $\mathbb{1}$ , will have no impact on the dynamics that we will study, it simply amounts to a global phase, and hence is often neglected. For pedagogical reasons of completeness we will keep this term. It is easy to confirm that  $S^2$  is a conserved quantity and we can view the  $N$ -particle system as an effective  $(N+1)$ -level spin. We can solve to find the ground state of the model by mapping to *bosonic* operators following the Holstein-Primakoff approximation, which is valid for  $N \gg 1$ . If we restrict to  $g > 1$ , the individual spins all tend to align along the  $z$  direction, and therefore we can define our collective spin operators in reference to this axis,

$$\begin{aligned} S_+ &= \sqrt{N} a, & S_- &= \sqrt{N} a^\dagger, \\ \implies S_x &= \frac{1}{2}(S_+ + S_-), & S_y &= \frac{1}{2i}(S_+ - S_-), \\ S_z &= \frac{N}{2} - a^\dagger a. \end{aligned} \quad (100)$$

In this picture, a state with all spins pointing up along  $z$  corresponds to the vacuum, i.e.,  $\langle a^\dagger a \rangle = 0$ . Plugging these into Eq. (99) we find

$$\begin{aligned} H &= -\frac{2}{N} \left( \frac{1}{4} (S_+ + S_-)(S_+ + S_-) \right) - 2g S_z + \frac{1}{2} \mathbb{1} \\ &= -\frac{2}{N} \left( \frac{N}{4} (a + a^\dagger)(a + a^\dagger) \right) - 2g \left( \frac{N}{2} - a^\dagger a \right) + \frac{1}{2} \\ &= -\frac{1}{2} (a^2 + (a^\dagger)^2 + a^\dagger a + a a^\dagger) - 2g \left( \frac{N}{2} - a^\dagger a \right) + \frac{1}{2}, \end{aligned} \quad (101)$$

and collecting terms together and using the canonical commutation relation,  $a a^\dagger = 1 + a^\dagger a$ , we finally arrive at our Hamiltonian written only in terms of bosonic annihilation and creation operators

$$\begin{aligned} H &= -\frac{1}{2} (a^2 + (a^\dagger)^2) - \frac{1}{2} (a^\dagger a + a a^\dagger) + 2g a^\dagger a \\ &\quad + \frac{1 - 2gN}{2} \\ &= -\frac{1}{2} (a^2 + (a^\dagger)^2) + a^\dagger a (2g - 1) - gN. \end{aligned} \quad (102)$$

We now want to diagonalize this Hamiltonian in order to remove the difficult terms in  $a^2$  and  $(a^\dagger)^2$ . This can be achieved using a Bogoliubov transformation. For the symmetry broken phase ( $g > 1$ ) that we are focusing on, we start by making the following transformation:

$$\begin{aligned} a &= \sinh\left(\frac{\alpha}{2}\right) b^\dagger + \cosh\left(\frac{\alpha}{2}\right) b, \\ a^\dagger &= \sinh\left(\frac{\alpha}{2}\right) b + \cosh\left(\frac{\alpha}{2}\right) b^\dagger. \end{aligned} \quad (103)$$

While at first sight this would seem to complicate things, we will see that it is through a judicious choice of the argument,  $\alpha$ , that is determined by a suitable quotient of the coefficients in Eq. (102), which allows us to diagonalize the model. To proceed, we require the expressions  $(a^2 + (a^\dagger)^2)$  and  $a^\dagger a$  in terms of the new bosonic operators  $b$  and  $b^\dagger$ :

$$a^2 = \left( \sinh\left(\frac{\alpha}{2}\right) b + \cosh\left(\frac{\alpha}{2}\right) b^\dagger \right)^2 = \sinh^2\left(\frac{\alpha}{2}\right) (b^\dagger)^2 + \cosh^2\left(\frac{\alpha}{2}\right) b^2 + \cosh\left(\frac{\alpha}{2}\right) \sinh\left(\frac{\alpha}{2}\right) (1 + 2b^\dagger b), \quad (104)$$

$$(a^\dagger)^2 = \left( \sinh\left(\frac{\alpha}{2}\right) b^\dagger + \cosh\left(\frac{\alpha}{2}\right) b \right)^2 = \sinh^2\left(\frac{\alpha}{2}\right) b^2 + \cosh^2\left(\frac{\alpha}{2}\right) (b^\dagger)^2 + \cosh\left(\frac{\alpha}{2}\right) \sinh\left(\frac{\alpha}{2}\right) (1 + 2b^\dagger b), \quad (105)$$

$$\begin{aligned} \implies a^2 + (a^\dagger)^2 &= \left( \cosh^2\left(\frac{\alpha}{2}\right) + \sinh^2\left(\frac{\alpha}{2}\right) \right) (b^2 + (b^\dagger)^2) + 2 \cosh\left(\frac{\alpha}{2}\right) \sinh\left(\frac{\alpha}{2}\right) (1 + 2b^\dagger b) \\ &= \cosh(\alpha) (b^2 + (b^\dagger)^2) + \sinh(\alpha) (1 + 2b^\dagger b), \end{aligned} \quad (106)$$

$$\begin{aligned} a^\dagger a &= \left( \sinh\left(\frac{\alpha}{2}\right) b + \cosh\left(\frac{\alpha}{2}\right) b^\dagger \right) \left( \sinh\left(\frac{\alpha}{2}\right) b^\dagger + \cosh\left(\frac{\alpha}{2}\right) b \right) \\ &= \sinh\left(\frac{\alpha}{2}\right) \cosh\left(\frac{\alpha}{2}\right) (b^\dagger)^2 + \sinh\left(\frac{\alpha}{2}\right) \cosh\left(\frac{\alpha}{2}\right) b^2 + \cosh^2\left(\frac{\alpha}{2}\right) b^\dagger b + \sinh^2\left(\frac{\alpha}{2}\right) (1 + b^\dagger b) \\ &= \frac{1}{2} \sinh(\alpha) (b^2 + (b^\dagger)^2) + \cosh(\alpha) b^\dagger b + \sinh^2\left(\frac{\alpha}{2}\right), \end{aligned} \quad (107)$$

where we have used the relations  $\sinh(2x) = 2 \sinh(x) \cosh(x)$  and  $\cosh^2(x/2) + \sinh^2(x/2) = \cosh(x)$ . We therefore find the Hamiltonian, Eq. (102), transformed according to Eq. (103) is

$$\begin{aligned} H &= -\frac{1}{2} \left( \cosh(\alpha) (b^2 + (b^\dagger)^2) + \sinh(\alpha) (1 + 2b^\dagger b) \right) \\ &\quad + (2g - 1) \left( \frac{1}{2} \sinh(\alpha) (b^2 + (b^\dagger)^2) + \cosh(\alpha) b^\dagger b + \sinh^2\left(\frac{\alpha}{2}\right) \right) - gN. \end{aligned} \quad (108)$$

While indeed apparently more complicated than our original Hamiltonian, we now make use of the special choice for the argument

$$\tanh(\alpha) = \frac{1}{2g - 1}, \quad (109)$$

which we see is related simply to the ratio of the relevant coefficients appearing in Eq. (102). This then allows us to write Eq. (108) solely in terms of  $\sinh(\alpha)$  since

$$\cosh(\alpha) = \frac{\sinh(\alpha)}{\tanh(\alpha)} = (2g - 1) \sinh(\alpha).$$

Thus, we have

$$\begin{aligned} H &= -\frac{1}{2} \frac{\sinh(\alpha)}{\tanh(\alpha)} (b^2 + (b^\dagger)^2) - \frac{1}{2} \sinh(\alpha) (1 + 2b^\dagger b) + \frac{1}{2} (2g - 1) \sinh(\alpha) (b^2 + (b^\dagger)^2) \\ &\quad + (2g - 1) \frac{\sinh(\alpha)}{\tanh(\alpha)} b^\dagger b + (2g - 1) \sinh^2\left(\frac{\alpha}{2}\right) - gN \\ &= -\frac{1}{2} (2g - 1) \sinh(\alpha) (b^2 + (b^\dagger)^2) + \frac{1}{2} (2g - 1) \sinh(\alpha) (b^2 + (b^\dagger)^2) - \frac{1}{2} \sinh(\alpha) (1 + 2b^\dagger b) \\ &\quad + (2g - 1) (2g - 1) \sinh(\alpha) b^\dagger b + (2g - 1) \sinh^2\left(\frac{\alpha}{2}\right) - gN, \end{aligned} \quad (110)$$

where we now see that the terms involving  $b^2$  and  $(b^\dagger)^2$  cancel and, after a little algebra to collect together terms such that we can express  $H$  in the standard harmonic oscillator form, we are left with

$$H = (4g^2 - 4g) \sinh(\alpha) \left( b^\dagger b + \frac{1}{2} \right) - g(N + 1) + \frac{1}{2}. \quad (111)$$

From Eq. (109) it follows that  $\sinh(\alpha) = (4g^2 - 4g)^{-1/2}$ , finally giving us our desired mapped Hamiltonian

$$H = 2\sqrt{g(g-1)} \left( b^\dagger b + \frac{1}{2} \right) - g(N+1) + \frac{1}{2}. \quad (112)$$

The reason for doing this is that the precise form for the counterdiabatic term is known in the case of a quantum harmonic oscillator [73] with a time-dependent frequency,  $\omega$ , and is given by

$$A = i \frac{\dot{\omega}}{4\omega} (b^2 - (b^\dagger)^2). \quad (113)$$

From Eqs. (104) and (105) we see that

$$\begin{aligned} a^2 - a^{\dagger 2} &= \left( \cosh^2 \left( \frac{\alpha}{2} \right) - \sinh^2 \left( \frac{\alpha}{2} \right) \right) b^2 \\ &\quad + \left( \sinh^2 \left( \frac{\alpha}{2} \right) - \cosh^2 \left( \frac{\alpha}{2} \right) \right) b^{\dagger 2} \\ &= b^2 - b^{\dagger 2}. \end{aligned} \quad (114)$$

Therefore we can readily determine the operators that the counterdiabatic term depends on in the collective spin picture by simply reverting using Eq. (100)

$$\begin{aligned} A_{\text{GS}} &\propto \frac{1}{N} (S_+^2 - S_-^2) \\ &= \frac{1}{N} \left( (S_x + iS_y)^2 - (S_x - iS_y)^2 \right) \\ &= \frac{2i}{N} (S_x S_y + S_y S_x). \end{aligned} \quad (115)$$

The counterdiabatic term for the ground state of the LMG model then follows by identifying the frequency,  $\omega = 2\sqrt{g(g-1)}$ , where we remind that  $g$  is the time-dependent parameter. We therefore finally find that the control term for  $g > 1$  is given by

$$A_{\text{GS}} = -\frac{(2g-1)\dot{g}}{4Ng(g-1)} (S_x S_y + S_y S_x). \quad (116)$$

We again stress that due to the regime of validity of the Holstein-Primakoff approximation, this term is an approximation for any finite  $N$  and only strictly applies to the ground state; nevertheless, it provides useful insight insofar as it indicates the operators that are likely to be highly relevant in achieving control over the system. As it is an approximation, it is prudent to test its effectiveness. For a ramp restricting to the symmetry broken phase, taking  $g(t) = g_0 + g_d(t/T)$  we show in Fig. 4(a) the performance, as determined by the final target state fidelity, of this ramp as we increase system size. For comparison we show unitary dynamics with no counterdiabatic driving

term (blue) and the exact, numerically calculated counterdiabatic term (black), i.e., Eq. (14) evaluated explicitly for the finite sized system. We clearly see that since Eq. (116) is an approximation it does not perfectly suppress the non-adiabatic transitions. However, crucially we see its performance improve as we increase the system size (notice that the apparent plateauing behavior is discussed by Takahashi [74]).

Given that Eq. (116) is an approximation, it leaves open the question as to whether we can do better by leveraging the information we have learned from the mapping, specifically the choice of operators that enter the control term, and examine whether there is a better choice for the driving profile. As a simple demonstration of this concept, somewhat inline with Ref. [43], we consider the functional form of the coefficient in front of the operators in Eq. (116). We can replace this with a simple pulse with a similar profile

$$f(t) = \frac{\exp(t') - b}{N}. \quad (117)$$

The task is then to determine what combination of the parameters  $a$  and  $b$  will give the best final state fidelity while still assuming the same ramping profile for  $g(t)$ , i.e., the linear protocol. For this simple case, we show the landscape in Fig. 4(b) where we consider a fixed system size,  $N = 50$ . We see that remarkably high target fidelities can be achieved. While the choice of pulse here was for ease of demonstration, it is clear that this hybrid approach is one which optimal control techniques are particularly suited to, especially if we allow for a more general pulse profile.

While the above analysis holds for the symmetry broken phase,  $g > 1$ , a qualitatively similar calculation can be done for the symmetric phase. This results in a control Hamiltonian with the same general form as Eq. (116) but with a different coefficient. In doing the mapping, the key difference is that the axis around which the Holstein-Primakoff approximation, Eq. (100), (or equivalently the Bogoliubov transformation) is performed depends on the exact value of  $g$ . We refer to Refs. [43,69,70,74] for more details. A remark regarding the simulation of the model in the symmetric phase: due to the quasi-degeneracy of the different parity sub-sectors, care must be taken when numerically implementing the control since the small energy gaps can lead to numerical instabilities in the dynamics.

The above example demonstrates a key message of the present tutorial: hybrid approaches to control offer a powerful framework to coherently manipulate complex quantum systems. In particular, we see that the  $N \rightarrow \infty$  limit serves as a good approximation for sufficiently large systems that we can use the functional form that the exact control term predicts and supplement it with optimal control techniques to achieve almost perfect performance. Naturally, one could have directly examined the problem from

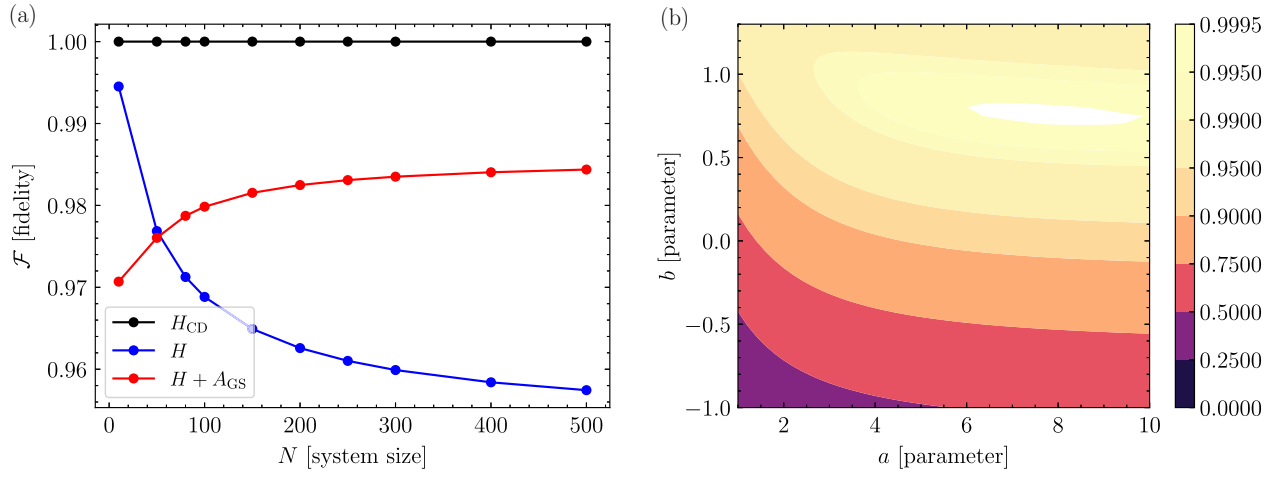


FIG. 4. (a) Ramping within the symmetry broken phase of the all-to-all model, with  $g(t) = 2 - 0.9t$  where the aim is to connect the ground state of  $H$  at  $t = 0$  with the ground state at  $t = 1$ . We show the final target state fidelity at  $t = 1$  as a function of system size for the exact, numerically evaluated CD (black), bare evolution (blue), and using the QHO approximation, Eq. (116), (red). (b) We fix  $N = 50$  and use the guess pulse Eq. (117) for the control given by the QHO approximation and sweep over the achievable fidelities.

an optimal control stand point following, e.g., Ref. [75], dictating from the outset what operators to include in the control Hamiltonian and then seeking the best pulse, and for a model as relatively simple as Eq. (98) this would be feasible. The advantage of the above approach however is that it allows us to immediately see why a particular subset of operators are important in suppressing non-adiabatic transitions, a point that will be again explored in the next example.

The Ising and LMG models demonstrate that the exact CD terms for many-body systems necessitate non-local interactions. While some non-local interactions can be tuned in some systems, e.g., atoms coupled to cavity modes [76], clearly there is an extremely high overhead in terms of the control needed over various ranges of interactions to implement the counterdiabatic term, cf. Eq. (87) for the Ising model. In general, we find that interaction ranges necessary to achieve exact CD driving are not fixed and can be of higher order than the terms in the bare Hamiltonian. Indeed, this is the case with the Ising model in Eq. (87): when we map the non-local two-body terms obtained for spinless fermions back to the spin-1/2 basis, we find that we generate  $N/2$ -body terms with periodic boundaries and  $N$  being the system size [42]. While one approach to tackle this issue is to come up with new and novel ways in which to realize exact CD, e.g., by shaking the system in a particular way [50], another is to ask how important such long-range terms truly are, cf. the discussion around Fig. 3. In most scenarios, especially when keeping away from critical points, to obtain a substantial speed-up through a CD control protocol it may be sufficient to consider only local approximations of the exact CD [46,77]. In scenarios where it is possible to derive the exact counterdiabatic Hamiltonian then it is possible to take truncations of this

in a sensible manner by considering dominant terms or taking a finite cut-off in the range of the exactly derived counterdiabatic terms, as demonstrated in Ref. [46] for the Ising model wherein it is shown that the variational adiabatic gauge potential provides the most insightful way to construct the control term. We next demonstrate this while also showing how the variational approach can be utilized for non-integrable models.

#### 4. Example: Variational counterdiabatic driving

The exact CD driving approach is powerful when feasible but limited in two central ways: (i) it requires the instantaneous Schrödinger equation to be solvable at all times of a given protocol, and (ii) it often results in a highly demanding non-local driving term to be implemented which can be beyond the capabilities of the hardware. The approach of variational CD driving and the use of the AGP, as detailed in Sec. III E was introduced in order to tackle both of these limitations at the cost of constructing only an approximate CD term.

We will take as an example the 1D Ising model, but will also include terms to make it non-integrable so that the Hamiltonian is given by

$$H_0(\lambda) = -J(\lambda) \sum_j \sigma_j^z \sigma_{j+1}^z - Z(\lambda) \sum_j \sigma_j^z + X(\lambda) \sum_j \sigma_j^x, \quad (118)$$

with  $\lambda$  being the time-dependent driving field. In this example, we will consider in detail the case of periodic boundary conditions for  $N > 2$ , however, we will outline where the boundary conditions become important and what happens in the case of  $N = 2$  for which we will derive the

exact CD utilizing the variational approach. In the limit of  $N = 1$  this model becomes equivalent to the Landau-Zener model considered in Sec. II F 1.

From the commutator ansatz of Eq. (53) we can find all possible operators on which the AGP can have non-zero support. However, this can be an exponentially large number of operators for an arbitrary  $N$  size system. In this example, we will be interested in applying approximate CD driving up to two-body terms and we can turn to the commutator ansatz to obtain the correct set of operators for which the AGP can have support for up to this limit. Before moving further forward it is useful to write out the commutation rules

$$\begin{aligned} [dA, B] &= d[A, B] \\ [A, B + C] &= [A, B] + [A, C] \\ [A, BC] &= B[A, C] + [A, B]C \\ [AB, C] &= A[B, C] + [A, B]B \end{aligned} \quad (119)$$

with  $A, B, C$  being operators and  $d$  a scalar. Since they will be heavily used in the following, we explicitly state the commutator relations for the Pauli matrices

$$\begin{aligned} [\sigma_j^x, \sigma_k^y] &= 2i\sigma_j^z \delta_{j,k} \\ [\sigma_j^y, \sigma_k^z] &= 2i\sigma_j^x \delta_{j,k} \\ [\sigma_j^z, \sigma_k^x] &= 2i\sigma_j^y \delta_{j,k} \\ [\sigma_j^z, \sigma_k^z] &= [\sigma_j^x, \sigma_k^x] = [\sigma_j^y, \sigma_k^y] = 0 \end{aligned} \quad (120)$$

with  $\delta_{j,k}$  the Kronecker delta, both of which we will use frequently in this example. We start by considering the first term of the commutator expansion, where we will drop the  $\lambda$ -dependence of each parameter for ease of reading,

$$\begin{aligned} A_\lambda^{(1)} &= i\alpha_1 [H_0(\lambda), \partial_\lambda H_0(\lambda)] = i\alpha_1 \sum_{j,k} \left( -J\dot{X} [\sigma_j^z \sigma_{j+1}^z, \sigma_k^x] - Z\dot{X} [\sigma_j^z, \sigma_k^x] - X\dot{J} [\sigma_j^x, \sigma_k^z \sigma_{k+1}^z] - X\dot{Z} [\sigma_j^x, \sigma_k^z] \right) \\ &= i\alpha_1 \sum_{j,k} \left( (X\dot{J} - J\dot{X}) (\sigma_j^z [\sigma_{j+1}^z, \sigma_k^x] + [\sigma_j^z, \sigma_k^x] \sigma_{j+1}^z) + (X\dot{Z} - Z\dot{X}) [\sigma_j^z, \sigma_k^x] \right) \\ &= -2\alpha_1 \sum_j \left( (X\dot{J} - J\dot{X}) (\sigma_j^z \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^z) + (X\dot{Z} - Z\dot{X}) \sigma_j^y \right). \end{aligned} \quad (121)$$

Now we want to utilize only the information about what operators can be generated by the commutator ansatz up to two-body terms. As a result, we will not utilize the coefficients derived in front of said operators and from Eq. (121) we would write the ansatz for the AGP as

$$A_\lambda^{(1)} = \alpha \sum_j \sigma_j^y + \gamma \sum_j (\sigma_j^z \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^z), \quad (122)$$

with  $\alpha$  and  $\gamma$  being the coefficients we aim to determine. The next step would be to derive the second term of the commutator expansion, but since we are only interested in up to two body terms and the calculation of the second term is tedious, we will simply quote the result by stating that the full ansatz for the AGP including up to two-body terms is given by

$$A_\lambda^{(2)} = \alpha \sum_j \sigma_j^y + \beta \sum_j (\sigma_j^x \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^x) + \gamma \sum_j (\sigma_j^z \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^z), \quad (123)$$

where we also now need to find  $\beta$ . For comparison, we will write the ansatz for the AGP up to one-body terms as

$$A_\lambda^{(1)} = \alpha \sum_j \sigma_j^y, \quad (124)$$

and note that the  $\alpha$  coefficients for the one-body and two-body terms are not necessarily equivalent.

We now need to derive the distance we will minimize according to Sec. II E, first, we need to obtain  $G_\lambda(A_\lambda^{(2)})$

$$\begin{aligned}
G_\lambda(A_\lambda^{(2)}) &= \partial_\lambda H_0 + i \left[ A_\lambda^{(2)}, H_0 \right] \\
&= -J \sum_j \sigma_j^z \sigma_{j+1}^z - \dot{Z} \sum_j \sigma_j^z + \dot{X} \sum_j \sigma_j^x + i\alpha \sum_{j,k} \left[ \sigma_j^y, -J \sigma_k^z \sigma_{k+1}^z - \dot{Z} \sigma_k^z + \dot{X} \sigma_k^x \right] \\
&+ i\beta \sum_{j,k} \left[ \sigma_j^x \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^x, -J \sigma_k^z \sigma_{k+1}^z - \dot{Z} \sigma_k^z + \dot{X} \sigma_k^x \right] + i\gamma \sum_{j,k} \left[ \sigma_j^z \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^z, -J \sigma_k^z \sigma_{k+1}^z - \dot{Z} \sigma_k^z + \dot{X} \sigma_k^x \right].
\end{aligned} \tag{125}$$

At this point, it is useful to consider each one of the commutators in turn and we will look at each of these in the order they are presented above. Starting with the one corresponding to  $\sigma_j^y$ ,

$$\begin{aligned}
&i\alpha \sum_{j,k} \left[ \sigma_j^y, -J \sigma_k^z \sigma_{k+1}^z - \dot{Z} \sigma_k^z + \dot{X} \sigma_k^x \right] \\
&= i\alpha \sum_j \left( -J \sum_k \left( \sigma_k^z \left[ \sigma_j^y, \sigma_{k+1}^z \right] + \left[ \sigma_j^y, \sigma_k^z \right] \sigma_{k+1}^z \right) - Z \sum_k \left[ \sigma_j^y, \sigma_k^z \right] + X \sum_k \left[ \sigma_j^y, \sigma_k^x \right] \right) \\
&= i\alpha \sum_j \left( -J \sum_k \left( \sigma_k^z \sigma_j^x \delta_{j,k+1} (2i) + \sigma_j^x \delta_{j,k} (2i) \sigma_{k+1}^z \right) - Z \sum_k (2i) \delta_{j,k} \sigma_j^x + X \sum_k (-2i) \delta_{j,k} \sigma_j^z \right) \\
&= 2J\alpha \sum_j \left( \sigma_k^z \sigma_j^x + \sigma_j^x \sigma_{k+1}^z \right) + 2Z\alpha \sum_j \sigma_j^x + 2X\alpha \sum_j \sigma_j^z.
\end{aligned} \tag{126}$$

Note that above, and in other places in the derivation of this example we use a subtle trick of labeling, the use of which can be frustrating for those who are not aware of it. For two operators labeled by  $j$  and  $k$ , where  $j$  and  $k$  are summed independently over the same values, then we can always swap  $j$  and  $k$ . This can be seen in a simple example

$$\sum_{j,k} (A_j B_k + A_k B_j) = \sum_{j,k} (A_j B_k + A_j B_k) = 2 \sum_{j,k} A_j B_k, \tag{127}$$

while this trick is rather obvious when stated so bluntly, it is surprising how amid long analytical derivations such a trick can be missed, and how software, e.g., Mathematica and Maple, can not easily implement this in general.

We now look at the commutator corresponding to  $\sigma_j^x \sigma_{j+1}^y$  terms, which is rather involved, and requires the application of the commutation relation rules many times over. To start, we expand the commutator to get every term into the form of the commutation of two Pauli operators

$$\begin{aligned}
&i\beta \sum_{j,k} \left[ \sigma_j^x \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^x, -J \sigma_k^z \sigma_{k+1}^z - \dot{Z} \sigma_k^z + \dot{X} \sigma_k^x \right] \\
&= i\beta \sum_j \left( -J \sum_k \left( \sigma_j^x \sigma_k^z \left[ \sigma_{j+1}^y, \sigma_{k+1}^z \right] + \sigma_j^x \left[ \sigma_{j+1}^y, \sigma_k^z \right] \sigma_{k+1}^z + \sigma_k^z \left[ \sigma_j^x, \sigma_{k+1}^z \right] \sigma_{j+1}^y + \left[ \sigma_j^x, \sigma_k^z \right] \sigma_{k+1}^z \sigma_{j+1}^y \right. \right. \\
&\quad \left. \left. + \sigma_j^y \sigma_k^z \left[ \sigma_{j+1}^x, \sigma_{k+1}^z \right] + \sigma_j^y \left[ \sigma_{j+1}^x, \sigma_k^z \right] \sigma_{k+1}^z + \sigma_k^z \left[ \sigma_j^y, \sigma_{k+1}^z \right] \sigma_{j+1}^x + \left[ \sigma_j^y, \sigma_k^z \right] \sigma_{k+1}^z \sigma_{j+1}^x \right) \right. \\
&\quad \left. - Z \sum_k \left( \sigma_j^x \left[ \sigma_{j+1}^y, \sigma_k^z \right] + \left[ \sigma_j^x, \sigma_k^z \right] \sigma_{j+1}^y + \sigma_j^y \left[ \sigma_{j+1}^x, \sigma_k^z \right] + \left[ \sigma_j^y, \sigma_k^z \right] \sigma_{j+1}^x \right) \right. \\
&\quad \left. + X \sum_k \left( \sigma_j^x \left[ \sigma_{j+1}^y, \sigma_k^x \right] + \left[ \sigma_j^y, \sigma_k^x \right] \sigma_{j+1}^x \right) \right).
\end{aligned} \tag{128}$$

We can then insert the known commutation relations and gather like terms

$$\begin{aligned}
& i\beta \sum_{j,k} \left[ \sigma_j^x \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^x, -J \sigma_k^z \sigma_{k+1}^z - \dot{Z} \sigma_k^z + \dot{X} \sigma_k^x \right] \\
&= i\beta \sum_j \left( -J \sum_k \left( \sigma_j^x \sigma_k^z (2i\delta_{j,k}) \sigma_{j+1}^x + \sigma_j^x (2i\delta_{j+1,k}) \sigma_{j+1}^x \sigma_{k+1}^z + \sigma_k^z (-2i\delta_{j,k+1}) \sigma_j^y \sigma_{j+1}^y + (-2i\delta_{j,k}) \sigma_j^y \sigma_{k+1}^z \sigma_{j+1}^y \right. \right. \\
&\quad \left. \left. + \sigma_j^y \sigma_k^z (-2i\delta_{j,k}) \sigma_{j+1}^y + \sigma_j^y (-2i\delta_{j+1,k}) \sigma_{j+1}^y \sigma_{k+1}^z + \sigma_k^z (2i\delta_{j,k+1}) \sigma_j^x \sigma_{j+1}^x + (2i\delta_{j,k}) \sigma_j^x \sigma_{k+1}^z \sigma_{j+1}^x \right) \right. \\
&\quad \left. - Z \sum_k \left( \sigma_j^x (2i\delta_{j+1,k}) \sigma_{j+1}^x + (-2i\delta_{j,k}) \sigma_j^y \sigma_{j+1}^y + \sigma_j^y (-2i\delta_{j+1,k}) \sigma_{j+1}^y + (2i\delta_{j,k}) \sigma_j^x \sigma_{j+1}^x \right) \right. \\
&\quad \left. + X \sum_k \left( \sigma_j^x (-2i\delta_{j+1,k}) \sigma_{j+1}^z + (-2i\delta_{j,k}) \sigma_j^z \sigma_{j+1}^x \right) \right) \\
&= 2J\beta \sum_j \left( \sigma_j^x \sigma_j^z \sigma_{j+1}^x + \sigma_j^x \sigma_{j+1}^z \sigma_{j+1}^x - \sigma_j^y \sigma_j^z \sigma_{j+1}^y - \sigma_j^y \sigma_{j+1}^z \sigma_{j+1}^y \right) \\
&\quad + 2J\beta \sum_j \left( \sigma_j^x \sigma_{j+1}^x \sigma_{j+2}^z + \sigma_j^z \sigma_{j+1}^x \sigma_{j+2}^x - \sigma_j^z \sigma_{j+1}^y \sigma_{j+2}^y - \sigma_j^y \sigma_{j+1}^z \sigma_{j+2}^z \right) \\
&\quad + 4Z\beta \sum_j \left( \sigma_j^x \sigma_{j+1}^x - \sigma_j^y \sigma_{j+1}^y \right) + 2X\beta \sum_j \left( \sigma_j^x \sigma_{j+1}^z + \sigma_j^z \sigma_{j+1}^x \right). \tag{129}
\end{aligned}$$

Now, we can use the fact that we can always combine Pauli operators acting on the same site, as is present above, in this case we use that  $\sigma_j^x \sigma_j^z = i\sigma_j^y$  and  $\sigma_j^z \sigma_j^y = i\sigma_j^x$  to obtain

$$\begin{aligned}
& i\beta \sum_{j,k} \left[ \sigma_j^x \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^x, -J \sigma_k^z \sigma_{k+1}^z - \dot{Z} \sigma_k^z + \dot{X} \sigma_k^x \right] \\
&= 2iJ\beta \sum_j \left( \sigma_j^y \sigma_{j+1}^x + \sigma_j^x \sigma_{j+1}^y - \sigma_j^x \sigma_{j+1}^y - \sigma_j^y \sigma_{j+1}^x \right) \\
&\quad + 2J\beta \sum_j \left( \sigma_j^x \sigma_{j+1}^x \sigma_{j+2}^z + \sigma_j^z \sigma_{j+1}^x \sigma_{j+2}^x - \sigma_j^z \sigma_{j+1}^y \sigma_{j+2}^y - \sigma_j^y \sigma_{j+1}^z \sigma_{j+2}^z \right) \\
&\quad + 4Z\beta \sum_j \left( \sigma_j^x \sigma_{j+1}^x - \sigma_j^y \sigma_{j+1}^y \right) + 2X\beta \sum_j \left( \sigma_j^x \sigma_{j+1}^z + \sigma_j^z \sigma_{j+1}^x \right) \\
&= 2J\beta \sum_j \left( \sigma_j^x \sigma_{j+1}^x \sigma_{j+2}^z + \sigma_j^z \sigma_{j+1}^x \sigma_{j+2}^x - \sigma_j^z \sigma_{j+1}^y \sigma_{j+2}^y - \sigma_j^y \sigma_{j+1}^z \sigma_{j+2}^z \right) \\
&\quad + 4Z\beta \sum_j \left( \sigma_j^x \sigma_{j+1}^x - \sigma_j^y \sigma_{j+1}^y \right) + 2X\beta \sum_j \left( \sigma_j^x \sigma_{j+1}^z + \sigma_j^z \sigma_{j+1}^x \right). \tag{130}
\end{aligned}$$

In the final stretch, we can turn to the commutator corresponding to  $\sigma_j^z \sigma_{j+1}^y$  terms, following a similar procedure to above we first expand the commutator

$$\begin{aligned}
& i\gamma \sum_{j,k} \left[ \sigma_j^z \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^z, -J \sigma_k^z \sigma_{k+1}^z - \dot{Z} \sigma_k^z + \dot{X} \sigma_k^x \right] \\
&= i\gamma \sum_j \left( -J \sum_k \left( \sigma_j^z \left[ \sigma_{j+1}^y, \sigma_k^z \sigma_{k+1}^z \right] + \left[ \sigma_j^y, \sigma_k^z \sigma_{k+1}^z \right] \sigma_{j+1}^z \right) - Z \sum_k \left( \sigma_j^z \left[ \sigma_{j+1}^y, \sigma_k^z \right] + \left[ \sigma_j^y, \sigma_k^z \right] \sigma_{j+1}^z \right) \right. \\
&\quad \left. + X \sum_k \left( \sigma_j^z \left[ \sigma_{j+1}^y, \sigma_k^x \right] + \left[ \sigma_j^y, \sigma_k^x \right] \sigma_{j+1}^z + \sigma_j^y \left[ \sigma_{j+1}^z, \sigma_k^x \right] + \left[ \sigma_j^z, \sigma_k^x \right] \sigma_{j+1}^y \right) \right)
\end{aligned}$$

$$\begin{aligned}
&= i\gamma \sum_j \left( -J \sum_k \left( \sigma_j^z \sigma_k^z \left[ \sigma_{j+1}^y, \sigma_{k+1}^z \right] + \sigma_j^z \left[ \sigma_{j+1}^y, \sigma_k^z \right] \sigma_{k+1}^z + \sigma_k^z \left[ \sigma_j^y, \sigma_{k+1}^z \right] \sigma_{j+1}^z + \left[ \sigma_j^y, \sigma_k^z \right] \sigma_{k+1}^z \sigma_{j+1}^z \right) \right. \\
&\quad \left. - Z \sum_k \left( \sigma_j^z \left[ \sigma_{j+1}^y, \sigma_k^z \right] + \left[ \sigma_j^y, \sigma_k^z \right] \sigma_{j+1}^z \right) + X \sum_k \left( \sigma_j^z \left[ \sigma_{j+1}^y, \sigma_k^x \right] + \left[ \sigma_j^y, \sigma_k^x \right] \sigma_{j+1}^z + \sigma_j^y \left[ \sigma_{j+1}^z, \sigma_k^x \right] + \left[ \sigma_j^z, \sigma_k^x \right] \sigma_{j+1}^y \right) \right). \tag{131}
\end{aligned}$$

We can then insert the known Pauli commutation relations

$$\begin{aligned}
&i\gamma \sum_{j,k} \left[ \sigma_j^z \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^z, -J \sigma_k^z \sigma_{k+1}^z - \dot{Z} \sigma_k^z + \dot{X} \sigma_k^x \right] \\
&= i\gamma \sum_j \left( -J \sum_k \left( \sigma_j^z \sigma_k^z (2i\delta_{j,k}) \sigma_{j+1}^x + \sigma_j^z (2i\delta_{j+1,k}) \sigma_{j+1}^x \sigma_{k+1}^z + \sigma_k^z (2i\delta_{j,k+1}) \sigma_j^x \sigma_{j+1}^z + (2i\delta_{j,k}) \sigma_j^x \sigma_{k+1}^z \sigma_{j+1}^z \right) \right. \\
&\quad \left. - Z \sum_k \left( \sigma_j^z (2i\delta_{j+1,k}) \sigma_{j+1}^x + (2i\delta_{j,k}) \sigma_j^x \sigma_{j+1}^z \right) \right. \\
&\quad \left. + X \sum_k \left( \sigma_j^z (-2i\delta_{j+1,k}) \sigma_{j+1}^z + (-2i\delta_{j,k}) \sigma_j^z \sigma_{j+1}^z + \sigma_j^y (2i\delta_{j+1,k}) \sigma_{j+1}^y + (2i\delta_{j,k}) \sigma_j^y \sigma_{j+1}^y \right) \right). \tag{132}
\end{aligned}$$

It is at this point that a peculiar thing happens, as the square of any Pauli matrix if the identity,  $\mathbb{1}_j$ , we get the following:

$$\begin{aligned}
&i\gamma \sum_{j,k} \left[ \sigma_j^z \sigma_{j+1}^y + \sigma_j^y \sigma_{j+1}^z, -J \sigma_k^z \sigma_{k+1}^z - \dot{Z} \sigma_k^z + \dot{X} \sigma_k^x \right] \\
&= 4J\gamma \sum_j \sigma_j^z \sigma_{j+1}^x \sigma_{j+2}^z + 2J\gamma \sum_k \left( \sigma_j^x \mathbb{1}_{j+1} + \mathbb{1}_j \sigma_{j+1}^x \right) + 2Z\gamma \sum_j \left( \sigma_j^z \sigma_{j+1}^x + \sigma_j^x \sigma_{j+1}^z \right) + 4X\gamma \sum_j \left( \sigma_j^z \sigma_{j+1}^z - \sigma_j^y \sigma_{j+1}^y \right) \\
&= 4J\gamma \sum_j \sigma_j^z \sigma_{j+1}^x \sigma_{j+2}^z + 4J\gamma \sum_k \sigma_j^x + 2Z\gamma \sum_j \left( \sigma_j^z \sigma_{j+1}^x + \sigma_j^x \sigma_{j+1}^z \right) + 4X\gamma \sum_j \left( \sigma_j^z \sigma_{j+1}^z - \sigma_j^y \sigma_{j+1}^y \right). \tag{133}
\end{aligned}$$

Note, we highlight this step involving the  $\mathbb{1}_j$  terms, as this single-body operator in the second term is truly being obtained from a two-body operator. This has consequences dependent on the boundary conditions, and if  $N = 2$  then the resulting one-body term in the final line should have a coefficient of  $2J$  instead of  $4J$ . We can now return to Eq. (125) and, gathering like with like terms from Eqs. (126), (130) and (133), we obtain

$$\begin{aligned}
G_\lambda \left( A_\lambda^{(2)} \right) &= (4X\gamma - J) \sum_j \sigma_j^z \sigma_{j+1}^z + (2J\alpha + 2X\beta + 2Z\gamma) \sum_j \left( \sigma_j^x \sigma_{j+1}^z + \sigma_j^z \sigma_{j+1}^x \right) + (2Z\alpha + 4J\gamma + \dot{X}) \sum_j \sigma_j^x \\
&\quad + (2X\alpha + \dot{Z}) \sum_j \sigma_j^z + 2J\beta \sum_j \left( \sigma_j^x \sigma_{j+1}^x \sigma_{j+2}^z + \sigma_j^z \sigma_{j+1}^x \sigma_{j+2}^z \right) - 2J\beta \sum_j \left( \sigma_j^y \sigma_{j+1}^y \sigma_{j+2}^z + \sigma_j^z \sigma_{j+1}^y \sigma_{j+2}^z \right) \\
&\quad + 4Z\beta \sum_j \sigma_j^x \sigma_{j+1}^x - (4Z\beta + 4X\gamma) \sum_j \sigma_j^y \sigma_{j+1}^y + 4J\gamma \sum_j \sigma_j^z \sigma_{j+1}^z + \sigma_j^z. \tag{134}
\end{aligned}$$

Note, we just subjected ourselves and the reader to a lot of algebra with the tracking of seemingly countless commutation relations and operators. This is done in part to, of course, educate the reader on these techniques, but also as a cautionary tale that this approach is not scalable and is therefore prone to human error. We will present current leading alternatives to this pen and paper method outlined here in Sec. VI.

Following the approach of Sec. II E, we now want to minimize  $\text{Tr}[G_\lambda^2(A_\lambda^{(2)})]$  which at first can appear quite daunting as we now need to square the long expression given above. However, Pauli matrices are traceless, i.e.,

$$\text{Tr} \left[ \sigma_j^x \right] = \text{Tr} \left[ \sigma_j^y \right] = \text{Tr} \left[ \sigma_j^z \right] = 0,$$

which means that any term in  $G_\lambda^2(A_\lambda^{(2)})$  with a single Pauli matrix will not contribute to  $\text{Tr}[G_\lambda^2(A_\lambda^{(2)})]$  and that only terms with Pauli matrices squared will be present, and that these will have a trace of 2. Therefore, to write the  $\text{Tr}[G_\lambda^2(A_\lambda^{(2)})]$  we skip calculating  $G_\lambda^2(A_\lambda^{(2)})$  and instead square the coefficients in front of each orthogonal operator giving

$$N2^{-N}\text{Tr}[G_\lambda^2] = (4X\gamma - j)^2 + (2J\alpha + 2X\beta + 2Z\gamma)^2 + (2Z\alpha + 4J\gamma + \dot{X})^2 + (2X\alpha + \dot{Z})^2 + 8J^2\beta^2 + 16Z^2\beta^2 - (4Z\beta + 4X\gamma)^2 + 16J^2\gamma^2. \quad (135)$$

Note, we have dropped the  $A_\lambda^{(2)}$  dependence of  $G_\lambda$  for ease of reading. However, it is also at this step that boundary conditions become important, as each of the two- and three-body terms will be present a different number of times and will have coefficients of, e.g.,  $N - 1$  for two-body terms with open boundary conditions, instead of the  $N$  for the case of periodic boundary conditions. It is simpler to think of the  $N = 2$  case as explicitly different, as there will be no three-body terms and any sum over two-body terms is irrelevant.

The final step in obtaining the coupled set of equations that we must solve to obtain  $\alpha, \beta, \gamma$  is to differentiate and minimize with respect to each coefficient of our ansatz, i.e., consider

$$\frac{\partial \text{Tr}[G_\lambda^2]}{\partial \alpha} = 0, \quad \frac{\partial \text{Tr}[G_\lambda^2]}{\partial \beta} = 0, \quad \text{and} \quad \frac{\partial \text{Tr}[G_\lambda^2]}{\partial \gamma} = 0.$$

We can then write the obtained coupled set of linear equations as

$$\begin{pmatrix} 4J^2 + 2X^2 + 2Z^2 & 4JX & 8JZ \\ JX & 2J^2 + X^2 + 4Z^2 & 3XZ \\ 4JZ & 6XZ & 8J^2 + 8X^2 + 2Z^2 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix} = \begin{pmatrix} \dot{X}Z - X\dot{Z} \\ 0 \\ \dot{X}J - X\dot{J} \end{pmatrix}. \quad (136)$$

For completeness, we will also state the coupled set of equations for the case of  $N = 2$

$$\begin{pmatrix} 2J^2 + 2X^2 + 2Z^2 & 2JX & 4JZ \\ 2JX & 2X^2 + 4Z^2 & 6XZ \\ 4JZ & 6XZ & 2J^2 + 8X^2 + 2Z^2 \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix} = \begin{pmatrix} \dot{X}Z - X\dot{Z} \\ 0 \\ \dot{X}J - X\dot{J} \end{pmatrix}, \quad (137)$$

as we will utilize these in the numerical example below, and it serves as a good test of the approach as we expect the variational CD with up to two-body terms to be the exact CD as there are only two bodies.

We now turn to a concrete example of the implementation of variational CD in the Ising model. We will consider a system described by Hamiltonian Eq. (118) with the following parameters:

$$J(\lambda) = J, \quad Z(\lambda) = J, \quad X(\lambda) = 2J\lambda,$$

with  $\lambda$  being the time-dependent linear function

$$\lambda(t) = \frac{t}{T}, \quad (138)$$

with  $t$  going from 0 to  $T$  where  $T$  is the total driving time. This allows us to explore both the adiabatic  $T \gg J^{-1}$  and diabatic  $T \ll J^{-1}$  regimes. First, we look to confirm that the variational CD terms are correctly derived by considering the case of  $N = 2$  in Fig. 5(a). Indeed, we observe that for all total driving times,  $T$ , the fidelity is perfect

to single precision, whereas the one-body corrections only compensate some transitions, bringing the fidelity to  $\sim 0.9$ . We note here, that there are limited other states for any excited dynamical state to populate with the Hilbert space consisting of only four states, as a result, the bare dynamical protocol with no variational CD still has an overlap of  $\sim 0.6$  with the ground state. Moving to larger systems of  $N = 3$  and  $N = 4$  sites in Figs. 5(b) and 5(c) we observe a clear decrease in the impact of variational CD to only two-body terms. This should be expected as the exact CD will include up to three- and four-body terms in the respective cases. The improvement that two-body variational CD can deliver in these larger systems is still substantial compared to the bare dynamical protocol, showing the utility of low-order variational CD when possible. In Fig. 5(d) we investigate the scaling of the fidelity for the bare and variational CD protocols for a fixed time  $T = 0.01J^{-1}$ . We also briefly note that this approach of variational CD can be enhanced via the introduction of control fields [46,53], thus combining the methodologies described in this section with that of Sec. III.

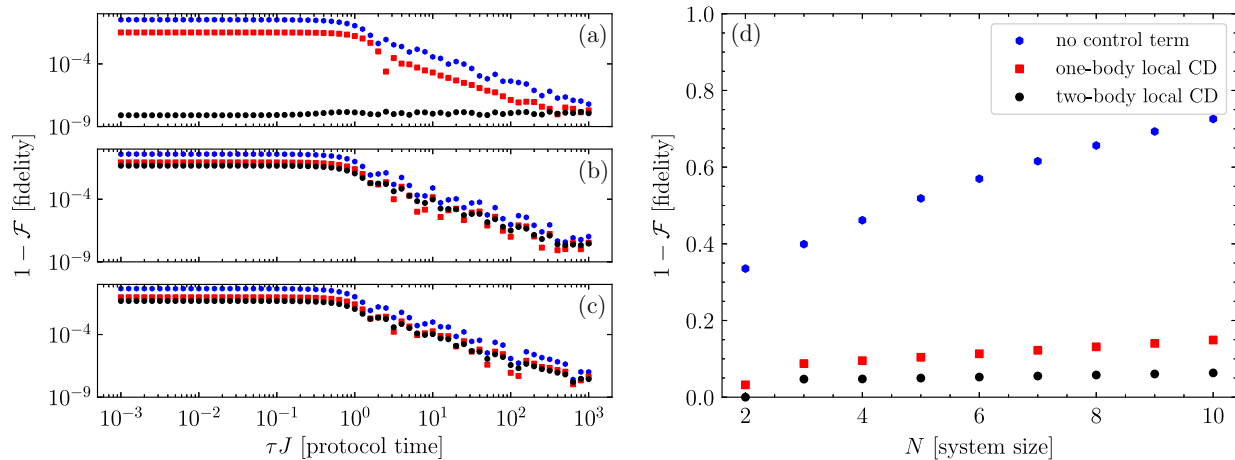


FIG. 5. Implementation of local CD in the Ising spin model. The case of no control term(s) being implemented (blue) is compared to implementing one-body local CD (red) and two-body local CD (black). (a)–(c) The achieved fidelity ( $1 - \mathcal{F}$ ) for different total protocol times for (a)  $N = 2$ , (b)  $N = 3$ , and (c)  $N = 4$  spins. Note, for the case of (a)  $N = 2$  the two-body local CD gives unit fidelity to below single precision. (d) The achieved final fidelity for a fixed total time,  $T = 0.01J^{-1}$ , for various system sizes showing the saturation of the fidelity for the one- and two-body implementations.

### G. Experimental implementations of counterdiabatic driving

The first experimental demonstrations of such shortcuts to adiabaticity (STA) were performed using trapped Bose-Einstein condensates (BECs) [78,79]. In particular, control of the LZ model was achieved using a BEC trapped in an optical lattice and an effective two-level system corresponding to Eq. (54) was accessed, in quasi-momentum space, by accelerating the gas [79]. In this setting the exact CD dynamics, Eq (59), can be achieved in two ways: either directly introducing a second optical lattice that induces the required control field or by making a suitable transformation to the Hamiltonian, in effect absorbing the control field into modulations of the bare Hamiltonian parameters. The latter approach being clearly more experimentally preferable as it requires the manipulation of only a single optical potential. Indeed performing such a transformation to achieve the control is characteristic of many experimental realization of STA protocols [80–82]. The realized protocol was shown to be remarkably robust. Subsequently control of the LZ model was demonstrated in another experimental platform consisting of a single nitrogen-vacancy (NV) center in diamond, where the effective two-level system was achieved by driving a particular subspace spanned by states of the electron spin [83]. This implementation allowed us to demonstrate the remarkable speed-up that can be attained with CD driving, achieving an improvement of between 2 and 3 orders of magnitude.

Counterdiabatic techniques have since been applied to more complex settings starting with three-level systems where CD control can significantly speed up stimulated Raman adiabatic passage protocols [80,82,84], continuous

variable control of a trapped ion [81], and superconducting Xmon set ups to achieve high fidelity quantum gates [85] and characterizing the thermodynamics of the control protocol [86]. The energetics of STA protocols has been recently explored in the context of the Landau-Zener model, Eq. (59), using a single trapped  $^{40}\text{Ca}^+$  ion [87] in order to experimentally demonstrate the trade-off between resource consumption and achievable control [60].

Approximate CD driving has been demonstrated in controlling transport of population in a synthetic lattice of momentum states in ultracold  $^{87}\text{Rb}$  with a 2 orders of magnitude improvement in population transfer [88] and in ground state preparation in a nuclear-magnetic-resonance setup [89]. The commutator expansion form of the AGP has led to the development of a new family of variational quantum algorithms, so-called CD quantum approximate optimization algorithm (QAOA) [90,91] which has been implemented on current hardware, including examining potential applications in protein folding [92] and finance [93].

### III. QUANTUM OPTIMAL CONTROL

In the previous section we discussed the basics of adiabatic control and the toolbox of shortcuts to adiabaticity. That framework assumed that the initial and target states of a control problem are given in terms of eigenstates of our Hamiltonian for different values of the control parameter. In some scenarios, however, this might not be the case. In this section we introduce the tools of quantum optimal control, which will allow us to tackle arbitrary quantum control problems by systematically searching for the optimal form of the control fields that drive a quantum

TABLE III. Summary of notation for the most common quantities used in quantum optimal control.

Object	Proposed symbol
Target state	$ \psi_*\rangle$
Fidelity (state)	$\mathcal{F}_S$
Target unitary	$ U_*\rangle$
Fidelity (unitary)	$\mathcal{F}_U$
Control field	$\mathbf{a}(t)$
Cost functional	$J[\mathbf{a}(t)]$
Control parameters	$\boldsymbol{\alpha} = \alpha_1, \dots, \alpha_M$
Number of time steps	$M$
Time step length	$\Delta t$
Total evolution time	$T$
Quantum speed limit time	$\tau_{\text{QSL}}^*$

system to given target configuration. Typical optimal control implementations require repeated numerical evaluations of the full dynamics of a quantum system, and thus are often, in practice, restricted to small systems. Nevertheless, the success of optimal control techniques in solving control problems in few-body settings is a testament to its importance in the field. In the following we introduce and illustrate these tools by revisiting some of the models introduced in Sec. II and analyzing them through the lens of optimal control. First, we will study optimal state control for a single qubit system obeying the Landau-Zener Hamiltonian (cf. Sec. II F 1), and then its extension to unitary gate preparation with a model of phase-only control. Finally, we will apply optimal control to the problem of generating entangled states in an extended version of the all-to-all Ising model introduced in Sec. II F 3.

### A. Basic statement of an optimal control problem

In any quantum control problem, the goal is to maximize the success of a control protocol in driving a quantum system towards a predefined target configuration. For problems of state control, like the ones discussed in the previous section, this can be theoretically quantified by the quantum state fidelity, which measures how accurately the state at the final time  $|\psi(T)\rangle$  approaches the target state  $|\psi_*\rangle$  (typically, up to a global phase). This reads

$$\mathcal{F}_S = |\langle \psi_* | \psi(T) \rangle|^2, \quad (139)$$

where  $i\hbar(d/dt)|\psi(t)\rangle = H(t)|\psi(t)\rangle$ . A more demanding control problem is that of unitary control, where we aim for the complete unitary evolution  $U(T)$  to approach a target transformation  $U_*$  (again up to a global phase). This will be the case, for instance, when the goal is to implement a quantum gate. In such cases, the success of a protocol can be measured by

$$\mathcal{F}_U = \frac{1}{d^2} |\text{Tr}(U_*^\dagger U(T))|^2, \quad (140)$$

where now  $i\hbar(d/dt)U(t) = H(t)U(t)$  and  $d$  is the Hilbert space dimension. In this section, we introduce the tools of *quantum optimal control* (QOC), which give a systematic approach to engineer  $H(t)$  such that the target is approximately achieved, i.e.,  $\mathcal{F} \simeq 1$ , where  $\mathcal{F}$  symbolizes the state or unitary fidelities, depending on the problem at hand. Note that, depending on the specific control problem, alternative measures of protocol success can be defined. For instance, one could be interested in minimizing the energy of a system, or in achieving a specific expectation value of an observable. Here, for concreteness, we will focus on the use of the fidelity as such a measure.

We start by considering the system Hamiltonian  $H(t)$  to be a function of a set of real *control fields*  $\mathbf{a}(t) = \{a_k(t)\}$ ,  $j = 1, \dots, K$ . For instance, we could consider the Hamiltonian to be written as

$$H(t) = H_0 + H_C(t) = H_0 + \sum_{k=1}^K a_k(t)H_{C,k}, \quad (141)$$

which makes explicit the existence of a drift (or free) Hamiltonian  $H_0$ , which we cannot manipulate directly, and the control Hamiltonian  $H_C(t)$ , which depends on time through the control fields  $\mathbf{a}(t)$  modulating the individual operators  $H_{C,k}$ . The properties of the drift and control Hamiltonians determine which states and transformations are in principle reachable to the system. Well-established methods to characterize this aspect exist [94,95], and use tools rooted in Lie group theory to determine the degree of controllability of a given model.

Here we take a pragmatic approach, and establish a systematic way to find the control functions  $\mathbf{a}(t)$  that minimize a *cost functional*

$$J[\mathbf{a}(t)] = 1 - \mathcal{F}[\mathbf{a}(t)], \quad (142)$$

where  $\mathcal{F}$  refers to either the state or unitary fidelity, and we have indicated the functional dependence of the fidelity explicitly. Note that here we follow the convention of defining a functional to be *minimized*; one could similarly define a functional to be maximized instead. We will focus on solving this problem via numerical optimization, for which we introduce a parametrization of the control fields in terms of real numbers

$$\boldsymbol{\alpha} = (\alpha_1, \alpha_2, \dots, \alpha_M) \rightarrow \mathbf{a}(t), \quad (143)$$

which serve as our control variables, and will allow us to treat the problem of minimizing Eq. (142) with standard numerical optimization tools. Note that, while we do not treat them in this tutorial, there are approaches in optimal control theory which rely on functional optimization (see Krotov's method [96,97] and the Pontryagin maximum principle [5]).

Generically, a numerical minimization routine involves initializing the controls with a guess  $\alpha_0$ , and then iteratively updating it following a set of rules that depend on the optimization algorithm. These can be broadly classified as gradient-based and gradient-free depending on whether they compute derivatives of the cost function  $J(\alpha)$  at various points to decide the most appropriate way to update the control field  $\alpha(t)$  at each iteration. The choice of parametrization in Eq. (143) is in principle arbitrary, but can be tailored to fit computational or experimental constraints. For instance, a Fourier parametrization would take the form

$$\begin{aligned} \alpha &= (A_1, \dots, A_{n_F}, B_1, \dots, B_{n_F}) \rightarrow a(t) \\ &= \sum_{m=1}^{n_F} [A_m \cos(\omega_m t) + B_m \sin(\omega_m t)], \end{aligned} \quad (144)$$

where  $n_F$  is the number of Fourier components and  $\omega_m$  could correspond to the harmonics of some base frequency, or be chosen randomly (as in the chopped randomized basis, CRAB, algorithm [98]). For concreteness, here we focus on parametrizing the field as a piecewise-constant function,

$$\begin{aligned} \alpha &= (\alpha_1, \alpha_2, \dots, \alpha_M) \rightarrow a(t) = \alpha_j \\ \text{if } (j-1)\Delta t \leq t < j\Delta t, \text{ where } j &= 1, \dots, M. \end{aligned} \quad (145)$$

where we have chosen a uniform time step  $\Delta t$  for convenience. This choice is conceptually simple and allows us to analytically compute the gradient of the cost function of Eq. (142) with respect to the control variables, which greatly improves the efficiency of the numerical search. We will work out this derivation explicitly in the next section.

The definition of the control Hamiltonian, e.g., Eq. (141), and parametrization, cf. Eq. (143), together with the choice of initial and target states ( $|\psi_0\rangle$  and  $|\psi_*\rangle$  for state control,  $\mathbb{1}$  and  $U_*$  for unitary control), completely determine the cost function  $J(\alpha)$ , which can be seen as an optimization *landscape*. Each point on the landscape corresponds to a different control field, with optimal protocols being global minima. The goal of any optimization algorithm is to traverse this landscape, starting from the initial guess  $\alpha_0$ , to eventually reach the global extremum via a series of updates of the control parameters  $\alpha_0 \rightarrow \alpha_1 \rightarrow \dots \rightarrow \alpha_{\text{opt}}$ . Here, we assume that such updates are done via a local search (instead of a global one), where only properties about the current point and its surrounding region informs the update rule. This process is depicted in Fig. 6. The properties of the landscape determine the hardness of finding a solution to the control problem. In a generic optimization problem, the landscape is rugged and filled with suboptimal minima or traps, as can be seen in Fig. 6(a), which often increases the computational expense of finding a near-optimal solution. In this scenario, it is

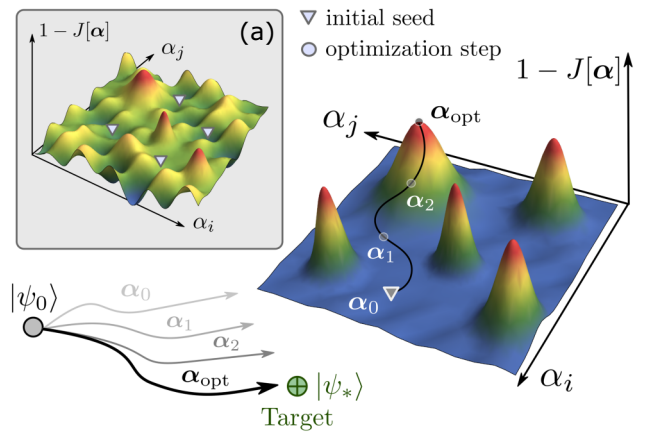


FIG. 6. Quantum control landscapes are determined by the dependence of the cost function  $J(\alpha)$  on the control variables  $\alpha$ . Main figure depicts how the optimal solutions, initial seeds, and optimization steps would appear in a two-dimensional landscape (note that real control problems typically have much more than  $M = 2$  dimensions). An optimization routine starts with an initial guess for the control variable  $\alpha_0$  which typically lands far from an extrema of the landscape, meaning that the evolution of the quantum state is far from the target configuration at the final time. The optimization consists in updating the control variables sequentially until an optimum is achieved  $\alpha_{\text{opt}}$ . In the presence of constraints, quantum control problems can have rugged landscapes with suboptimal extrema, as depicted in subfigure (a). In such cases, using many different initial seeds helps explore various parameter regions of the landscape.

typically convenient to balance the computational cost of the optimization between the exploration and the exploitation of the landscape. An optimization landscape can be explored by starting the search multiple times with different initial guesses (or “seeds”) which aim to cover different regions of parameter space. Each run is then exploited by allowing the search of a sufficient number of iterations to converge to a good solution. For quantum control problems, remarkably, it has been shown that optimization landscapes tend to have benign features [99]; in particular, provided there are no constraints to the control problem [100], quantum control landscapes can be shown to be generically devoid of traps in fully controllable systems (see Fig. 6). In practice, however, constraints on control problems (which we discuss below) will typically affect the topology of control landscapes. Interestingly, recent developments in quantum computing, particularly variational quantum algorithms, have revealed new properties of quantum control landscapes, like the phenomenon of barren plateaus, which we discuss briefly in Sec. VI K.

Often, constraints need to be taken into account when solving practical quantum control problems, and thus understanding their effect on the optimization search is an important task. Constraints can be broadly divided into two categories:

- (1) Constraints which are **built into** the control model. For instance, a fixed set of control Hamiltonians in the expansion of Eq. (141), a limited bandwidth set by the number of Fourier components in Eq. (144), or a limited evolution time  $T = M\Delta t$  in Eq. (145);
- (2) Constraints which are imposed **a posteriori** to the problem as additional terms to the cost function. For instance, bounding the amplitude of driving field to limit its energetic cost [97] or demanding the control processes to be robust to certain external perturbations [101,102].

We highlight here the role of one of these constraints: limiting the total evolution time  $T$ . Quantum mechanics imposes fundamental limitations to the rate at which one state can evolve into a different state. These are formalized as quantum speed limits (QSLs) [12], which provide lower bounds on the total evolution time, cf.  $T \geq \tau_{\text{QSL}}$ . For unitary dynamics, QSLs can be traced back to time-energy uncertainty relations, and take several forms [103–106]. In particular, we consider the bound derived from the work of Mandelstam and Tamm [107,108]:

$$T \geq \frac{\hbar}{\Delta E} \arccos(|\langle \psi_* | \psi_0 \rangle|) \equiv \tau_{\text{QSL}}, \quad (146)$$

where

$$\overline{\Delta E} = \frac{1}{T} \int_0^T dt \Delta E(t) \quad (147)$$

and  $\Delta E(t)$  is the variance of the Hamiltonian  $H(t)$ ,

$$\Delta E(t)^2 = \langle \psi(t) | H(t)^2 | \psi(t) \rangle - \langle \psi(t) | H(t) | \psi(t) \rangle^2, \quad (148)$$

with  $H(t)$  the total Hamiltonian, e.g., as defined in Eq. (141). The existence of QSLs fundamentally impacts quantum control problems as any optimization is destined to fail if one fixes the evolution time  $T$  to be below  $\tau_{\text{QSL}}$ . Furthermore, it is usually desirable to derive control processes which are as fast as possible since long evolution times allow for more sources of errors and decoherence to affect the dynamics. While in some cases optimal control solutions exist exactly at the fundamental bound  $T = \tau_{\text{QSL}}$  [109], typically constraints mean we can achieve solutions only at longer times  $\tau_{\text{QSL}}^* > \tau_{\text{QSL}}$  [95,110]. By definition, the time  $\tau_{\text{QSL}}^*$  is the minimum possible evolution time for which a given optimal control problem has a solution (i.e., one that achieves a fidelity  $\mathcal{F} = 1$ ). In Sec. III C 1 we will illustrate how to use optimal control methods to systematically search for the shortest evolution time  $\tau_{\text{QSL}}^*$  for simple control problems, and explore how constraining the evolution time affects the optimization landscape. We point out that geometric QSL bounds like Eq. (146), while universal, tend to become loose as the size of the

system increases (see Ref. [11] for an exception). This is because the energy variance in the denominator is typically an extensive quantity.

## B. Gradient-based optimal control using GRAPE

In this section we discuss in more detail how to approach an optimal control problem numerically. We focus on the simplest case of having a single control field  $a(t)$  parametrized with the piecewise constant ansatz of Eq. (145), leading to control variables grouped in a vector  $\alpha = \{\alpha_1, \alpha_2, \dots, \alpha_M\}$ . For a given choice of the number of steps  $M$  and total evolution time of  $T$ , the associated time step is  $\Delta t = T/M$ . This choice allows us to write the final-time evolution operator in closed form

$$U(T) = U_M U_{M-1} \dots U_2 U_1,$$

$$\text{where } U_j = \exp(-iH_j \Delta t) \text{ and } H_j = H(\alpha_j). \quad (149)$$

To search for a choice of  $\alpha$  that minimizes the cost function  $J(\alpha)$ , one can resort to a variety of numerical optimization algorithms. Because we expect the cost function to have a smooth dependence on the control variables in general, it is convenient to use gradient-based methods, which use information about the derivatives  $\partial J / \partial \alpha_k$  to inform the search at each step. Examples of these include gradient descent, Newton-Raphson and the Broyden-Fletcher-Goldfarb-Shannon (BFGS) algorithms. These are widely used optimization methods which are easily available in PYTHON, MATLAB, and other languages using standard packages (note that some quantum optimal control approaches also benefit from using gradient-free methods, like Nelder-Mead or Powell [46,98]). For these methods, the computation of the gradient can be done numerically (e.g., via finite-difference approximations, which are typically carried out automatically by these optimization routines) or it can be provided explicitly such that it only has to be numerically evaluated. The piecewise-constant parametrization of the control field we are using here actually allows us to compute the gradient of  $J(\alpha)$  analytically in a relatively straightforward way. The use of this parametrization together with the analytical form of the cost gradient in a quantum control problem constitutes the GRAPE (gradient-ascent pulse engineering) algorithm, first proposed by Khaneja *et al.* [111].

To compute the gradient, we first note that for both state and unitary control problems the cost function can be written as  $J(\alpha) = 1 - \mathcal{F}(\alpha) = 1 - |z(\alpha)|^2$  where

$$\text{state control: } z(\alpha) = \langle \psi_* | U(T) | \psi_0 \rangle \quad (150)$$

$$\text{unitary control: } z(\alpha) = \frac{1}{d} \text{Tr}(U_*^\dagger U(T)), \quad (151)$$

and thus

$$\frac{\partial J}{\partial \alpha_j} = \frac{\partial}{\partial \alpha_j} (z^* z) = -2 \operatorname{Re} \left\{ z^* \frac{\partial z}{\partial \alpha_j} \right\}. \quad (152)$$

As the control dependence is entirely contained within  $U(T)$ , to obtain  $\frac{\partial z}{\partial \alpha_j}$  we need to calculate

$$\frac{\partial U(T)}{\partial \alpha_j} = U_M U_{M-1} \dots \frac{\partial U_j}{\partial \alpha_j} U_{j-1} \dots U_1 \quad (153)$$

and thus we need to compute  $\partial U_j / \partial \alpha_j$ . Here, one should proceed with care because  $U_j = \exp(-iH_j \Delta t)$  does not necessarily commute with  $\partial H_j / \partial \alpha_j$ . In general, we have

$$\frac{\partial U_j}{\partial \alpha_j} = \frac{\partial}{\partial \alpha_j} \sum_{k=0}^{\infty} \frac{(-i)^k}{k!} \Delta t^k H_j^k \quad (154)$$

$$= \sum_{k=0}^{\infty} \frac{(-i)^k}{k!} \Delta t^k \sum_{i=1}^k H_j^{i-1} \frac{\partial H_j}{\partial \alpha_j} H_j^{k-i}. \quad (155)$$

In the limit where the time step  $\Delta t$  is small, and thus the piecewise constant field resembles a continuous function, the expression in Eq. (155) can be simplified, see for instance Refs. [17,111]. For that limit to be relevant, however, we need to deal with a large number  $M$  of control variables, which can be undesirable. Here we focus on obtaining an exact closed form for  $\partial U_j / \partial \alpha_j$ , valid for arbitrary time step sizes  $\Delta t$  [112,113]. For this, we introduce the spectral decomposition of  $H_j = \sum_{l=1}^d e_l |l\rangle \langle l|$  (note that both the eigenvalues and eigenvectors will be different for each  $j$ , i.e., for each time step). Inserting this into Eq. (155), we obtain

$$\begin{aligned} \frac{\partial U_j}{\partial \alpha_j} &= \sum_{k=0}^{\infty} \frac{(-i)^k}{k!} \Delta t^k \sum_{i=1}^k \sum_{l,m} e_l^{i-1} e_m^{k-i} |l\rangle \langle m| \left\langle l \left| \frac{\partial H_j}{\partial \alpha_j} \right| m \right\rangle \\ &= \sum_{l,m} |l\rangle \langle m| \left\langle l \left| \frac{\partial H_j}{\partial \alpha_j} \right| m \right\rangle \sum_{k=0}^{\infty} \frac{(-i)^k}{k!} \Delta t^k \sum_{i=1}^k e_m^{k-1} (e_l/e_m)^{i-1}. \end{aligned} \quad (156)$$

The infinite series appearing in the above expression can be evaluated exactly. For the case  $e_l = e_m$ , we obtain

$$\begin{aligned} \sum_{k=0}^{\infty} \frac{(-i)^k}{k!} \Delta t^k k e_m^{k-1} &= -i \Delta t \sum_{k=0}^{\infty} \frac{(-i)^{k-1}}{(k-1)!} \Delta t^{k-1} e_m^{k-1} \\ &= -i \Delta t \exp\{-ie_m \Delta t\}. \end{aligned} \quad (157)$$

If  $e_l \neq e_m$ , we can proceed as follows:

$$\begin{aligned} &\sum_{k=0}^{\infty} \frac{(-i)^k}{k!} \Delta t^k e_m^{k-1} \sum_{i=1}^k (e_l/e_m)^{i-1} \\ &= \sum_{k=0}^{\infty} \frac{(-i)^k}{k!} \Delta t^k \frac{e_m^k}{e_m} \sum_{n=0}^{k-1} (e_l/e_m)^n \\ &= \sum_{k=0}^{\infty} \frac{(-i)^k}{k!} \Delta t^k \frac{e_l^k - e_m^k}{e_l - e_m} \\ &= \frac{1}{e_l - e_m} (e^{-ie_l \Delta t} - e^{-ie_m \Delta t}). \end{aligned} \quad (158)$$

Combining expressions (156), (157), and (158) yields an analytical expression for all the matrix elements of  $\partial U_j / \partial \alpha_j$ . Notice that this requires us to diagonalize the operator  $H_j$  for each  $\alpha_j$ , which will need to be done numerically in most cases. For a single qubit, the eigenvectors and eigenvalues of  $H_j$  have a analytic form for an arbitrary  $\alpha_j$ . Here we do not use this form to illustrate the more general procedure, which applies to more complex systems. Once the factor from Eq. (156) is computed, the gradient can be constructed from Eq. (153),

$$\text{state control: } \frac{\partial z}{\partial \alpha_j} = \langle \psi_* | U_{j+1,M} \frac{\partial U_j}{\partial \alpha_j} U_{1,j-1} | \psi_0 \rangle \quad (159)$$

$$\text{unitary control: } \frac{\partial z}{\partial \alpha_j} = \frac{1}{d} \operatorname{Tr} \left( U_*^\dagger U_{j+1,M} \frac{\partial U_j}{\partial \alpha_j} U_{1,j-1} \right), \quad (160)$$

where we introduced the notation  $U_{k,l} = U_l U_{l-1} \dots U_{k+1} U_k$ . Note that in the case of state control, the cost function gradient at step  $j$  depends on the initial state  $|\psi_0\rangle$  forward-evolved until time step  $j-1$ , and on the target state  $|\psi_*\rangle$  backwards-evolved until time step  $j+1$ .

We can now construct a basic optimal control routine by following these steps:

- (1) Set the final evolution time  $T$  and the number of time steps  $M$  (which defines the number of control variables). This suffices to define a discretized time variable  $t \rightarrow \{t_k = k\Delta t\}$ , with  $k = 0, 1, \dots, M$ , and  $M\Delta t = T$ .
- (2) Choose an initial guess or ansatz  $\alpha_0$  for the control variable.
- (3) Implement a numerical minimization routine. Typical inputs for such a routine are as follows:
  - (a) The cost function  $J(\alpha)$ .
  - (b) The particular minimization method to be used (BFGS, gradient descent, ...)

- (c) The associated gradient function  $g(\boldsymbol{\alpha}) = (\partial/\partial\boldsymbol{\alpha}) J(\boldsymbol{\alpha})$ , Eq. (152).
- (d) The initial guess  $\boldsymbol{\alpha}_0$ , defined in step 2.
- (e) Bounds on the control variables such that the search is restricted to  $\alpha_{\min} \leq \alpha_j \leq \alpha_{\max}$  (optional)
- (f) Numerical tolerances or thresholds on the cost and/or the gradient, which determine at what level of precision the search is allowed to stop.

In the remainder of this section, we consider two examples in which we illustrate this procedure for both state and unitary control in a two-level system.

### C. Examples

#### 1. Example: optimal control in the Landau-Zener model

We begin by considering control of the Landau-Zener Hamiltonian, already introduced in Sec. II F 1, which has the form

$$H(t) = H_0 + v(t)H_C = \Delta\sigma^x + v(t)\sigma^z. \quad (161)$$

We also recall here the form of the ground state of  $H(t)$  as a function of the instantaneous value of  $v(t)$ :

$$|\phi_g(v)\rangle \equiv \cos\left(\frac{\theta}{2}\right)|0\rangle + \sin\left(\frac{\theta}{2}\right)|1\rangle, \quad \text{with } \tan\theta = \frac{\Delta}{v}. \quad (162)$$

For this model we are interested in the problem of driving the system from the initial state

$$|\psi_0\rangle = |\phi_g(-v_0)\rangle \quad (163)$$

to the target state

$$|\psi_*\rangle = |\phi_g(+v_0)\rangle \quad (164)$$

for some  $v_0 > 0$ . Note that, as  $v_0 \rightarrow \infty$ , we have  $\theta \rightarrow 0$  and the desired evolution is the one that connects  $|\psi_0\rangle = |0\rangle$  to  $|\psi_*\rangle = |1\rangle$ . We consider the field  $v(t)$  to be piecewise-constant with  $M = 10$  steps and denote its values  $\{\alpha_1, \dots, \alpha_M\}$  to make notation consistent with the previous section. Note that this example obeys the decomposition of Eq. (141). Even in this very simple case, one finds that  $H_j = H_0 + \alpha_j H_C$  and  $\partial H_j / \partial \alpha_j = H_C = \sigma^z$  do not commute, which motivates the need for the calculation of the gradient shown in Sec. III B.

To test the optimization procedure of steps 1–3 above, we start by fixing  $T = \pi/\Delta$ , and using a limited-memory variant of the BFGS method (L-BFGS-B) as our optimization routine. We also restrict the search to  $|\alpha_j| \leq c = 2\Delta$  to avoid control fields with large amplitudes. With

these specifications, we run the optimal control routine for two choices of the parameter  $v_0$  determining the initial and target states ( $v_0 = 1$  and  $v_0 = 5$ ), and two choices of initial ansatz; one where the  $\alpha_j$ 's are sampled from a uniform distribution between  $[-1, 1]$ , and another where we choose  $\alpha_j = 0$  for all  $j = 1, \dots, M$ . Results are shown in Fig. 7, where we display the initial guess for the control field (gray dashed lines) and the optimal field obtained with the optimization routine (colored full lines). In all cases we observe that the cost function is initially of order  $J(\boldsymbol{\alpha}_0) \in (0.5, 1)$ , and that the optimization is able to achieve  $J(\boldsymbol{\alpha}_{\text{opt}}) \in (10^{-16}, 10^{-13})$  in any standard run. Comparison of Figs. 7(a)–7(b) and 7(c)–7(d) reveals that for a given control problem (i.e., for a fixed  $v_0$ ), different initial guesses perform similarly. The comparison also illustrates two generic facts of quantum optimal control. When the system is controllable and given enough control resources, control problems typically have multiple possible solutions which are (roughly) equally good. Even for fixed, finite values of  $T$  and  $M$  (which effectively impose constraints on the control problem) these multitude of solutions can be accessed by exploration of the control landscape with different initial guesses  $\boldsymbol{\alpha}_0$  (see illustration in Fig. 6). In addition, the optimal fields often inherit properties of the initial guesses. This is because in presence of multiple minima, a local optimization routine like the one considered here will lead to a solution close to the initial point. In our example, we see that the optimal fields obtained following the random ansatz look random themselves, while the optimal fields obtained from the constant initial guess are reveal to be structured and symmetric. For simple systems like the one studied here, these features might be easily traced back to symmetries of the model [114] but, in more complex cases, the structure of the optimal field can be useful to identify non-trivial properties of the system [115–117].

*Identifying quantum speed limits using optimal control.*—Following on from previous discussions on quantum speed limits, it is natural to ask what is the minimum evolution time  $\tau_{\text{QSL}}^*$  required to achieve a good solution for this control problem. To explore this question, we set out to solve the optimization problem for different values of the total evolution time  $\{T_1, T_2, \dots, T_R\}$ , arranged in increasing order. In principle, each choice of  $T$  leads to an independent optimization procedure, given by the steps 1–3 above. Here we describe a slightly more elaborate approach, which has been routinely used in previous studies [8, 109, 118–123]. We start by solving the optimal control problem for the largest value of evolution time  $T_R$  with a given choice of initial ansatz  $\boldsymbol{\alpha}_0^{(R)}$  (say, a uniformly random one). If  $T_R > \tau_{\text{QSL}}^*$ , then we expect to be able to obtain an optimal field  $\boldsymbol{\alpha}_{\text{opt}}^{(R)}$  yielding a very small cost function. We then move on to solve the optimization for  $T_{R-1}$ , but now we choose as an initial guess the optimal field

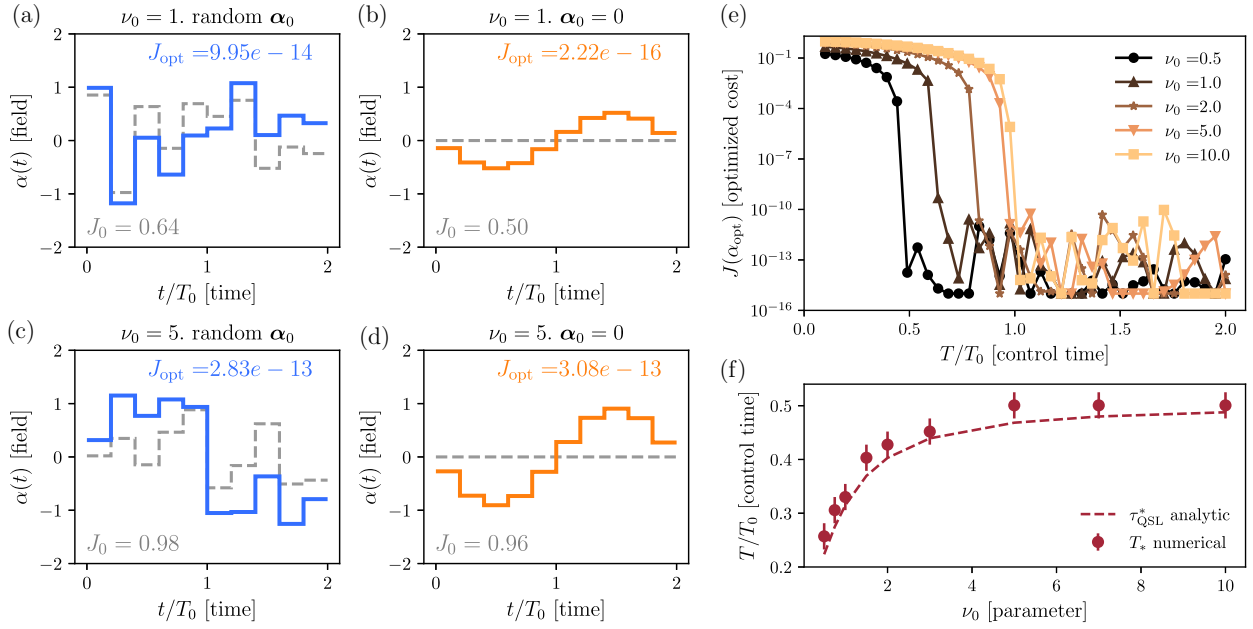


FIG. 7. Quantum optimal state control for the Landau-Zener model. (a)–(d) Show various control fields related to realizing a state transfer process between ground states of the Landau-Zener Hamiltonian, cf. Eq. (161). Shown are the initial guesses (gray dashed lines) and optimized fields (thick colored lines) for four different cases. (a) Initial state given by  $\nu_0 = 1$ , and initial guess for control field  $\alpha_0$  chosen randomly. (b)  $\nu_0 = 1$  and  $\alpha_0 = 0$ . (c)  $\nu_0 = 5$  and  $\alpha_0$  random. (d)  $\nu_0 = 5$  and  $\alpha_0 = 0$ . In all cases the initial  $J_0 = J(\alpha_0)$  and optimized  $J_{\text{opt}} \equiv J(\alpha_{\text{opt}})$  values of the cost function are displayed. Parameters used  $M = 10$ ,  $T = 2T_0$  where  $T_0 = \pi/(2\Delta)$ . (e) Optimized cost functions achieved by the optimization as a function of total evolution time  $T$ , for various choices of the parameter  $\nu_0$  determining the initial and final state. (f) Minimum control time for the Landau-Zener Hamiltonian, as a function of  $\nu_0$ . Circles indicate the numerical estimates obtained from the data shown in (e). Full line corresponds to the analytical solution derived in Ref. [110].

from the previous run, i.e.,

$$\begin{aligned} \text{initial guess for } T = T_{j-1} &\rightarrow \alpha_0^{(j-1)} \\ &= \alpha_{\text{opt}}^{(j)} \leftarrow \text{optimal field from } T = T_j. \end{aligned} \quad (165)$$

When repeating this procedure for every choice of  $T$  in our list, we are “helping” each new optimization by seeding a field which we expect to lead to a small cost (as long as  $T_j - T_{j-1}$  is sufficiently small). In Fig. 7(e) we plot the optimized cost as a function of the evolution time  $T$  for this procedure, for various choices of  $\nu_0$  (i.e., for various choices of initial and target states). Analyzing the curves starting from the right (large values of  $T$ ), in all cases we observe nearly optimal fidelities, corresponding to costs of orders between  $10^{-12}$  and  $10^{-15}$ . As the allowed evolution time decreases, the optimization is able to find good solutions up until a critical time, after which the optimized cost grows rapidly and takes values  $\sim 10^{-1}$ . This critical time  $T_*$  divides the controllable regime where  $T > T_*$  and the uncontrollable regime where  $T < T_*$ . Because  $T_*$  depends on details of the optimization procedure (for instance, on the number of time steps  $M$ ), it can only provide an upper bound to the true minimum control time  $\tau_{\text{QSL}}^*$ . Nevertheless, in many situations of interest, we can take  $T_*$  as a good approximation of  $\tau_{\text{QSL}}^*$ . For the problem treated here,

the actual optimal time is known analytically [110,124]. In Fig. 7(f) we compare the numerically obtained  $T_*$  and the analytical  $\tau_{\text{QSL}}^*$  as a function of  $\nu_0$ . The results indicate that the numerical optimal control procedure is able to obtain good estimates of the fundamental optimal control time for this problem. In particular, we see how as  $\nu_0 \rightarrow \infty$  and thus  $|\psi_0\rangle \rightarrow |0\rangle$  and  $|\psi_*\rangle \rightarrow |1\rangle$ , the control time reaches  $\tau_{\text{QSL}}^* = \pi/2\Delta$ . This is the time required by  $H = \Delta\sigma_x$  to rotate  $|0\rangle$  to  $|1\rangle$ .

## 2. Example: optimal control of a single-qubit gate

Here we consider a different model of a two-level system where we aim at generating unitary transformations, or gates, using optimal control. To this effect, we will study the Hamiltonian

$$H(t) = \frac{\Omega}{2} (\cos(\alpha(t))\sigma_x + \sin(\alpha(t))\sigma_y). \quad (166)$$

This model corresponds to a single spin-1/2 particle being driven by a field with constant amplitude  $\Omega$  (i.e., Rabi frequency) but variable direction in the  $x$ - $y$  plane, determined by  $\alpha(t)$ . When compared to the Landau-Zener Hamiltonian, cf. Eq. (161), this model has two important features. First, the energy of the system is bounded by  $\Omega/2$  irrespective of the control field. Thus, increasing the amplitude or

frequency of the field will not necessarily make the system evolve faster in time. Second, it does not follow the decomposition of Eq. (141). In particular, one finds that  $\partial H_j/\partial \alpha_j$  is now a function of  $\alpha_j$  due to the nonlinear dependence on the control field. Nonetheless, this case is naturally contemplated in the general calculation of the gradient of Sec. III B.

In a unitary control problem, we need to define the target transformation. We consider two families of target gates,

$$U_x(\phi) = \exp(-i\sigma_x\phi/2), \quad (167)$$

$$U_z(\phi) = \exp(-i\sigma_z\phi/2). \quad (168)$$

Note that for the targets  $U_x(\phi)$ , the problem has a straightforward solution, obtained by setting  $\alpha(t) = 0$  and evolving the system for a time  $T = \phi/\Omega$ . On the other hand, for the case  $U_z(\phi)$  the control field needs to play a non-trivial role to generate the desired gate.

We analyze this problem numerically by choosing the field  $\alpha(t)$  to be piecewise-constant with  $M = 10$ , as in the previous example. We start considering an evolution time of  $T = 2\pi/\Omega$  and two instances of the target gates  $U_z(\phi)$  with  $\phi = \pi/2$  and  $\phi = \pi$ . To illustrate the results of the search, we again consider two different types of initial guess: control parameters randomly sampled from a

uniform distribution in  $[-\pi, \pi]$ , or increasing linearly in time from  $-\pi$  at  $t = 0$  to  $\pi$  at  $t = T$ . Results are shown in Fig. 8. Similarly to the results displayed in the previous section, we observe how the optimization routine is able to decrease the value of the cost by over 10 orders of magnitude and achieve very high fidelities. We also observe how different initializations of the search lead to different optimal fields, which inherit properties of the initial seed.

Finally, we can use our optimization procedure to explore the quantum speed limit to achieve the indicated target gates. We do this in exactly the same way as described for the state control problem in the previous section. In Fig. 8(e) we show the optimized cost values as a function of evolution time  $T$  for the gates  $U_z(\phi)$  and various choices of  $\phi$ . We observe a clear separation between the controllable regime for  $T > \tau_{\text{QSL}}^*$  where very low cost values are achievable, and the uncontrollable regime where  $T < \tau_{\text{QSL}}^*$  and the optimal cost is highly constrained. The critical evolution time  $T_*$  increases with  $\phi$ , which is intuitive as gates with larger  $\phi$  are further away from the identity, i.e., the initial state for unitary control. A deeper analysis of this feature is presented in Fig. 8(f), revealing a square-root-like growth of  $T_*$  with  $\phi$ . This is consistent with the actual minimum evolution time  $\tau_{\text{QSL}}^{*(z)}$  which for

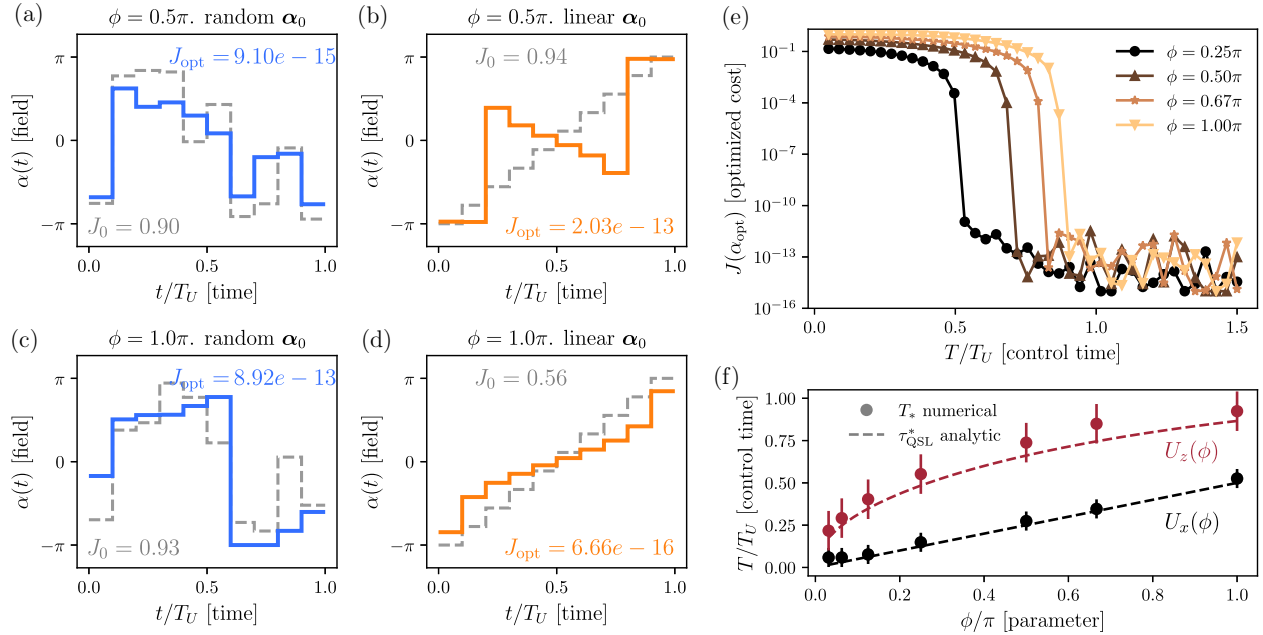


FIG. 8. Quantum optimal control for a single-qubit gate. (a)–(d) Show control fields related to implementing the target gates  $U_z(\phi)$  in Eq. (168) using the control Hamiltonian in Eq. (166). We display initial guesses (gray dashed lines) and optimized fields (thick colored lines) for four different cases. (a) Target gate given by  $\phi = \pi/2$ , and initial guess for control field  $\alpha_0$  chosen randomly. (b)  $\phi = \pi/2$  and  $\alpha_0$  linear. (c)  $\phi = \pi/2$  and  $\alpha_0$  random. (d)  $\phi = \pi$  and  $\alpha_0$  linear. In all cases the initial  $J_0 \equiv J(\alpha_0)$  and optimized  $J_{\text{opt}} \equiv J(\alpha_{\text{opt}})$  values of the cost function are displayed. Parameters used  $M = 10$ ,  $T = T_U$ , where  $T_U = 2\pi/\Omega$ . (e) Analysis of optimal cost functions achieved by the optimization as a function of total evolution time  $T$ , for various choices of the target  $U_z(\phi)$ . (f) Minimum control time as a function of  $\phi$  for the two families of target gates  $U_x(\phi)$  and  $U_z(\phi)$ . Circles indicate the numerical estimates obtained from the optimal control procedure. Dashed lines correspond to analytical solutions. For  $U_x$ ,  $\tau_{\text{QSL}}^*/T_U = \phi/(2\pi)$ . For  $U_z$  the optimal evolution time has been derived in Ref. [125].

this problem has been analytically obtained in Ref. [125]. For completeness, we include results corresponding to the family of gates  $U_x(\phi)$  for which, as mentioned before, optimization was not really necessary. Nonetheless, we verify with numerical optimization that it is not possible to break the minimum evolution time given by  $\tau_{\text{QSL}}^{*(x)} = \phi/\Omega$ , which scales as  $\phi$  instead of  $\sqrt{\phi}$ . Similarly to what happens with state control problems and the quantum speed limit, lower bounds to the minimum control times can be obtained via geometric arguments [95].

### 3. Example: generation of many-body entanglement

The examples shown in Secs. III C 1 and III C 2 show the application of QOC to the simplest model of a two-level system or qubit. One obvious extension of these cases is the generation of two-qubit gates, which has been thoroughly studied for various platforms, including superconducting qubits, trapped ions, and neutral atoms. Nevertheless, the formalism described in Sec. III B is general and can be directly applied to more complex scenarios.

Here we consider the problem of many-body state preparation, which is closely related to the control tasks treated in Sec. IV with the STA formalism. While the direct application of QOC to many-body systems is naturally limited by the computational complexity of numerically simulating these systems, we will show that even for modest system sizes QOC can provide interesting insights into the resources needed to perform control on a high-dimensional Hilbert space. This, in turn, will allow us to address a key point raised in Sec. I: the use of quantum control tools to bridge the gap between implementability and scalability of quantum operations.

We consider a model of  $N$  qubits undergoing global control and all-to-all interactions according to the Hamiltonian

$$H(t) = \frac{\Omega_x(t)}{2} \sum_{i=1}^N \sigma_i^x + \frac{\Omega_y(t)}{2} \sum_{i=1}^N \sigma_i^y + \frac{\beta}{4N} \sum_{i,j=1}^N \sigma_i^z \sigma_j^z. \quad (169)$$

Note that, as opposed to the single-qubit example, here the model has a drift term which is time-independent (thus, not directly controllable), which corresponds to the interaction between the qubits. The relevant time-scale of the problem is set by  $T_\beta = 2\pi/\beta$ . We have also chosen to normalize the interaction term to ensure that the Hamiltonian remains extensive as  $N$  increases.

We consider the following control task. Starting from the state

$$|\psi_0\rangle = |0\dots 0\rangle, \quad (170)$$

we want our system to reach an entangled state  $|D_k^{(N)}\rangle$  characterized by the symmetric superposition of all possible  $k$

bit-flips. For instance, if  $N = 4$  some of the target states we are interested in are

$$|D_1^{(N=4)}\rangle = \frac{1}{2} (|1000\rangle + |0100\rangle + |0010\rangle + |0001\rangle) \quad (171)$$

$$|D_2^{(N=4)}\rangle = \frac{1}{\sqrt{6}} (|1100\rangle + |1010\rangle + |1001\rangle + |0110\rangle + |0101\rangle + |0011\rangle). \quad (172)$$

These types of states display multipartite entanglement, and are regarded as a resource in quantum sensing [126] and cavity quantum optics [127], among others.

To tackle this problem, we first note that the Hamiltonian in Eq. (169) is invariant under any permutation of the qubits. This allows us to write this operator elegantly in terms of collective spin operators

$$J_\alpha = \frac{1}{2} \sum_{i=1}^N \sigma_i^\alpha, \quad (173)$$

leading to

$$H(t) = \Omega_x(t)J_x + \Omega_y(t)J_y + \frac{\beta}{N}J_z^2. \quad (174)$$

In this form, it also becomes clear that at all times  $[H(t), \mathbf{J}^2] = 0$ , where  $\mathbf{J}^2 = J_x^2 + J_y^2 + J_z^2$  is the total spin of the system. In addition, the set of initial and target states all correspond to joint eigenstates of  $J_z$  and  $\mathbf{J}^2$  (note that  $|\psi_0\rangle = |D_0^{(N)}\rangle$ ). These are also called Dicke states, which obey [128]

$$J_z |D_k^{(N)}\rangle = \left(\frac{N}{2} - k\right) |D_k^{(N)}\rangle \quad (175)$$

$$\mathbf{J}^2 |D_k^{(N)}\rangle = \frac{N}{2} \left(\frac{N}{2} + 1\right) |D_k^{(N)}\rangle. \quad (176)$$

Thus, our control problem takes place in a reduced subspace spanned by these states, which has a dimension  $N + 1$ , as the value of excitations  $k$  goes from 0 to  $N$ . This reduced subspace is actually exactly equivalent to the Hilbert space of a single spin- $J$  particle with  $J = N/2$ . In fact, the Hamiltonian in Eq. (174) accurately models the dynamics of electronic states in individual atoms, where the  $J_z^2$  now corresponds to a non-linear light-shift induced by non-resonant optical driving [129]. These systems are now routinely being considered as qudits for quantum computing platforms, and the results shown here directly carry over to that application as well [130] (see also Sec. VIK for further information).

With this model we can explore how control resources change as the complexity of the control task increases,




FIG. 9. Quantum optimal control for many-body entangled state generation. Panels show the optimized cost obtained from QOC on the protocol described in the main text, as a function of the allotted control time  $T$ . In all cases, both fields  $\Omega_x(t)$  and  $\Omega_y(t)$  are discretized with  $M = 15$  steps and constrained to maximum amplitude of  $\Omega_{\max} = 3\beta$ .

both in terms of increasing the number of target excitations  $k$  and the number of total particles  $N$ . To do this, we run the QOC procedure outlined in Sec. III B to this problem for  $N = 6, 8, 10$  particles, and in each case to reach a target state with  $k = 1, 2, \dots, N/2$  excitations. We note that the case  $k > N/2$  turns out to be equivalent to the cases considered here, up to global rotations.

Similar to the analysis shown in, e.g., Fig. 7(e), we run the optimization for various choices of the total evolution time, in order to identify an approximate quantum speed limit required to produce these various transformations. Results are shown in Fig. 9. For each value of  $N$  considered, we find that the time required for the control to find high-fidelity solutions does not depend significantly on the number of excitations  $k$  in the target state. This indicates that even when states with higher  $k$  require more bit-flips, the optimal solution is able to efficiently navigate Hilbert space and takes a shortcut to generate the target state directly. In turn, by comparing the curves for different values of  $N$ , we do observe that the total amount of time increases with system size in a seemingly linear fashion. This is expected given the linear increase of accessible Hilbert space. We note that this is also a consequence of the choice of normalization in Eq. (169).

#### D. Experimental implementations of quantum optimal control

Quantum optimal control tools have been successfully applied to various experimental platforms involving manipulation of quantum systems. Interestingly, the flexibility of QOC has led each community to adapt these tools to tackle and mitigate errors and limitations which are native to each implementation. Following the discussion in Sec. III A, this translates into imposing different constraints to the base QOC problem. Examples of this include leaking reduction in superconducting qubits [131] and inhomogeneous broadening mitigation in atomic and nuclear magnetic resonance (NMR) systems [132].

Here we name some examples of successful implementations of QOC in different platforms, in a manifestly non-exhaustive way; the reader is referred to Reviews such as Refs. [1,13] for further examples. In atomic systems, optimal control has been used to coherently control large manifolds of electronic states in individual neutral atoms. This includes both Zeeman [133–135] and Stark states [136]. In these cases both radiofrequency and microwave fields are modulated in time to achieve the desired control. Demonstrated tasks include the generation of circular Rydberg and non-classical superposition states, as well the implementation of arbitrary unitary operations within the corresponding electronic state manifold.

Likewise, QOC has become a standard tool to design entangling gates between neutral atoms for quantum computing. In this case, time-modulation of laser intensity and detuning can be optimized to produce operations which are highly robust to imperfections like light-shift broadening or spontaneous decay. Recent demonstrations which leverage QOC tools include Refs. [137–139]. A similar approach has been applied to trapped-ion quantum processors, with entangling gates enabled by optimal design of frequency comb structure [140]. These systems have also benefited from QOC in the task of transporting individual ions along a trap [141].

Superconducting circuits have also successfully implemented optimal control tools for various tasks. In these systems, control is exerted by microwave signals; recently demonstrated tasks include the implementation of a universal gate set on a logical cat qubit [142] and high-fidelity qudit gates [143]. Since superconducting qubits are two-levels embedded in a larger Hilbert space, it is crucial to minimize population leaking outside of the relevant computational subspace. The DRAG (derivative removal by adiabatic gate) method is a QOC approach that has been proposed to address this issue [131], and which has been experimentally demonstrated in Ref. [144]. Following its demonstrated success and flexibility, quantum optimal control is now becoming an industrial standard

for quantum computing, with many quantum hardware and software providers utilizing quantum control tools in various layers of the computing stack for tasks like circuit compilation, gate synthesis, state preparation and readout, and error mitigation.

Finally, we briefly mention that quantum optimal control tools have a long history beyond the quantum information processing platforms mentioned here. In particular, optimization is a standard tool in composite pulse design in NMR experiments [145]. Also, the control of chemical reactions using shaped laser pulses was one of the first applications of quantum optimal control theory, see for instance Ref. [146].

#### IV. REINFORCEMENT LEARNING FOR OPTIMAL QUANTUM CONTROL

In the previous sections, we introduced the concepts of counterdiabatic driving and optimal control. Here we take on a different approach to quantum control using reinforcement learning (RL). We also discuss how ideas from CD driving can be used in combination with RL to control quantum systems, leading to powerful hybrid approaches. While the previous sections would be useful to grasp the overarching theme, the material below can be considered a self-consistent standalone tutorial on reinforcement learning for quantum control. We begin with a short overview of reinforcement learning, defining the basic notions and concepts, and elaborating on the details behind a simple RL algorithm, cf. Sec. IV A. We then move on to discuss three examples that illustrate some of the features of RL algorithms applied to quantum control: in particular, in Sec. IV B 1, we discuss how to train an RL agent to learn how to prepare a fixed target starting from an arbitrary initial state. Section IV B 2 combines ideas from shortcuts to adiabaticity with RL that give rise to hybrid RL-OC algorithms. Finally, in Sec. IV B 3 we discuss how to extract

TABLE IV. Summary of notation for the most common quantities used in RL.

Object	Proposed symbol
States and state space	$s \in \mathbf{S}$
Actions and action space	$a \in \mathbf{A}$
Rewards and reward space	$r \in \mathbf{R}$
Trajectory	$\tau$
Expected return	$\mathbf{G}(\tau)$
Episode length (number of time steps)	$T$
Transition probability	$p(s' s, a)$
Policy (strategy)	$\pi(a s)$
$Q$ function	$Q_\pi(s, a)$
Probability for a trajectory	$P_\pi(\tau)$
RL objective	$\mathbf{J}$
Variational (learning) parameters	$\theta$
Monte Carlo samples	$N$

partial information about the state of a controlled qubit with the help of an ancilla to train an RL agent using feedback control.

Modern machine learning (ML) stands on three pillars [147–150]: (i) *Supervised learning* deals with labeled data, and aims at finding the most probable label for new, unseen datapoints. Notable examples include image recognition or classification tasks based on linear and logistic regression. (ii) The goal of *unsupervised learning*, on the other hand, is to produce an approximate model for the probability distribution that generates a given dataset; once such a model is trained, we can use it to produce more data with the desired properties. This includes, for instance, generative models that can compose music, draw paintings, or write text. (iii) Last, in *reinforcement learning* [151], an agent learns to execute a given task by interacting with its environment. Note that in RL the term “environment” refers to the physical system the agent interacts with [see Fig. 10]; this is in contrast to the notion of environment in open systems. RL is designed for problems that require an interactive feedback loop. The main objective is to learn a strategy (a.k.a. policy) that, based on observations of the environment, helps the agent manipulate or steer the environment. Reinforcement learning and optimal control share many common features and can be thought of as the two sides of the same coin [152]. Therefore, when leveraging ML for quantum control [153,154], RL appears as a natural choice for many (though certainly not all) quantum control setups.

In RL, the environment is initialized in a certain state  $s_0$ , and the agent is presented with a set of actions  $\mathbf{A}$  to choose from. Depending on the action  $a_t$  selected, the environment reacts and changes state:  $s_t \rightarrow s_{t+1}$ . After each such step  $t$ , the agent observes the new state of the environment, and is

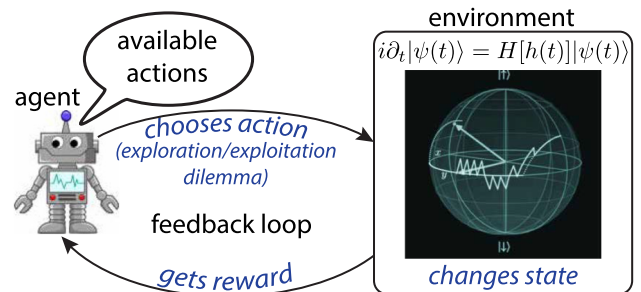


FIG. 10. RL in a nutshell: the agent observes its environment in a given state, and is presented with a choice of actions to pick from, each of which changes the state of the environment. As a consequence of its decision, the agent is given a reward, depending on the new environment state. Similar to animals being rewarded for achieving a task during training, the goal in RL is to maximize the long-term expected return. This is accomplished through learning about those properties of the environment that are particularly relevant for maximizing the cumulative reward. During the training process, RL agents autonomously figure out an optimal policy to perform the task.

given a reward  $r_{t+1}$  that quantifies the quality of the action taken with respect to the objective of the task. Based on this information, the agent then makes an educated guess to select the next action, and so forth, cf. Fig. 10. The goal of the agent is to bring the environment into a desired state. To do so, its objective is to maximize the cumulative expected reward  $\mathbf{G}$ , called return. The expectation is taken over the policy (i.e., the strategy) of the agent and the environment dynamics if the latter is stochastic. Such optimization in expectation brings out as a clear advantage of RL the ability to learn in uncertain or stochastic environments. The learning process in RL is often divided into episodes, which typically comprise a fixed number of steps; once the episode comes to an end, the environment is reset to its initial state  $s_0$  and the agent starts over again. Importantly, however, the agent keeps building on the experience gained in previous episodes.

To select an action, the agent uses a probability distribution  $\pi$ , called a policy (or strategy). Given a state of the environment, the policy  $\pi(a|s)$  defines the probability of taking each of the available actions. If the policy is a delta-function over the actions, we call it deterministic; such a case corresponds to selecting specific actions with certainty, e.g., when applying a fixed external control protocol to the environment. One issue with deterministic policies is that they always produce the same output, and hence the RL agent cannot cause the environment to behave differently, to learn from. To explore new actions (and ultimately determine which one is optimal for the given state), the policy in RL is typically non-deterministic: for instance, if the action space is discrete, one can think of the policy as a normalized state-dependent histogram (distribution) over the available actions; for continuous action spaces, the policy is a continuous probability density, such as a Gaussian or a Lorentzian. However, if the agent explores too much (i.e., its policy is close to a uniform distribution), it cannot learn a better policy either. This tension between exploring new actions and exploiting the already gathered knowledge is known as the “exploration-exploitation dilemma.”

RL algorithms iteratively improve the policy until it becomes (close to) optimal. Optimal policies, by definition, maximize the total expected return  $\mathbf{G}$ . In an alternative formulation of the learning procedure, we can look at the state of the environment, and try to predict what the expected return from this state onward (following a fixed policy) should be. This leads us to the concept of action-value (or  $Q$ -) functions  $Q_\pi(s, a)$ ; they assign an expected return value to each state-action pair. Once the  $Q$ -function is known, one can extract a policy from it by, e.g., looking up the values of different actions for a fixed state, and taking the action that maximizes the action-value function.

Reinforcement learning is not a single algorithm; instead, it should be viewed as a framework that comprises a collection of optimization algorithms, broadly divided

into two classes: In policy gradient methods one approximates the policy  $\pi \approx \pi_\theta$  with a set of variational parameters  $\theta$ . For instance,  $\theta$  can be the weights and biases of a deep neural network, or  $\theta = (\mu, \sigma^2)$  contain the mean and the variance of a Gaussian distribution in the case of continuous actions [151], but other variational ansätze (e.g., a matrix product state or a tensor network more generally [155,156]) might be better suited to the particular problem at hand. The parameters  $\theta$  are then updated using the gradient of the total expected return, estimated, e.g., using Monte Carlo sampling. Examples of widely used algorithms include policy gradient (PG) [157], proximal policy optimization (PPO) [158], trust-region policy optimization (TRPO) [159], and actor-critic (AC) methods [160]. Alternatively, value-function methods allow us to encode information about the policy; they aim at improving the  $Q$ -function itself which allows us to indirectly improve the policy [151]. Such methods include, among others,  $Q$ -learning [161] and its deep-learning version—DQN [162], Deep deterministic policy gradient (DDPG) [163], etc. Policy gradient methods are typically less sample-efficient, since updating the policy renders the old data obsolete; thus, new data have to be collected more frequently. Value-function methods, on the other hand, are often off-policy: this feature allows them to learn from old data and makes them more sample efficient. Sample efficiency determines the suitability of RL algorithms for quantum control. This is particularly relevant for experimental applications. Depending on the details, collecting rewards from the environment can bring a substantial overhead. In a simulator, if it is easier and faster to generate data than training the agent, policy gradient and actor-critic algorithms are typically a good starting point. In potential applications to experiments, deep  $Q$ -learning may be a good choice since it is off-policy. The relative sample efficiency is shown in Fig. 11; the runtime (wall clock time) usually goes in the opposite direction. Deep value iteration methods align with the objective of developing algorithms for both theory and experiment. Ultimately, which algorithm to use depends on the concrete problem and setup at hand: in some cases, it is easier to learn the policy directly, whereas in others—the policy may be too complex, but the value function may be easier to approximate.

When it comes to applications of RL for the manipulation of quantum systems, we can broadly identify four mutually overlapping areas: quantum state manipulation, or preparation, [155,164–171], quantum metrology [172], quantum gate design [173–177], quantum circuit design [178–182], and quantum error correction [183–187]. Out of these, by far the most common set of problems relates to quantum state manipulation. As exemplified in Sec. IV A below, the idea here is to use an RL agent to learn a protocol that transfers the population from an initial to a target state. This has been done extensively for both

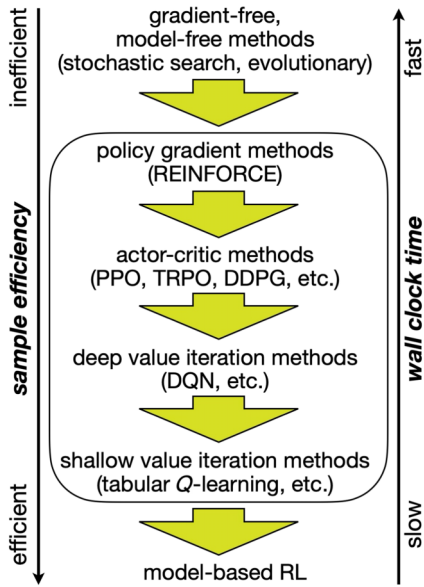


FIG. 11. Heuristic schematic of sample efficiency and runtime of some RL algorithms.

few-particle and many-body systems, in the context of bosonic, fermionic, and spin systems; RL agents have been shown to perform on par with, or better than, some optimal control algorithms in rugged optimization landscapes. An important concern when applying RL to quantum systems is the lack of accessibility to the exact quantum state (quantum states are unmeasurable mathematical constructs). While this is not an issue when training in simulators, it quickly poses a formidable challenge when applications to real experiments come into sight. For this reason, when applying RL to quantum systems one typically defines the RL state as a set of accessible observables, or a sequence of previously taken actions [cf. Sec. IV B 3]. Whereas in many cases RL algorithms work using such partially observable environments, extra care needs to be taken to make sure the learning problem remains well-posed. In turn, the reward function in quantum control problems ranges from maximizing a fidelity (or the negative log-fidelity, in the case of many-body states) through minimizing the energy (as is common when targeting ground states), to optimizing the expectation value of a specific observable; RL agents can also be used to manipulate quantum entanglement [10].

### A. Optimal control via policy gradients

In RL, an agent is confronted with the problem of learning to execute a task by interacting with its environment. The environment is described by a state  $s$ . The agent observes the state of the environment, and is given a choice of actions  $a$  to select from, that alter the state of the environment. The environment can react deterministically or stochastically to these actions, following a set of prescribed

rules, the so-called transition probabilities  $p(s'|s, a)$ , that contain the laws governing the physics of the system. The agent is then given a reward  $r(s, a)$ , which determines how well it performs on the task. Based on this information, the agent then selects a new action, in an attempt to maximize its total expected return  $\mathbf{G}$ . The RL agent selects actions based on its policy  $\pi(a|s)$ , which defines the probability of choosing action  $a$  from state  $s$ . This is repeated for a fixed number of time steps,  $T$ , that define a learning episode. The sequence of state-action-reward triples  $(s, a, r)$  concatenated over the episode, is called a trajectory  $\tau$ .

In the following, we will take a look at some simple-to-implement RL algorithms. Let us start by deriving REINFORCE—the simplest policy gradient algorithm. The reinforcement learning objective  $\mathbf{J}$  is to maximize the expected total return, starting from an initial state and taking actions according to the policy  $\pi(a|s)$  thereafter. Let us denote the probability for the environment to transition from state  $s$  to state  $s'$  upon taking action  $a$  by  $p(s'|s, a)$ ; for an episode of  $T$  steps, denoting the initial state distribution by  $p(s_0)$ , the probability for the trajectory  $\tau = (s_0, a_0, r_1, s_1, a_1, \dots, s_{T-1}, a_{T-1}, r_T, s_T)$  to occur is given by

$$P_\pi(\tau) = p(s_0) \prod_{t=1}^T \pi(a_t|s_t) p(s_{t+1}|s_t, a_t). \quad (177)$$

The RL objective can be defined formally as

$$\mathbf{J} = \mathbb{E}_{\tau \sim P_\pi} [\mathbf{G}(\tau) | s_{t=0} = s_0], \quad \mathbf{G}(\tau) = \sum_{t=1}^T r(s_t, a_t), \quad (178)$$

where  $\mathbf{G}$  is the total return for a fixed trajectory  $\tau$ . The notation  $\mathbb{E}_{\tau \sim P_\pi}$  means that we draw a trajectory  $\tau$  according to the probability  $P_\pi$ ; in practice, we do this by sampling the initial state distribution  $p(s_0)$  and then the policy  $\pi(a|s)$  for each episode step at a time (this is similar to generating runs when playing a video game). The expectation value of the return  $\mathbf{G}(\tau)$  over  $P_\pi$  requires a sum over all possible trajectories  $\tau$ , each weighted by its probability to occur  $P_\pi(\tau)$ . While this sum can formally be written down, it often contains either exponentially or even infinitely many trajectories. This is similar to the concept of partition functions of Feynman path integrals; we can thus borrow the notation:

$$\mathbf{J} = \mathbb{E}_{\tau \sim P_\pi} [\mathbf{G}(\tau) | s_{t=0} = s_0] = \int \mathcal{D}\tau P_\pi(\tau) G(\tau), \quad (179)$$

where  $\mathcal{D}\tau$  is just a formal notation that allows us to do algebraic manipulations. Notice that RL is designed to optimize in expectation, which allows it to handle

non-differentiable or even discontinuous reward functions  $r(s, a)$ .

Policy gradient methods in RL directly approximate the policy  $\pi$  using a variational ansatz  $\pi_\theta$ , parametrized by the unknown parameters  $\theta$ ,  $\pi \approx \pi_\theta$ . The goal is then to find those optimal parameters  $\theta$ , which optimize the RL objective  $\mathbf{J}(\theta)$  (think of the agent that needs to win the game by maximizing the score). To define an update rule for  $\theta$ , we may use gradient ascent (or any of its generalizations [147]). This requires us to evaluate the gradient of the RL objective with respect to the parameters  $\theta$ :

$$\begin{aligned} \nabla_\theta \mathbf{J}(\theta) &= \nabla_\theta \mathbb{E}_{\tau \sim P_\pi} \left[ \sum_{t=1}^T r(s_t, a_t) | s_{t=0} = s_0 \right] \\ &= \nabla_\theta \int \mathcal{D}\tau P_{\pi_\theta}(\tau) \mathbf{G}(\tau) = \int \mathcal{D}\tau \nabla_\theta P_{\pi_\theta}(\tau) \mathbf{G}(\tau), \end{aligned} \quad (180)$$

where we brought the derivative inside the path integral; note that the return  $\mathbf{G}(\tau)$  is independent of  $\theta$  and hence it is not affected by the derivative.

Computing the path integral in Eq. (180) exactly is out of reach because we cannot exhaust the space of all trajectories (similar to the difficulty with computing partition functions in statistical mechanics). In addition, in a model-free setting, we do not know the laws of physics governing the dynamics of the system; as a result, our agent does not have access to the transition probabilities  $p(s'|s, a)$ . Hence, computing the path integral exactly is not a viable approach. These issues require us to find alternative ways, for instance, by estimating the gradient  $\nabla_\theta \mathbf{J}(\theta)$  from samples. This is feasible because our agent interacts with the environment, and hence we can sample trajectories from it. To this end, we need to write the expression for the gradient  $\nabla_\theta \mathbf{J}(\theta)$  as an expectation value over  $P_{\pi_\theta}(\tau)$ . This can be accomplished by using the logarithmic derivative rule  $\nabla_\theta P_{\pi_\theta} = P_{\pi_\theta} \nabla_\theta \log P_{\pi_\theta}$  almost everywhere (i.e., up to a set of measure zero), in Eq. (180):

$$\begin{aligned} \nabla_\theta \mathbf{J}(\theta) &= \int \mathcal{D}\tau \nabla_\theta P_{\pi_\theta}(\tau) \mathbf{G}(\tau) \\ &= \int \mathcal{D}\tau P_{\pi_\theta}(\tau) \nabla_\theta \log P_{\pi_\theta}(\tau) \mathbf{G}(\tau) \\ &= \mathbb{E}_{\tau \sim P_\pi} [\nabla_\theta \log P_{\pi_\theta}(\tau) \mathbf{G}(\tau)]. \end{aligned} \quad (181)$$

The next step is to evaluate  $\nabla_\theta \log P_{\pi_\theta}(\tau)$ . Since the initial state distribution and the transition probabilities are independent of  $\theta$ , using the definition of  $P_{\pi_\theta}$  in Eq. (177), we have  $\nabla_\theta \log P_{\pi_\theta}(\tau) = \nabla_\theta \log \pi_\theta(\tau)$  where  $\pi_\theta(\tau) = \prod_{t=1}^T \pi_\theta(a_t | s_t)$ . We can therefore use Monte Carlo sampling to estimate the gradients directly from a sample of trajectories  $\{\tau_j\}_{j=1}^N$  using the policy  $\pi_\theta(\tau)$  (together with

the initial state and transition probabilities):

$$\begin{aligned} \nabla_\theta \mathbf{J}(\theta) &= \mathbb{E}_{\tau \sim P_\pi} [\nabla_\theta \log \pi_\theta(\tau) \mathbf{G}(\tau)] \\ &\approx \frac{1}{N} \sum_{j=1}^N \nabla_\theta \log \pi_\theta(\tau_j) \mathbf{G}(\tau_j) \\ &= \frac{1}{N} \sum_{j=1}^N \left( \sum_{t=1}^T \nabla_\theta \log \pi_\theta(a_t^j | s_t^j) \sum_{t'=1}^T r(a_{t'}^j, s_{t'}^j) \right). \end{aligned} \quad (182)$$

Here, we first replace the expectation value over trajectories by its Monte Carlo approximation,  $\mathbb{E}_{\tau \sim P_\pi}[\cdot] \approx N^{-1} \sum_{j=1}^N (\cdot)$ , where each trajectory  $\tau_j$  is drawn from the probability  $P_\pi$ . Then, in the last equality, we used  $\log \pi_\theta(\tau) = \log \prod_{t=1}^T \pi_\theta(a_t | s_t) = \sum_{t=1}^T \log \pi_\theta(a_t^j | s_t^j)$ , together with the definition of the total return  $\mathbf{G}(\tau)$  from Eq. (178).

Since, in practice, we always deal with a finite number  $N$  of trajectories in our sample, the estimate of the gradient for some parameters  $\theta_k$  may have a large variance. To alleviate this issue, we use a two-fold strategy: first, notice that the  $t'$ -sum in Eq. (182) runs over all time steps  $t' = 1, \dots, T$ ; however, only times  $t' > t$  should contribute to the gradient since rewards obtained at earlier times  $t' < t$  do not affect the action at step  $t$ . Keeping a smaller number of summands naturally reduces the variance of the estimate. Second, we introduce a so-called baseline  $b = N^{-1} \sum_j \sum_{t=1}^T r(a_t^j | s_t^j)$ , defined as the average return over the sample. This so-called centering of the training data is a common trick used in ML. The PG update then takes the form

$$\nabla_\theta \mathbf{J}(\theta) \approx \frac{1}{N} \sum_{j=1}^N \sum_{t=1}^T \nabla_\theta \log \pi_\theta(a_t^j | s_t^j) \left[ \sum_{t'=t}^T r(a_{t'}^j, s_{t'}^j) - b \right]. \quad (183)$$

One can show that these tricks do not change the expectation value  $\nabla_\theta \mathbf{J}(\theta) = \mathbb{E}_{\tau \sim P_\pi} [\nabla_\theta \log \pi_\theta(\tau) \mathbf{G}(\tau)]$ . The corresponding gradient ascent update rule reads as

$$\theta \leftarrow \theta + \alpha \nabla_\theta \mathbf{J}(\theta), \quad (184)$$

for some step size (sometimes called a learning rate)  $\alpha$ .

## B. Examples

Let us now apply the REINFORCE algorithm to the single-qubit quantum control problem. We will do this in a series of three examples of increasing complexity. In Example 1, we consider fully observable qubit control: while this is somewhat unrealistic (as we explain later on), it represents the simplest quantum environment, and we use it to gain intuition about how RL works in

practice. Subsequently, in Example 2 we use RL in combination with insights from CD driving to design a simple algorithm for variational quantum control. Unlike CD driving, our agent aims at preparing the target only at the final time step. Finally, in Example 3 we consider a realistic quantum control problem; in contrast to Example 1, the environment here is stochastic and minimally observable. Importantly, as will become clear below, the challenge of applying RL to quantum systems lies with the definition of the RL environment—the choice of actions, states, and rewards—rather than with the RL algorithm itself. We therefore place significant focus on discussing how to set up the environment, and refer the interested reader to the accompanying Jupyter notebooks for a practical implementation of the PG algorithm.

### 1. Example: Universal single-qubit state preparation

Our goal in this example will be to train an RL agent to learn how to prepare the target state  $|\psi_*\rangle = (1, 0)^t$  starting from any single-qubit initial state  $|\psi\rangle$  by using a set of predefined gates.

We begin by defining the qubit environment that models the action of gates on the state of the two-level system (2LS). The state of a qubit  $|\psi\rangle \in \mathbb{C}^2$  is modeled by a two-dimensional complex-valued vector with unit norm:  $\langle\psi|\psi\rangle := \sqrt{|\psi_1|^2 + |\psi_2|^2} = 1$ . Every qubit state is uniquely described by two angles  $\theta \in [0, \pi]$  and  $\varphi \in [0, 2\pi)$ :

$$|\psi\rangle = \begin{pmatrix} \psi_1 \\ \psi_2 \end{pmatrix} = e^{i\alpha} \begin{pmatrix} \cos \frac{\theta}{2} \\ e^{i\varphi} \sin \frac{\theta}{2} \end{pmatrix}. \quad (185)$$

The overall phase  $\alpha$  of a single quantum state has no physical meaning. Thus, any single qubit state can be pictured as an arrow on the unit sphere (called the Bloch sphere) with spherical coordinates  $(\theta, \varphi)$ .

To operate on qubits, we use quantum gates. Quantum gates are represented as unitary transformations  $U \in \text{U}(2)$ , where  $\text{U}(2)$  is the unitary group in two dimensions. Gates act on qubit states by matrix multiplication transforming an input state  $|\psi\rangle$  into an output state  $|\psi'\rangle$ :  $|\psi'\rangle = U|\psi\rangle$ . For this problem, we consider four gates and their inverses:

$$\begin{aligned} U_0 &= \mathbb{1}, & U_x &= \exp(-i\delta t\sigma^x/2), & U_y &= \exp(-i\delta t\sigma^y/2), \\ U_z &= \exp(-i\delta t\sigma^z/2), \end{aligned} \quad (186)$$

where  $\delta t$  is a fixed time step,  $\exp(\cdot)$  is the matrix exponential,  $\mathbb{1}$  is the identity, and we remind that the Pauli matrices are defined

$$\begin{aligned} \mathbb{1} &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, & \sigma^x &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, & \sigma^y &= \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \\ \sigma^z &= \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \end{aligned} \quad (187)$$

We also allow our RL agent to use the inverse operations for  $U_x^\dagger, U_y^\dagger, U_z^\dagger$ .

To determine if a qubit, described by the state  $|\psi\rangle$ , is in a desired target state  $|\psi_*\rangle$ , we compute the fidelity

$$\mathcal{F} = |\langle\psi_*|\psi\rangle|^2 \in [0, 1]. \quad (188)$$

Physically for a single qubit, the fidelity corresponds to the angle between the arrows representing the state on the Bloch sphere (we want to maximize the fidelity but minimize the angle between the states).

Now, let us define an episodic RL environment, which contains the laws of physics that govern the dynamics of the qubit (i.e., the application of the gate operations to the qubit state). We recall that the environment in RL is different from the notion of environment in open systems. Our RL agent will interact with this environment by taking actions and observing its states to learn how to control the qubit to bring it from an initial state to a prescribed target state. To this end, we define the RL states  $s = (\theta, \varphi)$  as an array containing the Bloch sphere angles of the quantum state. For each step within an episode, the agent can choose to apply one out of the actions, corresponding to the seven gates  $(\mathbb{1}, U_x, U_y, U_z, U_x^\dagger, U_y^\dagger, U_z^\dagger)$ . We use the instantaneous fidelity with respect to the target state as a reward:  $r_t = \mathcal{F} = |\langle\psi_*|\psi(t)\rangle|^2$ . More formally, we have to define:

- (1) **State space:**  $\mathbf{S} = \{(\theta, \varphi) \mid \theta \in [0, \pi], \varphi \in [0, 2\pi)\}$ . There are no well-defined states to terminate episodes in this task. Instead, we consider a fixed number of time steps  $T$ , after which the episode terminates deterministically. The target state (i.e., the qubit state we want to prepare) is  $|\psi_*\rangle = (1, 0)^t$ : it has the Bloch sphere coordinates  $s_* = (0, 0)$ .
- (2) **Action space:**  $\mathbf{A} = \{\mathbb{1}, U_x, U_y, U_z, U_x^\dagger, U_y^\dagger, U_z^\dagger\}$ . Actions change the state of the controlled system; they act on RL states as follows:
  - (a) if the current state is  $s = (\theta, \varphi)$ , we first create the quantum state  $|\psi(s)\rangle$ .
  - (b) we apply the gate  $U_a$  corresponding to action  $a$  to the quantum state, and obtain the new quantum state  $|\psi(s')\rangle = U_a|\psi(s)\rangle$ .
  - (c) last, we compute the Bloch sphere coordinates, which define the next state  $s' = (\theta', \varphi')$ , using the Bloch sphere parametrization for qubits given in Eq. (185). Note that all actions are allowed from every state.
- (3) **Reward space:**  $\mathbf{R} = [0, 1]$ . We use the fidelity between the next state  $s'$  and the target state  $s_*$  as a reward at every episode step:

$$r(s, s', a) = \mathcal{F} = |\langle\psi_*|U_a|\psi(s)\rangle|^2 = |\langle\psi_*|\psi(s')\rangle|^2 \quad (189)$$

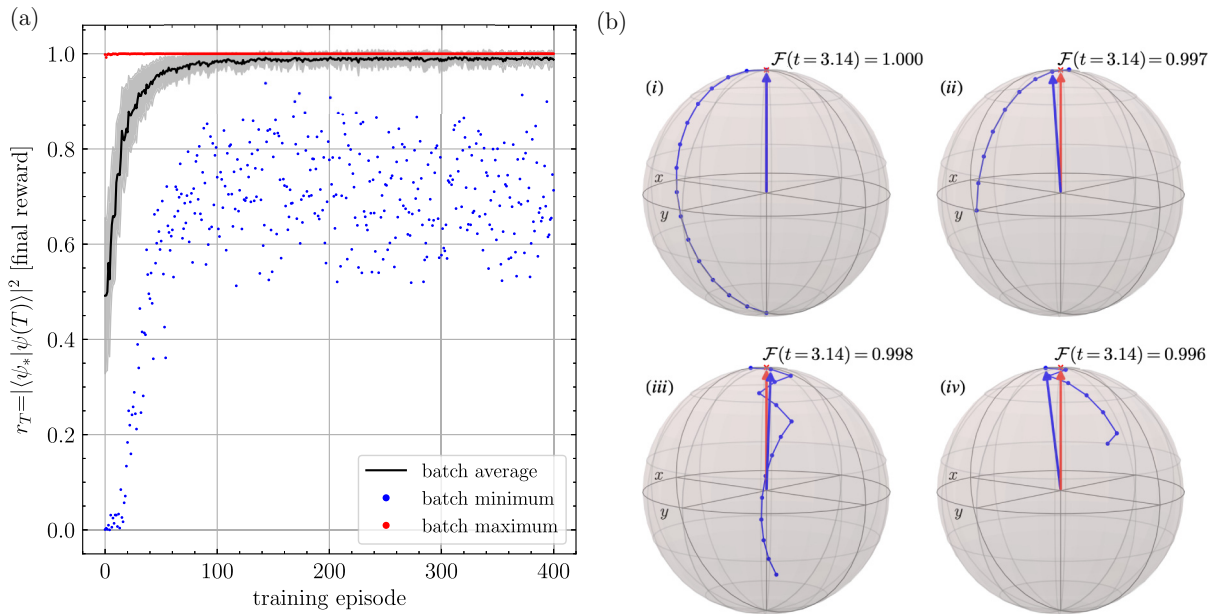


FIG. 12. Universal single-qubit control using reinforcement learning: (a) training curves for an RL agent learning to prepare the target  $|\uparrow\rangle$  states (north pole of Bloch sphere) starting from any initial state. The final reward at the end of the protocol is shown, averaged over a batch of  $N = 256$  random initial states (black), as well as the minimum and worst-case (blue dashed) and the maximum (red dashed) values within the batch. (b) (i-iv) trajectories learned by the RL agent for four different initial states; the target (final) state is shown in red (blue). The agent achieves a fidelity  $1 - \mathcal{F}$  within the 1% error range. We used a fully connected neural network with (512, 256, 7) neurons in each layer to learn the optimal strategy using the policy gradient algorithm. The time step is  $\delta t = T/\pi$  for a total of  $T = 15$  time steps.

for all states  $s, s' \in \mathcal{S}$  and actions  $a \in \mathbf{A}$ . This choice is motivated by the fact that the fidelity of the instantaneous state with the target state is a measure of the distance between the states. Alternatively, one can use other reward functions, such as the ground state expectation of the parent Hamiltonian for the target state.

Figure 12(a) shows the result of the training procedure. We use a uniform initial state distribution  $p(s_0)$  since this will expose the agent to a sample of the entire single-qubit Hilbert space; in turn, this will help it learn how to steer any initial state into the target. To train the agent, we employ the REINFORCE algorithm with a batch of  $N = 256$  trajectories; we invite the interested reader to explore the accompanying Jupyter notebook. To this end, the agent runs a large number of episodes during which it explores the control protocol space by taking actions. The reward at the end of the episode  $r_T = |\langle \psi_* | \psi(T) \rangle|^2$  tells us how close the final state is to the target. Monitoring this quantity as a function of the training episodes (i.e., iterations of the algorithm) shows a steady increase of the batch average (black solid line) to a value close to unity. For comparison, we also show the best (red) and worst (blue) trajectories from the batch. The spread in the worst trajectories, as training progresses, comes from exploration:

the RL agent takes actions probabilistically, and hence occasionally, suboptimal actions can be taken.

To test the performance of the trained agent, we consider a greedy policy that takes the action with the highest probability deterministically. Figure 12(b) shows the Bloch sphere trajectories of qubits controlled by the trained RL agent, starting from four different initial states (the time discretization is marked by the blue dots on the trajectory); the target state is shown in as a red arrow, and the final state in blue. Importantly, we emphasize that there is no further optimization required: once the training process has completed successfully, the same agent knows how to produce the optimal protocol starting from any initial state. This is a prime example of the learning (or generalization) capabilities of RL agents, which contrasts RL from optimal control. In all cases, the agent manages to bring the state close to the target; the small leftover deviation is caused by the finite time step  $\delta t$ . Movies within the Supplemental Material [188] show the trajectories in time. We invite the interested reader to check out the accompanying Jupyter notebook and explore the behavior of the agent further [19].

The example above serves as a stepping stone to learn how RL works in practice. In principle, all one has to do to control more complex quantum systems is redefine the environment, and the same implementation of the RL

algorithm can be re-used. However, for quantum many-body systems, this is a tricky task: for instance, there is no convenient Bloch sphere representation for the quantum states. In this case, one can resort to using the bare wavefunction amplitudes, or other exact representations of the state (e.g., different momentum modes in cases, such as the transverse-field Ising model, where the problem can be mapped to a free system, or Clifford tableaux for Clifford circuit dynamics, etc.). One may also resort to suitable approximations, e.g., matrix product states [155] whenever appropriate. We should also point out that the fidelity is typically exponentially small (in the number of qubits) for many-body states; in this case, one typically resorts to minimizing the energy of a parent Hamiltonian (i.e., a Hamiltonian whose ground state is the target state) [165]. It is important to keep in mind that, while RL provides a way to finding optimal protocols for quantum systems, it does not resolve any of the problems of many-body physics associated with representing quantum states or extracting information from them.

## 2. Example: RL vs. counterdiabatic driving in the presence of trotterization errors

We now build on the example above and apply RL to the problem of preparing the ground state of a target single-qubit Hamiltonian. We adopt a hybrid approach where the structure of the control terms in the Hamiltonian are designed using knowledge from CD driving (Sec. II B). In particular, we compare and contrast the behavior of the RL agent to counterdiabatic (CD) driving. Before we dive into the RL implementation, let us briefly mention that CD driving can be enhanced using other ML techniques, see, e.g., Ref. [189].

As we have seen in Sec. II, CD driving introduces additional control terms to suppress excitations during the protocol:

$$H_{\text{CD}}(t) = \Delta\sigma^x + v(t)\sigma^z - \frac{1}{2} \frac{\Delta \dot{v}(t)}{\Delta^2 + v^2(t)} \sigma^y. \quad (190)$$

As a result, evolving an eigenstate of the initial Hamiltonian  $H(0) = \Delta\sigma^x + v(0)\sigma^z$  in time with  $H_{\text{CD}}(t)$ , the system follows the instantaneous eigenstates of the Hamiltonian  $H(t)$  at all times.

We now initialize the system in the ground state of the Hamiltonian  $H(0)$  at  $v(0) = v_i = +2\Delta$ , and want to transfer the population in the ground state at  $H(t_*)$  at  $v(t_*) = v_* = -2\Delta$  in a finite amount of time  $T$  (referred to as the duration of the protocol). During the evolution, the system is subject to the Hamiltonian

$$H_g(t) = \Delta\sigma^x + v(t)\sigma^z - g(t)\sigma^y, \quad (191)$$

where  $v(t) = (v_* - v_i)t/T + v_i$  is a linear drift term that is always kept on (and the RL agent has no control over it),

while  $g(t)$  is the unknown counter protocol that we want our RL agent to find. The form of the Hamiltonian  $H_g(t)$  is motivated by the CD Hamiltonian in Eq. (190); however, we do not make any assumptions on the structure of the drive protocol  $g(t)$  whatsoever, and we leave it to the RL agent to find its shape.

The qubit evolves following  $H_g(t)$ , and its state at time  $t$  is determined by the time evolution operator, which can be expressed in terms of the time-ordered exponential of the Hamiltonian  $H_g(t)$  generating the dynamics:

$$|\psi(t)\rangle = U(T, 0)|\psi(0)\rangle = \mathcal{T} \exp\left(-i \int_0^t ds H_g(s)\right) |\psi(0)\rangle. \quad (192)$$

Once again, we use the fidelity to determine if a qubit, described by the state  $|\psi(t)\rangle$ , is in a desired target state  $|\psi_*\rangle$ :

$$\mathcal{F} = |\langle\psi_*|\psi(t)\rangle|^2, \quad \mathcal{F} \in [0, 1]. \quad (193)$$

Now that we have defined the control problem, we discuss how to phrase it within the RL framework. Similar to the example in Sec. IV B 1, we define an episodic RL environment, which contains the laws of physics that govern the dynamics of the qubit [i.e., the time evolution following the Hamiltonian  $H_g(t)$ ]. The environment will provide data for the RL agent to learn how to bring the qubit to the target ground state. To this end, we discretize the protocol into  $N_T$  time steps of size  $t$ , so that the total duration is  $T = N_T\delta t$ . At each step, the RL agent constructs the value of the control protocol by selecting an action out of the available action set. The total number of time steps in the protocol comprises a learning episode. Once an episode comes to an end, the environment is reset to its initial state, and the agent starts constructing the protocol again. In each subsequent episode, the agent uses the knowledge gained during previous protocols.

Now, for each time step of size  $\delta t$  within the episode, the agent has to determine the optimal value of the protocol function  $g(t)$ ; to this end, starting from some initial value  $g_i$ , we let the agent find the optimal relative change in the protocol (i.e., by how much the protocol value has to change at the given time step). In practice, we define a minimum protocol size  $\delta g$ , and consider  $2n + 1$  steps of size  $-n\delta g, -(n-1)\delta g, \dots, -\delta g, 0, \delta g, \dots, (n-1)\delta g, n\delta g$  that the agent has to choose from. We also set  $g_i = \frac{1}{2} \frac{\Delta \dot{v}(0)}{\Delta^2 + v^2(0)}$ , which fits our goal to directly compare the performance of the agent with CD driving, cf. Eq. (191).

- (1) **State space:**  $\mathbf{S} = \{(\theta, \varphi) \mid \theta \in [0, \pi], \varphi \in [0, 2\pi)\}$ .

There are no well-defined terminal states in this task. Instead, we consider a fixed number of time steps, after which the episode terminates deterministically.

(2) **Action space:**  $\mathbf{A} = \{-n, -(n-1), \dots, -1, 0, 1, \dots, n-1, n\}$ . Actions correspond to fixed discrete amounts by which the agent can change the protocol  $g(t)$  (see above). Actions act on RL states as follows:

- (a) if the state at time step  $n$  (corresponding to physical time is  $t = n\delta t$ ) is  $s = (\theta, \varphi)$ , we first create the quantum state  $|\psi(s)\rangle$ ;
- (b) we apply to the quantum state the unitary  $U_a$  corresponding to action  $a \in \mathbf{A}$ , defined by

$$U_a = \exp(-i\delta t H_{a \times \delta g}(n\delta t)),$$

$$H_{a \delta g}(n\delta t) = \Delta\sigma^x + v(n\delta t)\sigma^z - a\delta g\sigma^y, \quad (194)$$

where  $n$  labels the time steps, and  $\delta g$  ( $\delta t$ ) is a fixed minimum protocol (time) step size (parameters of our choice);

- (c) we obtain the new quantum state  $|\psi(s')\rangle = U_a|\psi(s)\rangle$ ;

(d) last, we compute the Bloch sphere coordinates which define the next state  $s' = (\theta', \varphi')$ , using the Bloch sphere parametrization for qubits given above. Note that all actions are allowed from every state.

- (3) **Reward space:**  $\mathbf{R} = [0, +\infty)$ . Since we want to incentivize the agent to prepare the target state at the end of the protocol, the reward is given only at the end of the episode. Note the difference to CD driving, where the system is in the instantaneous ground state at any given time during the protocol. In general, such reward functions are referred to as sparse rewards in RL. The return  $\mathbf{G} = r_T$  is then equal to the final reward.

Thus, the reward is zero at each step during the protocol:  $r_{t < T} = 0$ , except at the last step, where we use the fidelity between the final state  $|\psi(T)\rangle$  and the target state  $|\psi_*\rangle$  as a reward at the end of every protocol:

$$r(s, s', a) = r_t = \begin{cases} 0, & \text{for } t < T \\ -\log_{10}(1 - \mathcal{F}), & \text{for } t = T \end{cases}, \quad \mathcal{F} = |\langle \psi_* | \psi(T) \rangle|^2, \quad |\psi(T)\rangle = U(T, 0)|\psi(0)\rangle, \quad (195)$$

for all states  $s, s' \in \mathbf{S}$  and actions  $a \in \mathbf{A}$ .

Since we are interested in high-fidelity protocols, it is important to build a high resolution into the reward signal. One way to do this is to use the logarithmic fidelity  $1 - \mathcal{F}$  with respect to the target state at the end of the protocol as a reward. The logarithm allows us to learn on a logarithmic scale. Using a base-10 logarithm allows us to interpret the reward as the number of decimal places in the deviation of the fidelity from unity: e.g.,  $r_t = 1$  corresponds to  $\mathcal{F} = 0.9$ , while  $r_t = 2$  is equivalent to  $\mathcal{F} = 0.99$ , etc.

We now apply the REINFORCE algorithm to this RL problem. Figure 13(a) shows the training curves for the reward  $r_T$  at the end of the episode (equal to the total return in this case). Like before, we consider the batch average (black), and the best (red)/worst (blue) protocols within the batch. Notice that achieving a 6-digit reward (i.e., a fidelity of  $\mathcal{F} = 0.999999$ ) requires about  $8 \times 10^3$  training episodes. We also see that there exist protocols with  $r_T > 8$  which the agent encounters during the exploration but fails to learn; this can be due to various reasons, e.g., the corresponding protocols may be unstable to small perturbations (recall that the agent takes actions probabilistically), or they may correspond to very deep and isolated

minima of the control landscape (and hence, they are difficult to encounter often enough for the signal to accumulate and lead to stable learning).

Figure 13(b) shows a comparison between the exact CD protocol (red dashed line), its discretization with the environment time step  $\delta t$  (blue), and the protocol learned by the RL agent (green). Note that, while the continuous CD protocol is by construction perfect ( $r_T = \infty$ ), its discretized version has a reduced logarithmic fidelity  $1 - \mathcal{F}$  of  $r_T = 4.144$ . For comparison, the RL agent achieves  $r_T = 5.983$ . The RL agent achieves a better result, since it optimizes for preparing the target ground state only and pays no attention to the intermediate states of the system, see deviation from unity in the overlap with the instantaneous ground state [brown curve in Fig. 13(c)]. By contrast, exact CD driving keeps the system in the instantaneous GS of  $H(t)$  at *any* time  $t \in [0, T]$  with strictly unit fidelity. For example, at time  $t/T = 0.5$  in the middle of the protocol, the RL protocol achieves an instantaneous overlap with the GS of  $r_T = 2.678$  (brown curve at  $t = 0.6$ , inset). We invite the interested reader to check out the accompanying Jupyter notebook and explore the behavior of the agent further [19].

The idea of using hybrid RL-OC algorithms extends immediately to multiqubit systems. The discrete nature of the control setup (time split into discrete time steps




FIG. 13. Reinforcement learning the CD control problem. (a) Reward (logarithmic fidelity) at the end of the protocol against the number of training iterations (episodes): training is performed using a batch size of 256. RL agent achieves an average fidelity  $1 - \mathcal{F}$  of about  $10^{-6}$  (black line, with the shaded area showing the standard deviation); worst (best) protocols in the batch are shown in blue (red). (b) Comparison of protocols and their final reward (corresponding to the number of decimals in  $1 - \mathcal{F}$ ). The exact continuous CD protocol (dashed red) performs worse than the RL agent (green) when time is discretized (blue). (c) Main plot shows the fidelity of the protocol learned by the RL agent (green) and the instantaneous fidelity (brown); inset shows the corresponding infidelities on a logarithmic scale.

and protocol values taken from a predefined finite discrete action set) suggests a natural extension in terms of quantum circuits built out of unitary gates. It turns out, a suitable choice for actions are gate generators built out of the terms appearing in the expansion for the adiabatic gauge potential. For instance, in Ref. [165], this idea was used to prepare the many-body ground state of Heisenberg and Ising chains, by using the (negative) energy of the final state as a reward function.

### 3. Example: Continuous single-qubit feedback control using quantum data

While the previous two examples demonstrated the basics of using RL, they employ RL frameworks that make it difficult to apply the corresponding agents in realistic experiments. The issue comes from the data used to train the agent, and is two-fold: (i) at each time step, the agent is assumed to observe the state of the qubit; this is problematic since quantum states are not observable. While it is not unimaginable to do a full-state tomography for a single qubit, besides the extra measurement overhead this would still result in some shot noise in the training data; moreover, tomography requires many repetitions of the experiment which can become an issue in the presence of noise because each trajectory is likely to end up in a different quantum state. We therefore need to revisit the definition of RL states if we want to run our agents on real quantum experiments. (ii) The RL agent uses the fidelity as a reward; similar to the issue above, the fidelity needs to be estimated from a number of binary measurements (following the projective measurement, the system is either found in the ground state, or in the excited state) taken sequentially and, at the very least, will also be noisy.

To resolve these issues and design an RL framework that can interact with experiments, we have to replace the fully observable RL environment with a *partially observable* environment. Hence, our agent will only have access to (partial) observations of the true RL state, which drastically reduces the amount of information it receives, and hence makes training more challenging; in fact, this is further exacerbated by single-shot qubit measurements that only contain binary information. For the implementation of our simulator, this means that we will have an internal variable that defines the exact RL state (here the full quantum wavefunction) and which undergoes the exact quantum dynamics; however, the agent-environment interface will only have access to observations of that state. Below, we will design an RL framework where the RL observations (now playing the role of the RL states) are all binary.

More concerning is the quantum nature of the projective measurement itself, since it collapses the state of the qubit: this is similar to the state resetting procedure at the beginning of every episode, except now it occurs at each time step. The quantum Zeno effect then tells us that the state will barely evolve at all, and hence we may have a hard time bringing it to the desired target. To alleviate the measurement collapse problem while adhering to the laws of quantum mechanics, we introduce an ancilla qubit on which to perform measurements [see Fig. 14(a)]. Whenever the controlled qubit and the ancilla are weakly entangled, measuring the ancilla will only weakly perturb the state of the qubit; that said, the weaker the entanglement, the smaller the amount of information we can extract about the state of the qubit from measuring the ancilla.

Last, we turn to the issue of defining the reward. As already mentioned above, we can only afford to collapse the state of the qubit once per episode, and it is most natural

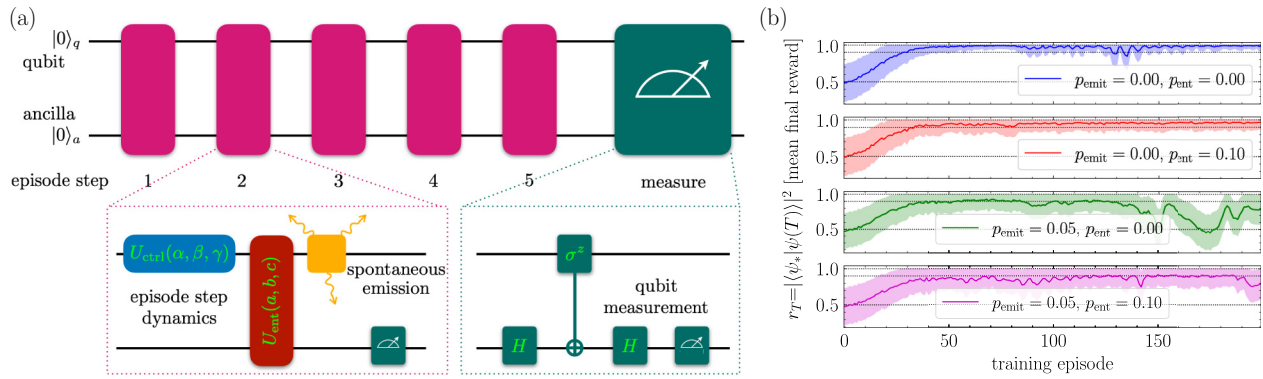


FIG. 14. Reinforcement learning from quantum data. (a) We couple a qubit to an ancilla which we can measure projectively. At each episode step, the agent can apply a control gate  $U_{\text{ctrl}}(\alpha, \beta, \gamma)$  (blue); however, the system is also subject to entangling noise  $U_{\text{ent}}(a, b, c)$  (red) and spontaneous emission (orange). After each step we project the ancilla on the ground state. The time step, the measurement output, and the detected spontaneously emitted photon represent the RL observations; the reward is given at the end of the circuit by the binary output of the qubit collapse on the target state. At the end of the circuit, we use the ancilla to project the state of the qubit to the target state; this produces a binary output that defines the return. (b) Training curves showing the final fidelity over a batch of 2048 trajectories for: (blue) the clean case (i.e., no entangling noise and no spontaneous emission); (red) in the presence of entangling noise only; (green) in the presence of spontaneous emission only; (magenta) in the presence of both noise and spontaneous emission. The shaded area marks the  $1\sigma$  window over the batch. The horizontal dotted lines mark the 0.5, 0.8, 1.0 fidelity thresholds, respectively. We use a single-layer fully-connected neural network with 8 hidden neurons.

to do this at the end (which sets the return to be equal to the final reward, a.k.a. sparse reward function). However, instead of using the exact fidelity as a return, we will use the binary measurement output.

At first sight, common sense may make it seem ludicrous that we can train any agent to learn from only binary information. Especially in realistic experiments where noise, decoherence, and spontaneous decay processes cause the qubit to end up in a different final state every time, even if we applied the same sequence of actions deterministically in an attempt to exactly repeat an episode. In such a scenario, estimating the qubit state and the fidelity requires using the theoretical framework of density matrices, whereas the RL agent is tasked to prepare a pure target state.

To demonstrate that this is not as crazy as it may first sound, let us define an experimentally-friendly RL framework for a noisy qubit, and use the REINFORCE algorithm to train an RL agent to prepare a pure target state. As already advertised, we consider a qubit that can easily be reset to the ground state  $|1\rangle_q$  at any time, and our goal is to transfer its population into a target state  $|\psi_*\rangle_q$ . To do this, this time we use continuous gates parametrized by the three angles  $\alpha, \beta, \gamma \in [-\pi, \pi)$  (note that while one can also use an Euler angle parametrization, we prefer to work with a Tait-Bryan parametrization since it avoids solutions being degenerate in the case of  $\beta = 0$ ):

$$\begin{aligned}
 U_{\text{ctrl}}(\alpha, \beta, \gamma) \\
 = \exp(-i\gamma\sigma^z/2) \exp(-i\beta\sigma^y/2) \exp(-i\alpha\sigma^x/2).
 \end{aligned}
 \tag{196}$$

The goal of the RL agent is to learn to select the optimal angles to prepare the target starting from the fixed initial state. In addition, we will require that the agent can keep the state in the target over some fixed time.

As we mentioned above, our qubit is paired with an ancilla initialized in the state  $|0\rangle_a$ . Crucially, any measurement on the system will be performed on the ancilla only, in order to avoid collapsing the state of the qubit; the binary output of these measurements will be used to define the observations and the return to train the RL agent.

To make the problem more realistic, we now assume that the state of the qubit can undergo spontaneous emission and decay into its ground state. This spontaneous decay should in principle also be present for the ancilla, but we leave that for the interested reader to implement as an exercise. We also assume that some noise can lead to a small entanglement between the qubit and the ancilla, which we cannot control. In the following, we describe how to model these processes which define the RL environment.

**Spontaneous emission of qubit.**—Let us assume that the qubit state can undergo spontaneous emission with some probability  $p_{\text{emit}}$ ; whenever it occurs, this process causes the qubit to decay to the ground state  $|1\rangle_q$ , no matter which state it was originally in.

We mention in passing that this is just one example of an error channel that is relevant to some (but not all) platforms; however, the procedure below can be generalized to other loss processes, depending on the particular system of interest.

Spontaneous emission is interesting for our RL setup since, by energy conservation, it is accompanied by a

photon that can in principle be detected. Thus, although spontaneous emission is a stochastic process, we have a means to detect when it happened. We can then pass this information to our RL agent in the form of an observation to improve learning.

We use a simplified model for spontaneous emission as follows:

$$|\psi\rangle_q \longrightarrow \begin{cases} \frac{P_q^z |\psi\rangle}{\sqrt{|\langle\psi|P_z|\psi\rangle_q|}}, & \text{with prob. } p_{\text{emit}} \\ |\psi\rangle_q, & \text{with prob. } 1 - p_{\text{emit}}, \end{cases} \quad (197)$$

where  $P_q^z = \frac{1}{2}(1 - \sigma^z)_q$  is the projector on the ground state  $|1\rangle_q$  of the qubit.

**Entangling noise between qubit and ancilla.**—In realistic situations, the qubit and the ancilla are not perfectly decoupled from one another. Leftover terms in the implementation of the gates can lead to small entanglement between the qubit and the ancilla.

We assume that this entangling noise is fixed, and can be parametrized by the parameters  $a, b, c$ , that we will refer to

as angles, as

$$\begin{aligned} |\psi\rangle_q |0\rangle_a &\longrightarrow \exp\left(-i(a\sigma_q^x \sigma_a^x + b\sigma_q^y \sigma_a^y + c\sigma_q^z \sigma_a^z)\right) \\ |\psi\rangle_q |0\rangle_a &= U_{\text{ent}}(a, b, c) |\psi\rangle_q |0\rangle_a, \end{aligned} \quad (198)$$

where  $\sigma_q^\alpha \sigma_a^\alpha$  ( $\alpha = x, y, z$ ) acts on the qubit-ancilla pair, respectively. One can check explicitly that these three operators commute and hence the exponential factors into three separate entangling gates that can be applied to the state consecutively (different Pauli matrices anti-commute for the same qubit but commute for different qubits).

As a result of this noisy operation, the qubit and ancilla states may get entangled. Hence, measuring the ancilla to produce a binary partial observation, will perturb the state of the qubit; the strength of the perturbation is determined by the angles  $a, b, c$  which are a property of the physical system (in our simulations, we use random numbers drawn uniformly from the interval  $[-p_{\text{ent}} \times \pi, p_{\text{ent}} \times \pi]$  that are fixed across different episodes). It is this back-action that is difficult to model over multiple steps in a stochastic setting, and where RL can have an edge compared to more traditional control algorithms.

**Measuring the state of the ancilla.**—Recall that we want our RL agent to receive as observations the binary output of ancilla measurements. We now discuss how to perform a measurement of the ancilla in the  $z$ -basis. Suppose the qubit-ancilla system is in the joint state  $|\psi\rangle$ . The ancilla measurement is then given by

$$|\psi\rangle \longrightarrow \begin{cases} \frac{P_a^z |\psi\rangle}{\sqrt{p}}, & \text{with prob. } p = |\langle\psi|P_a^z|\psi\rangle|^2 \quad \& \text{measurement outcome } -1 \\ \frac{(1 - P_a^z) |\psi\rangle}{\sqrt{1 - p}}, & \text{with prob. } 1 - p \quad \& \text{measurement outcome } +1 \end{cases} \quad (199)$$

Here the projector  $P_a^z = 1_q \otimes \frac{1}{2}(1 - \sigma^z)_a$  acts only on the ancilla subspace.

Note that, in contrast to the spontaneous emission of the qubit, the measurement occurs with a probability that depends on the state of the system. Moreover, the state after the measurement always collapses.

**Measuring the state of the qubit.**—Finally, we turn to the measurement that produces the reward data. To determine if the qubit at the end of the control circuit is in the target state, we can use the ancilla to perform a measurement of the operator

$$\sigma_* = \vec{n}_* \cdot \vec{\sigma}, \quad (200)$$

where  $\vec{\sigma}_q = (\sigma^x, \sigma^y, \sigma^z)_q$  is the vector of Pauli operators, and  $\vec{n}_* = \langle\psi_*|\vec{\sigma}|\psi_*\rangle_q$ . The output of the measurement returns the  $+1$ -eigenvalue of  $\sigma_*$  with probability

$\mathcal{F} = |\langle\psi_*|\psi\rangle_q|^2$ . After the measurement, the qubit is projected to the eigenstate corresponding to the measured eigenvalue.

In practice, the measurement is performed using the ancilla, which is initialized in the  $|0\rangle_a$  state. To this end, one first applies a Hadamard gate  $H_a = (\sigma^x + \sigma^z)_a / \sqrt{2}$  on the ancilla, then a controlled- $\sigma_*$  gate which we denote by  $C_{\sigma_*}$ , followed by a second Hadamard gate on the ancilla [see Fig. 14(a), green inset]. This circuit entangles the qubit state with the ancilla. Measuring the ancilla in the  $z$ -basis then returns the eigenvalue  $+1$  with probability  $\mathcal{F} = |\langle\psi_*|\psi\rangle_q|^2$ , and collapses the state of the qubit on  $|\psi_*\rangle_q$ .

To see how this works, let us expand the state of the qubit in the eigenbasis of  $\sigma_*$ :

$$|\psi\rangle_q = a|\psi_*\rangle_q + b|\psi_*^\perp\rangle_q, \quad (201)$$

where  $\sigma^*|\psi_*\rangle_q = +|\psi_*\rangle_q$  and  $\sigma^*|\psi_*^\perp\rangle_q = -|\psi_*^\perp\rangle_q$ . The state of the qubit-ancilla system before the measurement circuit is thus

$$|\psi\rangle_q|0\rangle_a = a|\psi_*\rangle_q|0\rangle_a + b|\psi_*^\perp\rangle_q|0\rangle_a, \quad (202)$$

where  $q, a$  label the qubit and ancilla, respectively. After the first Hadamard on the ancilla, the state becomes (we use the convention  $|0\rangle = (1, 0)^T$  corresponding to the excited state):

$$\begin{aligned} H_a|\psi\rangle_q|0\rangle_a &= \frac{a}{\sqrt{2}}|\psi_*\rangle_q(|0\rangle_a + |1\rangle_a) + \frac{b}{\sqrt{2}}|\psi_*^\perp\rangle_q(|0\rangle_a + |1\rangle_a) \\ &= \frac{1}{\sqrt{2}}(a|\psi_*\rangle_q + b|\psi_*^\perp\rangle_q)|0\rangle_a + \frac{1}{\sqrt{2}}(a|\psi_*\rangle_q \\ &\quad + b|\psi_*^\perp\rangle_q)|1\rangle_a. \end{aligned} \quad (203)$$

Next, we apply the controlled- $\sigma_*$  gate  $C_{\sigma_*}$ , i.e., we apply the  $\sigma_*$  gate on the qubit if the state of the ancilla is  $|1\rangle_a$ , and we do not take any action if the ancilla is in  $|0\rangle_a$ ; in doing so we recall that we decomposed the qubit state  $|\psi\rangle_q$  in the eigenbasis of  $\sigma_*$ :

$$\begin{aligned} C_{\sigma_*}H_a|\psi\rangle_q|0\rangle_a &= \frac{1}{\sqrt{2}}(a|\psi_*\rangle_q + b|\psi_*^\perp\rangle_q)|0\rangle_a \\ &\quad + \frac{1}{\sqrt{2}}(a|\psi_*\rangle_q - b|\psi_*^\perp\rangle_q)|1\rangle_a \\ &= \frac{a}{\sqrt{2}}|\psi_*\rangle_q(|0\rangle_a + |1\rangle_a) + \frac{b}{\sqrt{2}}|\psi_*^\perp\rangle_q(|0\rangle_a - |1\rangle_a). \end{aligned} \quad (204)$$

The second Hadamard gate on the ancilla then maps the state to

$$H_aC_{\sigma_*}H_a|\psi\rangle_q|0\rangle_a = a|\psi_*\rangle_q|0\rangle_a + b|\psi_*^\perp\rangle_q|1\rangle_a. \quad (205)$$

It now becomes clear that if after applying the circuit above we measure the ancilla in the  $z$ -basis, we obtain the measurement result  $+1$  with probability  $|a|^2 = |\langle\psi_*|\psi\rangle_q|^2$ , and  $-1$  with probability  $|b|^2$ , as expected. Moreover, due to the entanglement between the qubit and the ancilla introduced by the  $C_{\sigma_*}$  gate, the measurement of the ancilla automatically projects the state of the qubit in the corresponding eigenstate of  $\sigma_*$ . Note that this outcome is precisely what would come out if one could measure the operator  $\sigma_*$  in the state  $|\psi\rangle_q$ .

We are now fully set to define an episodic RL environment, which contains the laws of physics that govern the dynamics of the qubit-ancilla system (i.e., the application of the control gates and various sources of noise and

decoherence to the qubit state). Our RL agent will interact with this environment to learn how to control the qubit to bring it from the initial ground state  $|1\rangle_q$  to a prescribed target state  $|\psi_*\rangle_q = \cos\frac{\theta}{2}|0\rangle + e^{i\phi}\sin\frac{\theta}{2}|1\rangle$  with  $\theta = \pi/4$  and  $\phi = \pi/3$ . Each episode of the environment has  $T = 5$  steps; the agent applies a gate of the form in Eq. (196) at each step [see Fig. 14(a)].

As discussed above, note that we cannot use the quantum state as an RL state since it is not observable. We therefore define RL observations using three ingredients: (i) information about the time step (the agent should know how many steps there are until the end of the episode), (ii) the binary measurement output of the ancilla measurement after each step, and (iii) the binary output of the photon detector which tells us whether a spontaneous emission event occurred or not.

At each step within an episode, the agent will use its policy to generate the angles  $\alpha, \beta, \gamma$  that define the control unitary from Eq. (196). Because the angles can take on continuous values, we model the policy by an independent and uncorrelated Gaussian distribution for each angle (the agent will learn the mean and variance, respectively)  $\pi_\theta \propto \mathcal{N}(\mu_\alpha, \sigma_\alpha^2)\mathcal{N}(\mu_\beta, \sigma_\beta^2)\mathcal{N}(\mu_\gamma, \sigma_\gamma^2)$ ; the angles to be applied will then be sampled from this policy. This is an example of a continuous control problem. In practice, we use a neural network (or any other variational ansatz) parametrized by the parameters  $\theta$  to learn the optimal values of the means  $\mu_k$  and the standard deviations  $\sigma_k^2$ : for a given observation  $o_t$  which is an input to the neural network, the output is the tuples  $(\mu_k, \sigma_k^2)$ .

Finally, the reward is set to zero at each time step (due to the character of projective measurements that destroy the state), except at the very end of the episode, when we use the ancilla to produce the binary reward corresponding to the quantum measurement.

- (1) **Observation space:**  $\mathbf{S} = \mathbb{Z}_2^{3 \times T}$ . The time step will be parsed to the agent as a one-hot encoding: the corresponding vector is initially set to zero. The ancilla measurements take on the value  $+1$  if the ancilla is found in the state  $|0\rangle_a$ , and  $-1$  otherwise, and are initialized to  $+1$  for all episode steps. The ancilla is measured at the end of each step of the environment and is reset to  $|0\rangle_a$  afterwards. Finally, the photon detector gives  $-1$  (i.e., qubit found in the ground state) if a photon has been detected at a given step and is initialized to  $-1$  (no photon detected *a priori*). Hence, the observation consists of a  $3 \times T$  binary vector. There are no well-defined terminal states in this task; instead, we consider a fixed number of  $T = 5$  time steps, after which the episode terminates deterministically.
- (2) **Action space:**  $\mathbf{A} = [-\pi, \pi]^3$ . Note that the action space in this example is continuous; in particular, a single step is sufficient to reach any target state from

any other state on the Bloch sphere. However, the presence of different types of noise makes this more challenging. Actions act on RL states as follows:

- (a) sample the angles  $\alpha, \beta, \gamma$  from the Gaussian policy;
- (b) build the quantum gate  $U(\alpha, \beta, \gamma)$ , see Eq. (196);
- (c) compute new state:  $|\psi'\rangle = U(\alpha, \beta, \gamma)|\psi\rangle$ .

Note that all actions are allowed from each state.

- (3) **Reward space:**  $\mathbf{R} = \{\pm 1\}$ . We use sparse binary rewards obtained from the quantum measurement of the ancilla at the last episode step  $T$  using the protocol described above.

Figure 14(a) shows the control protocol applied to the qubit-ancilla system. The purple box represents one step, and is unpacked in the corresponding inset: first, the agent determines the optimal angles  $\alpha, \beta, \gamma$  and applies the control unitary; the state then undergoes the entangling noise [blue] (whose strength is random, but fixed during training), followed by an eventual spontaneous emission of the qubit [orange]. Finally, we measure the ancilla in the  $z$ -basis, and use the binary measurement output as an observation. This protocol repeats for a total of  $T = 5$  steps, before we project the qubit to the target state [green box]. To do so, we use the Hadamard- $C_{\sigma_*}$ -Hadamard sequence described above: measuring the ancilla in the  $z$ -basis then returns the correct output (as if we had measured the operator  $\sigma_*$ ) and collapses the state of the qubit; it is this measurement data that we use as a reward to train the RL agent.

We now parameterized the policy by a single-layer fully connected neural network with 8 hidden neurons; the output of the network are the three mean and standard deviation values; we can sample actions (i.e., angles) from three independent corresponding normal distributions. To train the agent, we use a batch of 2048 trajectories, over a total of 200 episodes. Figure 14(b) shows the corresponding training curves for four different types of noisy environments; (i) when the entangling noise and spontaneous emission are off ( $p_{\text{emit}} = 0, p_{\text{ent}} = 0$ ), the only source of noise is the shot noise when estimating the fidelity. We see that in this case, the agent learns to control the qubit after about 50 episodes, reaching an almost perfect fidelity. Moreover, the  $1\sigma$  uncertainty window (shaded area) shrinks as training progresses which indicates that the policy becomes more deterministic. (ii) We can then keep the spontaneous emission off and consider a weak entangling noise. This slightly reduces the achievable final fidelity to about 90%, but does not degrade the overall performance of the agent. (iii) In the opposite case, with spontaneous emission replacing the weakly entangling gate as a dominant noise source, we find that the

agent reaches about 80% fidelity. Following a stable learning stage, we notice large fluctuations in the training curve: these are presumably due to the small neural network used. We expect that training the agent longer will eventually stabilize the curve close to its maximum value. (iv) Finally, we turn all sources of noise on. The overall fidelity achieved by the agent in this case lies at around 80%. We invite the interested reader to check out the accompanying Jupyter notebook and explore the behavior of the agent further [19].

Generalizing this setup beyond single-qubit control is a topic of forefront research. For instance, there was a recent demonstration of how RL agents can discover error-correcting codes in few-qubit systems by using an ancilla to infer partial information about the state of the system [183]. Similarly, a recent experimental realization of a GKP code built from the states of a cavity resonator and manipulated using a transmon qubit as an ancilla [187], provides another real-world application of the framework introduced above. For details, we refer the interested readers to the growing body of literature on using RL for quantum error correction.

### C. Implementations of reinforcement learning in experiment

We close this section by mentioning three applications of RL in experiments with quantum systems. The field is rapidly evolving; hence, the selection below is to be taken as exemplary rather than exhaustive.

In Ref. [168] the authors used an RL agent to initialize the state of a single transmon qubit. This state preparation task requires real-time control on time scales much shorter than the coherence time of a qubit. Thus, for the purpose of efficiency, the RL agent has to be integrated with the experimental device requiring low-latency feedback. The authors achieved this by using a sub-microsecond-latency neural network on a field-programmable gate array (FPGA) which interacts directly with the transmon qubit. The agent was then trained using quantum measurements. This work laid the foundations for resolving the challenge of developing and training a reinforcement learning agent able to operate using real-time feedback.

A more difficult, yet closely related, problem is the preparation of quantum gates. Quantum gates are implemented using control pulses, much like those discussed in Sec. III. Gate synthesis can be viewed as a state preparation problem, but for a collection of a complete basis set for the Hilbert space; the challenge is to prepare the basis states with the correct relative phases among them. The agent typically has access to the control fields of a Hamiltonian, the time evolution of which generates the quantum gate. RL has proven useful in tasks involving the design of low-level pulse controls that can compensate for hardware errors. In Ref. [175], the authors demonstrate an

RL agent that can design a universal set of error-robust quantum logic gates. In particular, they also train their agent in runtime on a superconducting quantum computer to implement a cross-resonance two-qubit gate. The agent does not require knowledge of a Hamiltonian model for the system, its controls, or any underlying error processes. The pulse sequences produced by the RL agent for the cross-resonance gate are up to three times faster than state-of-the-art pulses. Moreover, RL gates were found to be robust against calibration drifts.

And finally, a generalization of the RL framework used in Sec. IV B 3 was recently used in an experiment to demonstrate quantum error correction for a cavity-QED qubit. The cavity mode defines a harmonic oscillator from which a logical qubit is built using a Gottesman-Kitaev-Preskill (GKP) code. The oscillator is coupled to a transmon which allows for readout and control of the logical qubit. The RL agent is trained to preserve the logical subspace using measurements of the code stabilizers. As a result, the trained RL agent increases the lifetime of the GKP qubit by more than a factor of 2, using an active quantum error correction protocol [187].

## V. ADVANTAGES, LIMITATIONS, AND HYBRID ALGORITHMS

As should be evident from the preceding sections, each technique has its own strengths and weaknesses: e.g., the simplicity in the formulation of counterdiabatic driving comes at the price of the often highly complex control terms it necessitates, optimal quantum control produces high-fidelity protocols so long as it does not get stuck in local minima, whereas the ability of RL agents to generalize typically requires large amounts of data. In Table V we provide our opinion on some of the key advantages and limitations for each of the approaches discussed in Secs. II–IV.

While we have presented and discussed each control technique in isolation, with an aim to provide a clear pedagogical introduction to the basic mathematical formulation of quantum control protocols, it is worth highlighting that significant benefits can be achieved by *combining* these techniques together. The effectiveness of such an approach can already be seen in fairly rudimentary settings, e.g., Sec. II F 3 where counterdiabatic driving informed the choice of operators to employ for the control, leaving the pulse shape to be optimized. More sophisticated approaches, although similar in spirit, have been used to control many-body systems by combining shortcuts-to-adiabaticity with optimal control to design more experimentally favorable control protocols. The most intuitive approach starts by constraining the available set of controllable operators. By optimizing the operators' time-dependence, this technique can allow for remarkably good performance when traversing a critical point [43,77]

so long as we stay away from the thermodynamic limit when the spectral gap closes and the norm of the AGP diverges. As we will discuss more below, variational counterdiabatic driving allows us to formalize and significantly expand this approach as shown in Ref. [46]. Beyond the clear practical advantage that hybrid approaches can provide vis-à-vis designing easier to implement control protocols, it can also allow to explore the fundamental limits of controllability. Hybrid control approaches allow determination of the true minimal control time under a given set of constraints and, therefore, explore the attainability of fundamental bounds such as the quantum speed limit and Lieb-Robinson bounds [190,191]. Significant efficiencies in the optimization procedure itself are also attainable by combining gradient-free with gradient-based optimization procedures [75], the crucial insight here being that once the need for numerical optimization enters the problem, it can be highly advantageous to ensure that a “good” initial seed is employed. This becomes particularly relevant when exploring rugged control landscapes [192–195].

For the other considered control techniques, as we have seen, both optimal control and reinforcement learning offer numerical tools to devise quantum control protocols, and a natural question arises as to when one gains an edge over the other. A major difference between RL and optimal control is that optimal control focuses on the optimal solution and any acquired information about the physical system used during optimization is lost. RL, on the other hand, extracts essential information for maximizing the reward, and stores it in the policy—a left-over product. Reinforcement learning algorithms are capable of re-using this information at a later stage, e.g., starting from a different initial state, self-correcting for exploration mistakes made during training or caused by noise and uncertainty, or even trying to control a related but different system. However, this comes at a cost: RL is not computationally as efficient as optimal control and takes more iterations to find an optimal solution, especially if it starts without prior knowledge of the controlled system (cf. model-free vs. model-based RL [151]). Another advantage of RL over established quantum control methods, which rely on built-in optimization algorithms, is the balance between exploitation of already obtained knowledge, and exploration in uncharted parts of the control landscape (a.k.a. the exploration-exploitation dilemma [151], which was mentioned also in Sec. III A). Below the quantum speed limit, decision-based exploration offers an alternative to multi-starting local gradient optimizers. Unlike such methods, the RL agent progressively becomes familiar with the control landscape, and finds protocols robust to sampling noise. Due to exploration noise, RL is not suitable for finding global minima of the cost function; in return, the minima it finds are more stable and robust to changes in the protocol. In difficult optimization problems, however, optimal control is never guaranteed to find the globally optimal protocol either, and

TABLE V. Comparative table highlighting some key advantages and limitations of the various control techniques presented throughout the tutorial and other closely related approaches.

Techniques	Advantages & focus	Limitations & Scope of applicability	References
<b>Shortcuts-to-adiabaticity—Sec. II</b>			
Counterdiabatic driving	Analytically derived which guarantees effectiveness and also allows us to provide physical insight into fundamental controllability of the system.	Only strictly applicable to exactly solvable models and often results in highly nonlocal control terms limiting direct experimental implementability.	Refs. [2,41]
Variational counterdiabatic driving	Can be approached semi-analytically or fully numerically and lends itself to a hybrid approach combined with quantum optimal control. Allows us insight into the adiabatic gauge potential in regimes where its exact calculation is difficult.	The local approximation of counterdiabatic terms tends to be poor when approaching critical and/or transition points in the spectrum, i.e., where energy gaps close.	Ref. [44,45]
<b>Quantum optimal control—Sec. III</b>			
Open-loop optimal control	Systematic approach that can be applied to any setting and allows us to impose physically motivated constraints, e.g., allowable operations, time scale etc. Benign landscape structure gives certain guarantees to find good, if not necessarily the absolutely optimal, solution.	Does not generalize as system parameters change—each configuration requires solving a new optimization problem. Numerically obtained control solutions are hard to interpret physically.	Ref. [13]
<b>Machine learning based quantum control—Sec. IV</b>			
Reinforcement learning	Versatile in scope of application and like OC can accommodate any constraints from the outset. Can generalize to unseen setups (e.g., starting from different initial states). (Typically) does not require gradients of the function to be optimized. Does not require information about the controlled quantum system.	May require a lot of training data. (Typically) does not come with formal convergence guarantees. May get stuck in a suboptimal local minimum. Does not require information about the controlled quantum system.	Ref. [15]
Other ML approaches	Train once, apply in diverse setups. Whenever optimization is time-consuming, learn the correct optimized output to avoid running it every time.	Generalization (i.e., extrapolation and interpolation) has limits and depends on the data used in training. Requires static training data distribution (i.e., characteristic system parameters should not change over time).	Refs. [147,148]

we usually have no way of verifying if it did. Anticipating possible future applications of RL in physics, we note that RL brings in three key ideas: (i) *model-free* RL requires no pre-knowledge of the controlled physical system; (ii) *adaptive* RL transfers acquired knowledge to new setups and can point to hidden relations between nonequilibrium phenomena; (iii) *autonomous* RL provides novel insights into automating complex manipulation protocols.

Hybridized procedures allow pitfalls associated with a particular approach to be circumvented, ultimately pushing the limits of controllability. As previously mentioned, a drawback of counterdiabatic driving (and associated analytical techniques) is the requirement to exactly solve the model, thus significantly limiting its range of applicability. However, it is evident that many physically relevant settings can be approximated by models with known

analytical solutions. From the above examples and discussion, it follows that using the information accessible by determining the control protocol for the analytically tractable approximation, can serve as a highly insightful seed which can then be optimized using, e.g., optimal control theory and/or machine learning. Such an approach has already been demonstrated to be highly effective in a diverse range of settings, we refer the reader to Refs. [90, 92, 93, 165, 196], for some further examples. To highlight the benefit that hybridized control strategies offer we mention three recent approaches that are showing particular promise:

- (1) *Enhanced shortcuts-to-adiabaticity (eSTA)*. The analytical nature of STA methods means that they have typically been applied to low-dimensional, idealized models. Thus, while the models serve as approximations to real physical systems, it is evident that directly applying the corresponding control terms to the actual system would, necessarily, not achieve perfect finite-time adiabatic dynamics. However, given that the formulation of STA methods guarantees that a counterdiabatic control term exists, it follows that if the Hamiltonian description of the real physical system can be considered a perturbation away from the idealized model, the correct control term is also likely to be closely related to the one derived for the idealized model. This insight served as the starting motivation for eSTA techniques [197, 198], which combines the knowledge learned from the STA methods together with gradient based optimization to achieve high fidelity control in several physically relevant settings, in particular transport [197–199] and expansion [200] in anharmonic trapping potentials, manipulation and transport in double well [201], driven qubits beyond the rotating wave approximation [197], and bosonic Josephson junctions [202]. The computed control terms have also been shown to be robust to systematic errors and noise sources [198, 202].
- (2) *Counterdiabatic local optimized driving*. The approach of variational counterdiabatic driving introduced in Sec. IIF 4 enabled the construction of approximate CD terms through minimization without the need to diagonalize the Hamiltonian [44, 45]. This minimization procedure can be tackled analytically for a limited set of scenarios, including the Ising model discussed in Sec. IIF 4, but often requires numerical minimization protocols to obtain the coefficients of the approximate CD for a particular protocol, akin to those discussed in the context of quantum optimal control in Sec. III. The most natural way to enhance the performance of variational CD protocols is to introduce higher-order CD terms with operators that involve more bodies,

i.e., by adding up to  $N$ -body interactions, but these are inherently non-local and often difficult to engineer in experiments or on quantum hardware. Note, these non-local and many-body terms are also what we are precisely wanting to avoid calculating when motivating the introduction of variational counterdiabatic driving. Counterdiabatic local optimized driving [46] takes a different approach by fixing the variational CD at low, local, orders of approximation while adding control fields that extend the family of dynamical Hamiltonians that can be explored. This allows us to use counterdiabatic optimized local driving to find the optimal dynamical path for the local approximated CD. Recent work has extended this to determine the additional control fields and associated optimal set of operators necessary for a given problem [53]. Furthermore, this approach allows one to optimize the dynamical protocol without direct simulation, by calculating the ignored non-local CD terms and minimizing their contribution, and this method could be further enhanced by combining it with the numerical approaches that have since been developed to calculate the CD terms, which we will discuss in Sec. VIF.

- (3) As a framework providing optimization algorithms capable of learning, RL presents an ideal playground for developing concepts for hybrid quantum control. Whereas RL does not magically produce high-fidelity protocols, this can be greatly facilitated by combining it with various STA methods, including CD driving, explored in Refs. [165, 169]. The key idea to apply RL to real-world noisy intermediate-scale quantum (NISQ) devices is to exploit the QAOA which defines a variational quantum circuit with unknown parameters. The RL agent has to learn either the gate angles (continuous) or the circuit architecture (discrete) and, ideally, both. Since solving both the continuous and discrete optimization involves mixed continuous-discrete action spaces which can be challenging to deal with, often one resorts to RL for one of these two problems, leaving the other to optimal control methods (Sec. III), such as gradient descent, or tree searches [203]. To design the circuit architecture ansatz, one can use the terms appearing in the nested commutator ansatz [50] for the variational AGP, see Sec. IIF 4. While this does not implement CD driving itself, these terms generate unitary dynamics that unlocks shortcuts in the Hilbert space, leading to efficient state preparation away from the adiabatic regime. Using this approach Ref. [165] developed RL agents to prepare ground states of many-body spin chains on NISQ devices, while the RL agent of Ref. [169] is trained to design optimal pulses

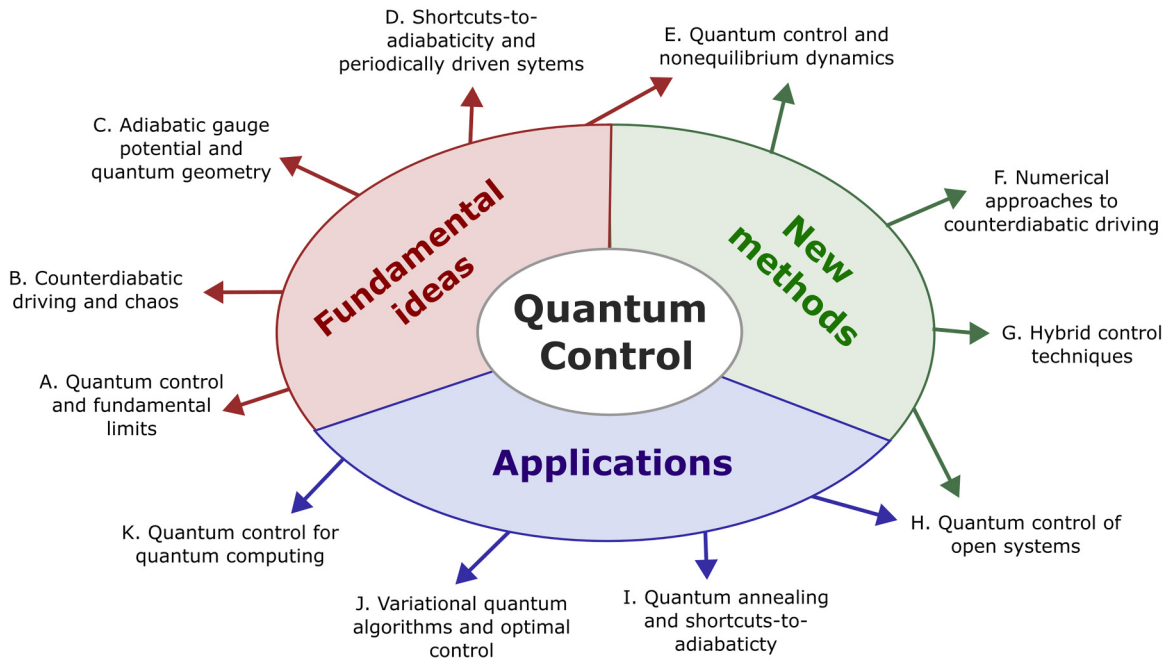


FIG. 15. Overview of the topics discussed in Sec. VI and their relation to fundamental ideas, new methods, and applications of quantum control tools.

for single-qubit gates applicable, e.g., on a superconducting qubit device. We have explored in detail such hybrid CD-RL control shortly, see Sec. IV B 2. We mention in passing that, besides CD driving, RL can be combined in a hybrid approach to optimize dynamical invariants—a different STA technique [171].

## VI. FUTURE CHALLENGES AND PROSPECTS FOR QUANTUM CONTROL

The ubiquity of control means that the techniques and tools presented above can find fertile application in a wide variety of places. In this section, we conclude this tutorial by briefly touching on several areas where quantum control is prevalent and where there are potentially significant benefits to be gained from the continued development of more advanced techniques for coherent control of complex quantum systems, see Fig. 15.

### A. Quantum control to push fundamental limits in quantum dynamics

Given a model for a quantum-dynamical system and a target state or unitary transformation, the tools of optimal control allow us to search for a protocol that achieves such a process. If the optimization systematically fails to find a solution, this could indicate that the process was not achievable, with the resources available, in the first place, thus revealing a fundamental limitation of the dynamics.

In the examples of Sec. III, this was related to restrictions placed by the quantum speed limit, and the fact that there is a minimum time required to perform a given quantum transformation, which in turn can be traced back to fundamental uncertainty relations. For quantum control problems of the form described in Sec. II, where the initial and target states can be adiabatically connected, it has been shown that quantum speed limit bounds can be significantly refined to yield more informative bounds on the systems dynamics [11,58,107,108,204–206], which in turn are linked with the geometric picture described in Sec. VIC and can help in developing alternative strategies to achieve control even for complex systems [207].

The example above suggests that optimal control could be a useful tool to reveal and probe other fundamental bounds related to quantum dynamics. One such example is the Cramér-Rao bound [127], which limits the precision with which a quantum state can act as a probe of a weakly encoded phase. Efforts towards using QOC for quantum sensing have focused on optimizing the quantum Fisher information [208–211], and related protocols based on variational quantum circuits (see also Sec. VIK) have looked at optimizing other types of metrics [212]. Instead of maximizing the sensitivity of the quantum state to an external perturbation, optimal control methods could also be applied to minimize such sensitivity, i.e., yielding *robust* quantum processes [102]. Naturally, robust quantum control is highly desirable in several applications, for instance, in quantum computing where extreme precision at the quantum level is typically required, or in quantum

thermodynamics where control techniques can be leveraged to designing protocols for the efficient charging of quantum batteries [213–218].

In the realm of many-body systems, different fundamental bounds limit the rate at which quantum information spreads through the degrees of freedom of a system. These include Lieb-Robinson bounds [219], which determine the existence of light-cone-like dynamics in a system with local interactions, as well as fast-scrambling bounds [220], related to the maximum speed of operator growth. It is an open question how driving and optimal control could be used to create protocols that scramble quantum information optimally. Even more fundamentally, optimal control tools for observable evolution (i.e., in the Heisenberg picture) have not been studied in depth up to now. Protocols that manipulate quantum information optimally could be of use for certain quantum algorithms, as well as for quantum simulation of complex models related to black hole information dynamics.

### B. Connections between counterdiabatic driving, the adiabatic gauge potential, and chaos

Adiabatic gauge potentials underlie quantum geometry, cf. Sec. VI C. The latter encodes how weakly the wavefunction changes, as measured by the fidelity susceptibility. It was pointed out that this connection can be used, e.g., to detect phase transitions without the need to define any order parameters [221,222]. Recently, it was proposed to use the sensitivity of eigenstates to external perturbations to define quantum (and classical) chaos [223]. Indeed, eigenstate deformations are quantified by the metric tensor  $g_{\lambda\mu} = \text{Re} \langle n | \mathcal{A}_\lambda \mathcal{A}_\mu | n \rangle_c$ ; in chaotic systems, one typically performs an extra average over the eigenstates which corresponds to an infinite-temperature ensemble. Hence, eigenstate deformations can be used as a probe for quantum chaos. To see this connection from a different perspective, recall that the correction to the eigenstates in the presence of perturbations is determined by the AGP  $\mathcal{A}$ . Moreover, as it satisfies the operator-valued equation  $[G(\mathcal{A}), \mathcal{H}] = 0$  [cf. Eq. (45)], one can associate adiabatic transformations (i.e., the transformations generated by the AGP) with emergent conservation laws whose generators also commute with the Hamiltonian  $\mathcal{H}$ . This leads to an intricate relationship between the AGP and integrability. Since chaos is a manifestation of the lack of integrability, it turns out it can be quantified by analyzing the stability of time-averaged trajectories to small deformations in the Hamiltonian [224]. We note in passing that chaos and ergodicity are, in general, different notions [225]; the AGP can help identify the presence of chaotic behavior. Finally, recall that the fidelity susceptibility can be measured in experiments since it is related to the response function of the system [45]; in this sense, the more slowly a system

relaxes and the stronger its sensitivity is to low-frequency response, the more chaotic it is.

### C. Adiabatic gauge potential and quantum geometry

One of the central objects of this tutorial—the adiabatic gauge potential—is inextricably related to quantum geometry [45]. Indeed, its connected correlation function in the eigenstate  $|n\rangle$  of the Hamiltonian,  $\chi_{\mu\nu} = \langle n | \mathcal{A}_\lambda \mathcal{A}_\mu | n \rangle_c$ , defines the quantum geometric tensor  $\chi_{\mu\nu}$ . The real part,  $g_{\mu\nu} = \text{Re} \chi_{\mu\nu}$ , is the metric tensor which determines the distance between states on the manifold parametrized by  $\lambda$  (the Fubini-Study metric). It can be shown that adiabatic evolution follows the geodesics (i.e., the minimal-distance curves) with respect to  $g_{\mu\nu}$ . The metric tensor is also known as the quantum Fisher information or fidelity susceptibility, and finds ample applications in quantum parameter estimation and quantum sensing. On the other hand, the imaginary part of the quantum metric tensor is proportional to the Berry curvature  $F_{\mu\nu} = -2 \text{Im} \chi_{\mu\nu}$ . It is often used to determine the Chern number—a topological invariant that lies at the heart of understanding topological insulators, and is invoked in the explanation of the integer quantum Hall effect [47]. Moreover, the eigenstate expectation value of the gauge potential,  $\langle n | \mathcal{A}_\lambda | n \rangle$  is the Berry connection which defines parallel transport on the parameter manifold; it is frequently used to compute the geometric (Berry) phase accumulated by the wavefunction upon parallel transport along a closed path in the parameter manifold. Unlike the Berry connection, both the quantum metric and the Berry curvature are invariant under changing the phase of the wavefunction, and are therefore measurable quantities. In particular, the metric tensor is related to the spectral function via linear response theory, while the Berry curvature can be measured from the leading order nonadiabatic response of observables [45].

### D. Shortcuts-to-adiabaticity and periodically driven (Floquet) systems

As we discussed in Sec. II F 4, real-valued Hamiltonians have purely imaginary-valued adiabatic gauge potentials. For non-interacting systems, these take the form of  $y$ -magnetic fields or current terms; interactions modify these currents and can even make them less local (think of interaction-dependent hopping terms). Thus, implementing even approximate CD driving requires experimental control over such generalized current operators. Moreover, in systems with broken time-reversal symmetry, the Hamiltonian is complex-valued and even more complicated terms appear in the AGP. It is therefore important to find ways to implement the AGP in experiments, at least approximately. One way to do this is to use periodic (Floquet) drives [226–231]. Recently, it was shown that the AGP can be generated by taking nested commutators of

the non-driven Hamiltonian  $\mathcal{H}(\lambda = 0)$  and the generalized force  $\partial_\lambda \mathcal{H}(\lambda)$  [50]. Incidentally, the same nested commutators appear in the Floquet-Magnus expansion (or its step-drive version, the Baker-Campbell-Hausdorff expansion), used to find approximations to the time-independent Floquet Hamiltonian  $\mathcal{H}_F$  which governs the stroboscopic dynamics of periodically driven systems. This opens up the intriguing possibility of implementing an effective stroboscopic dynamics generated by  $\mathcal{H}_F = \mathcal{H} + \dot{\lambda} \mathcal{A}_\lambda$  [50] [a related idea was used to enhance energy minimization algorithms on digital quantum devices using CD driving, giving rise to the CD quantum approximate optimization algorithm (CD-QAOA) [90]]. A proof of principle realization of counterdiabatic driving using periodic drives has recently been reported in an experiment using NV centers [232].

Floquet drives, on their own, are typically used to alter the properties of static systems via so-called Floquet engineering. The key idea is to use a strong time-periodic drive to dress the states of the system; the resulting stroboscopic dynamics generated by the Floquet Hamiltonian can be used to explore dynamically states and/or parameter regimes accessible only in so-called synthetic quantum matter (e.g., ultracold or Rydberg atoms, trapped ions, superconducting qubits, etc.) [226–231]. However, using periodic drives as an engineering toolbox brings up the question of how to control the population of stationary states in the presence of the periodic drive that created them in the first place; commonly used protocols in Floquet engineering include ramps of the drive amplitude, modulation of the drive frequency (so-called chirps), or protocols changing the phase of the periodic drive. Floquet control is a daunting challenge since any external control will violate the periodicity of the Hamiltonian and render Floquet’s theorem inapplicable; moreover, due to drive-induced resonances, the adiabatic limit for the effective Floquet Hamiltonian does not exist (though there can be adiabatic regimes at small but finite ramp speeds) [37,38]. Nevertheless, the concept of local CD driving can be extended to Floquet states: besides the contribution from the Floquet eigenstates, the AGP associated with the Floquet Hamiltonian contains an additional term due to the micromotion dynamics (i.e., the dynamics in between the stroboscopic times). Curiously, there also exists a generalization of the variational principle we derived in Sec. II E [36]. In that sense, finding the Floquet Hamiltonian (i.e., applying Floquet engineering) can be cast as an inverse counterdiabatic problem [233]. Floquet systems demonstrate that the ideas of CD driving are more general and apply to systems taken far out of equilibrium; moreover, CD driving proves instrumental in telling inherently non-equilibrium effects apart from equilibrium-like phenomena. Indeed, CD driving offers an entirely new, geometric perspective on Floquet theory [233]. The question of how to efficiently control generic nonequilibrium steady

states is one of the open frontiers of present-day quantum dynamics.

### E. Quantum control to characterize non-equilibrium dynamics

While the principal goal of quantum control is to achieve a particular aim, it is useful to acknowledge that this paradigm can be inverted and the control of a system can be used as a remarkably versatile tool for understanding the underlying physics. In the case of control via counterdiabatic techniques the intuition is simple: the control allows us to circumvent or suppress the non-equilibrium dynamics; therefore, the control itself contains all the relevant information about the dynamical response of the system to an, in principle, arbitrary perturbation. This can readily be seen for example in assessing characteristics of the controlled dynamics which allows us to study the Kibble-Zurek mechanism [35,61]. Furthermore, while counterdiabatic driving predicts that highly complex and non-local terms are generally necessary for many-body systems, the techniques outlined previously provide a versatile toolbox in order to build up the control while respecting given experimental constraints. This not only allows for these techniques to be scaled up, and thus enhance their practical utility, it also provides us with a framework to consider the relevance of particular interactions [46,51] or to determine the minimal operator set necessary for control [234], ultimately providing valuable insight into the fundamental controllability of complex systems.

### F. Numerical approaches to the calculation of counterdiabatic driving and the adiabatic gauge potential

In Sec. II, we exposed the reader to a number of approaches for obtaining the adiabatic gauge potential (AGP), or counterdiabatic driving term, which were analytic in nature. The final variational approach [44] outlined in Sec. II F 4 was until relatively recently the state-of-the-art, with the advancement of the known form of the ansatz given in Ref. [50]. However, with the advancement of adiabatic quantum computing, and the later connection of the AGP to the digital variational QAOA [90], there was growing interest in developing numerical procedures to calculate or approximate the AGP in many-body quantum systems beyond that which can be tackled with the techniques outlined in Sec. II and which are complex systems from the point of view of optimal control. To date, this has led to three approaches which tackle this issue from different angles. The orthogonal commutator expansion method [51] takes a similar approach to that of the variational adiabatic gauge potential of Sec. II F 4 but recasts solving for the AGP into that of solving a block-tridiagonal system of equations, which can be solved

efficiently numerically. This method is particularly targeted at obtaining the exact AGP numerically in complex scenarios. The second method is a Krylov approach which uses a standard approach to solving large quantum many-body problems numerically by introducing a Krylov basis which spans the minimal operator subspace to describe the AGP [50,234]. This method is particularly useful for finding approximations of the AGP at moderate orders of the commutator expansion of Ref. [50] for complicated many-body Hamiltonians. The third method uses matrix product operators—a tensor network technique [235,236], and allows one to compute the AGP in large quantum systems. It can outperform both the orthogonal commutator and the Krylov ansatz methods, although it is currently unclear if this advantage is generic. The three methods are equivalent in the limit in which they all solve for the exact AGP, if they can reach that limit, in a given problem. However, as all three methods have only recently been introduced, it is unclear how the approximate forms of each compare and this is a current area of active research.

### G. Hybrid control techniques

As discussed in Sec. V, it is clear that hybridized protocols offer the most versatile approach to achieve coherent control. The performance of such techniques also indicates that determining the ultimate limits of controllability is a multifaceted problem and one which will require input from the various sub-communities. This is particularly relevant as quantum devices continue to be developed and there is a particular need for effective and robust control. Indeed, the utility and scalability of novel quantum technologies is likely to be strongly impacted by our ability to coherently manipulate the components of the devices, particularly in the presence of unwanted, or potentially unknown, sources of noise, as will be discussed in the proceeding Sec. VIH.

### H. Quantum control of open systems

A simple fact is that no system is ever truly shielded from the surrounding environment and indeed virtually all experimental setups go to great lengths to offset, or at least mitigate, the unavoidable interaction between a system and, e.g., stray electric and magnetic fields or unwanted seismic activity (such as a subway passing by), etc. If the timescales of the noise are sufficiently slow compared to the system's evolution or if they host so-called decoherence-free subspaces or can be dynamically decoupled from the environment [237], then often the control protocol for the closed unitary dynamics are put to work, i.e., precisely the techniques that have been the focus of this tutorial. Indeed, speeding up the unitary dynamics such that it achieves a given evolution faster than the noise can act is one of the core drivers in the development of quantum control techniques. However,

increasingly quantum control is becoming an integral part of emerging quantum technologies themselves and this has resulted in a steady growth in examining control of genuinely open quantum systems [20]. The simplest such studies involve employing the same Hamiltonian control techniques already discussed while explicitly in the presence of unwanted and/or unavoidable environmental noise commonly modeled using phenomenological approaches, particularly Markovian master equations [20]. Control strategies, particularly those following the QOC theory and ML type prescriptions, are determined while explicitly including the open system effects as part of the optimization, aiming to find the optimal path in state space that minimizes the impact of the environment.

The most general picture for control is when both the system's Hamiltonian and interaction with its environment can be manipulated, so-called reservoir engineering approaches, see, e.g., Ref. [238]. Such a setting encompasses and extends the above to the non-Markovian regime which, owing to the inherent complexity of characterizing and understanding such open system dynamics, means comparatively little is known regarding the typically achievable levels of control. For control techniques based on shortcuts-to-adiabaticity the issue becomes perhaps even more fundamental since the very definition of adiabaticity must be carefully reconsidered when the system exchanges energy with the environment [239,240]. Nevertheless some progress in extending shortcuts-to-adiabaticity to open systems has been made [241,242]. In addition to these general considerations, it is worth noting that NISQ devices have recently opened up a new kind of non-unitary way to control quantum systems: these so-called feed-forward protocols rely on mid-circuit projective measurements the results of which are fed back to the system and determine the subsequent structure of the applied unitary circuit, giving rise to a mixed quantum-classical control. If accessible, feed-forward protocols can provide a way to cut down the preparation time for various many-body states since they can evade some constraints on information spreading imposed by locality, such as Lieb-Robinson bounds or quantum speed limits [243–248].

The complexity associated with developing general control protocols that are either robust to, or leverage, environmental noise indicates that hybrid approaches are likely to be the most fruitful path forward and therefore a greater integration of the various control sub-communities is vital moving forward.

### I. Quantum annealing and shortcuts-to-adiabaticity

Quantum annealers are analog quantum devices designed to access low-energy states of arbitrary Hamiltonians on specific interaction graphs which can encode solutions to practical problems, for example, in combinatorial optimization [28]. These devices largely rely on

being able to adiabatically transform simple Hamiltonians to more complex ones, in the way described in Sec. II B. The success of annealing protocols is predominantly limited by the coherence time of the device, and the scaling of the minimum energy gap for a given problem. Counterdiabatic driving can be a useful tool to lower the total evolution time used by a control protocol, to make it fit within the coherence time of the device. These tools have recently begun to be implemented in quantum annealing prototypes, as well as their gate-based extensions (like QAOA, see Sec. VI K).

Problems, where the minimum energy gap is too small or scales unfavorably with system size, represent a major challenge for quantum annealers. From a physics perspective, it is known that minimum gaps scaling polynomially with system size are typically associated with second-order (continuous) phase transitions, while exponentially vanishing gaps are linked to first-order (discontinuous) phase transitions [249]. To mitigate the gap scaling, it has been proposed to include so-called catalysts in the annealing protocol [250] or employing augmented counterdiabatic techniques [251]. These are additional terms in the Hamiltonian that provide intermediate stops in the annealing schedule. It has been shown that such strategies can substantially improve gap scaling. Unfortunately, it is hard to know *a priori* what kind of catalyst Hamiltonian will be needed for a given problem. Connecting CD driving to the catalyst formalism could help provide a more formal ground to these tools, and improve their usefulness [252]. More generally, the properties of CD protocols could be used to probe interesting connections between computational complexity, gap scaling, and quantum phase transitions (see also the discussion in Sec. VI E).

### J. Variational quantum algorithms and quantum optimal control

More generally, an important connection has been recognized between the formalism of quantum optimal control and variational quantum algorithms (VQAs) [253], which have led the charge in the search of useful applications of near-term NISQ computers. In a generic VQA, a circuit is formed of a series of parameter-dependent gates  $U_i(\{\theta_k\})$ , which act upon a fiduciary state, typically  $|0\rangle^{\otimes N}$ . Measuring the qubits at the end of the circuit allows us to evaluate some cost function, for instance, associated with the energy of a target model. This cost function is then fed to a classical optimizer, which searches for a new set of parameters  $\{\theta_k\}$  to try in the next cycle of the hybrid quantum-classical procedure.

The VQA framework is the foundation for potential applications such as the QAOA, the variational quantum eigensolver (VQE), and many quantum machine learning (QML) approaches. The connection between the QOC formalism laid out in Sec. III and VQAs is evident; the

main difference being that the quantum evolution necessary to evaluate the cost function is implemented in an actual quantum processor in the latter, instead of numerically simulated classically in the former. VQAs have been intensely explored in recent years and have revealed important properties of what we could expect in quantum optimal control when scaling up system size and in the presence of noise [254,255]. An example of these is the phenomenon of barren plateaus, which predict that certain highly expressible (i.e., controllable [94]) quantum circuits lead to cost function gradients that decrease exponentially with system size, meaning that the resulting optimization being performed by the classical computing resource becomes prohibitive at large scales. The presence of noise has been shown to aggravate this problem. Workarounds to this issue imply the use of less expressible VQA ansätze, which in quantum control language set the stage for studying complex quantum systems that are not fully, but partially controllable [256].

### K. Quantum control for quantum computing

As we have mentioned throughout this tutorial, quantum optimal control already plays a major role in the development of quantum computers, as it has allowed to improve the fidelities of entangling gates in many physical platforms.

One prominent example is the use of the so-called cross resonance (CR) gate [257,258]. This has been a driving protocol behind many advances in quantum computing based on superconducting technology [259]. The CR gate provides a protocol that requires only microwave control. This reduces the need for magnetic flux lines in superconducting circuit designs that could otherwise become sources of additional noise. By employing various types of optimization, very high fidelities for the CR gates have been achieved [260,261], as well as extension of the CR gate based on the same including the so-called cross-cross resonance gate [262]. This work has used optimal control in a closed loop that takes into account the hardware specifics [144]. At this point, the high CR gate fidelities are typically measured in systems with a limited number of qubits [263], and it is not given that the control protocols used here will also apply as one scales the system to a large number of qubits. This has led to proposals for new types of pulse shaping in terms of DRAG protocols designed to work in the presence of a large number of qubits [264].

Other major commercial team building superconducting quantum computers have also made strides towards high-fidelity gates by using quantum control. New ways to perform DRAG pulse shaping have been implemented, resulting in large reductions in single qubit gate errors for transmon systems [265]. Furthermore, new approaches to reduce errors associated with readout and reset using

quantum optimal control can be found in Ref. [266] with key examples from transmon systems.

A drawback of transmon systems is their low anharmonicity in the spectrum that leads to physical limitations on the speed of quantum gates. This has led to the pursuit of other qubit modalities, and a prominent recent example is flux qubits that may allow for very fast gates. The regime of interest is influenced by driving pulses in several important ways, including inducing ac Stark shifts of the energy levels and unintended leakage to levels outside the computational (two-level) space. Countering these effects is a main venue for pulse shaping techniques. Experimental work [267] has shown how to mitigate these effects and achieve very high single-qubit gate fidelities. However, this comes at the cost of additional hardware requirements. In a very recent paper [268] it was shown how one may avoid hardware complexity by rethinking the pulse shaping schemes. This requires a reworking of the rotating wave approximation that is applied widely in the superconducting circuit community, but that is not expected to hold when targeting very fast gates.

For trapped ion quantum computing, a major ingredient in many experimental implementations is the entangling gate proposed by Sørensen and Mølmer [269]. This has been hugely successful as a protocol for achieving high-fidelity gates in ion trap quantum computers [270–272]. In a recent work by Kirchhoff and co-workers [273], a new theoretical analysis of the scheme in Ref. [269] shows that there may be errors that have not previously been accounted for and that when included these properly, one finds a lower limit on the error that matches well with recent experimental errors in the different ion trap quantum computing implementations. The work of Ref. [273] then goes on to suggest how to proceed using control techniques to further reduce errors using, e.g., pulse shaping.

A key point in which control has yielded decisive benefits for quantum computing has been in the optimization of performance for near-term noisy quantum devices. In the context of variational quantum circuits (see the previous section), this has led to exploration of how to use parametrized quantum gates to shorten circuit depth and reduce noise for increased performance. One such effort has been directed at quantum computing to simulate the behavior of fermionic systems in Ref. [274], that took advantage of fast adiabatic control [275]. Further work has since demonstrated how these features can be used for quantum chemistry and quantum materials calculations on quantum computers [276,277]. Parametrized two-qubit gates are expected to be the better option for these quantum simulation tasks on quantum computers [278,279] and it is expected to bring further challenges and opportunities for quantum control [280].

Beyond the two-qubit gates, there are many instances in which gates acting on three or more qubits can be a major benefit in the implementation of quantum algorithms

on quantum computers [280]. The controlled operations of basic algorithms such as Grover search and other oracle-based designs, as well as key steps in Shor’s factoring algorithm, will imply an overhead in decomposition into single- and two-qubit gates. Even in the presence of error correction, this overhead will influence the execution time as the problem size grows. Here, the availability of  $N$ -qubit operations can be used to reduce the overhead, and in many cases by significant amounts. Furthermore, three-qubit gates are crucial ingredients of quantum error correction protocols [281], and as such are highly desirable operations to have in any quantum computing architecture with low overhead and high quality. Hence, major efforts have been put into realizing multi-qubit gates with high fidelity.

Optimized multi-qubit quantum gates have been considered in several platforms. As an example, the three-qubit Toffoli gate has been considered using a single-shot protocol within superconducting transmon qubits [282]. This can be generalized to multi-qubit gates beyond three qubits such as the  $n$ -bit Toffoli or related controlled gates that may be realized using Rydberg atoms [283], trapped ions [284], and superconducting circuits [283–285]. Other controlled three-qubit gates realized with quantum optimal control have been proposed that target time-optimal implementations [286] in neutral atoms that may be used for quantum computing [287] and quantum simulation tasks [288]. For trapped ion quantum computing, the Grover search algorithm has been realized on a small set of qubits [289] using three- and four-qubit gates [290].

The topic of multiqubit gates is intimately connected to the exploration of systems beyond qubits, i.e., systems in which the local Hilbert space has dimension  $d > 2$ , so-called qudits. One may also view this from the point of view of breaking the abstraction that the assumption of having qubits in the system binds us to [280]. Here one can implement quantum algorithms in a more versatile fashion [291,292]. The gates acting on qudits are more involved and can present new challenges for control techniques as multiple levels are now active and it may be advantageous to activate several transitions at once to realize qudit gates [280,293]. This may in turn introduce new issues as there are now multiple ways in which leakage can occur and it is can be more challenging to design quantum control protocols to mitigate these effects. Using different types of protocols it has by now been possible to realize qudit-based quantum processors in several platforms, including cold atoms [294], superconducting circuits [295], trapped ions [296], and photonics [297,298].

There are other important connections between quantum control and quantum computing development which have been more recently proposed, and are currently being exploited. A natural extension is to use quantum control to perform quantum tasks that are specific to different tasks in the quantum computing stack. For instance, developing

TABLE VI. Summary of abbreviations used throughout.

Abbreviation	Definition
2LS	two-level system
AC	actor-critic
AGP	adiabatic gauge potential
BEC	Bose-Einstein condensate
BFGS	Broyden-Fletcher-Goldfarb-Shannon
CD	counterdiabatic
CRAB	chopped randomized basis
DDGP	deep deterministic policy gradient
DQN	deep $Q$ -network
DRAG	derivative removal by adiabatic gate
FPGA	field-programmable gate array
GKP	Gottesman-Kitaev-Preskill
GRAPE	gradient-ascent pulse engineering
L-BFGS-B	limited memory variant of BFGS
LMG	Lipkin-Meshkov-Glick
LZ	Landau-Zener
ML	machine learning
NISQ	noisy intermediate-scale quantum
NMR	nuclear magnetic resonance
NV	nitrogen-vacancy
PG	policy gradient
PPO	proximal policy optimization
QAOA	quantum approximate optimization algorithm
QML	quantum machine learning
QOC	quantum optimal control
QSL	quantum speed limit
RL	reinforcement learning
STA	shortcuts to adiabaticity
STIRAP	stimulated rapid adiabatic passage
TRPO	trust-region policy optimization
VQA	variational quantum algorithm
VQE	variational quantum eigensolver

fault-tolerant quantum algorithms can benefit from using optimal control to generate magic states [299] or parity check unitaries [300] for quantum error correction, developing robust implementations of quantum gates [301], and for the manipulation of surface code defects [302]. It has also been demonstrated that quantum optimal control can be used to synthesize gates for digital quantum simulation circuits [135,303]. The lessons learned from STA, the impact of ML to train for general scenarios, and, importantly, the framework of quantum optimal control are crucial for the development and optimization of quantum error suppression, mitigation, and correction which are a critical part of the strategies and roadmaps being pursued throughout the quantum computing ecosystem, in both experiments and theory, across government, academia, and the growing quantum computing industry.

### JUPYTER NOTEBOOKS

Jupyter notebooks are available on GitHub [19] for interested readers to explore different code parameters in

the Examples discussed in the main text. These notebooks can also be used as a starting point to apply the control techniques introduced in this tutorial to open problems at the forefront of research.

### ACKNOWLEDGMENTS

We thank all the wonderful people we have worked with on quantum control over the years. All authors thank Andrew J. Daley, Christiane Koch, and Anatoli Polkovnikov for helpful discussions and advice throughout the writing of this tutorial. S.C. is particularly grateful to Eoin Carolan, Anthony Kiely, and George Mihailescu for valuable discussions and feedback. M.B. thanks Paul M. Schindler and F. Balducci for numerous related discussions. C.W.D. thanks Ewen D.C. Lawrence, Ieva Čepaitė, and Stewart Morawetz for numerous related discussions. P.M.P. is grateful to Anupam Mitra, Nathan Lysne, Sivaprasad Omanakuttan, Poul Jessen, and Ivan Deutsch for many helpful optimal control discussions over the years. N.T.Z. would like to thank Stig Elkjær Rasmussen, Kasper Sangild Christensen, Kasper Poulsen, and in particular Alan Costa dos Santos for collaboration and discussions on projects involving quantum control. S.C. acknowledges support from the Science Foundation Ireland Starting Investigator Research Grant “SpeedDemon” No. 18/SIRG/5508 and the Alexander von Humboldt Foundation. P.M.P. acknowledges support by U.S. National Science Foundation (Grant No. PHY-2210013). C.W.D. was supported by the Engineering and Physical Sciences Research Council through Grant No. EP/Y005058/2. M.B. was funded by the European Union (ERC, QuSimCtrl, 101113633). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. NTZ is supported in part by the Novo Nordisk Foundation, the Innovation Fund Denmark, as well as the Independent Research Fund Denmark.

### DATA AVAILABILITY

The data and the notebooks associated with this manuscript version are openly available on Zenodo [304].

- [1] C. P. Koch, U. Boscain, T. Calarco, G. Dirr, S. Filipp, S. J. Glaser, R. Kosloff, S. Montangero, T. Schulte-Herbrüggen, D. Sugny, and F. K. Wilhelm, Quantum optimal control in quantum technologies. Strategic report on current status, visions and goals for research in Europe, *EPJ Quantum Technol.* **9**, 19 (2022).
- [2] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, Shortcuts

- to adiabaticity: Concepts, methods, and applications, *Rev. Mod. Phys.* **91**, 045001 (2019).
- [3] I. Bloch, J. Dalibard, and W. Zwerger, Many-body physics with ultracold gases, *Rev. Mod. Phys.* **80**, 885 (2008).
- [4] A. M. Kaufman and K.-K. Ni, Quantum science with optical tweezer arrays of ultracold atoms and molecules, *Nat. Phys.* **17**, 1324 (2021).
- [5] U. Boscain, M. Sigalotti, and D. Sugny, Introduction to the Pontryagin maximum principle for quantum optimal control, *PRX Quantum* **2**, 030203 (2021).
- [6] L. T. Brady, C. L. Baldwin, A. Bapat, Y. Kharkov, and A. V. Gorshkov, Optimal protocols in quantum annealing and quantum approximate optimization algorithm problems, *Phys. Rev. Lett.* **126**, 070505 (2021).
- [7] B. Kraus and J. I. Cirac, Optimal creation of entanglement using a two-qubit gate, *Phys. Rev. A* **63**, 062309 (2001).
- [8] M. H. Goerz, T. Calarco, and C. P. Koch, The quantum speed limit of optimal controlled phasegates for trapped neutral atoms, *J. Phys. B: At. Mol. Opt. Phys.* **44**, 154011 (2011).
- [9] P. Watts, J. Vala, M. M. Müller, T. Calarco, K. B. Whaley, D. M. Reich, M. H. Goerz, and C. P. Koch, Optimizing for an arbitrary perfect entangler. I. Functionals, *Phys. Rev. A* **91**, 062306 (2015).
- [10] P. Tashev, S. Petrov, F. Metz, and M. Bukov, Reinforcement learning to disentangle multiqubit quantum states from partial observations, [arXiv:2406.07884](https://arxiv.org/abs/2406.07884).
- [11] M. Bukov, D. Sels, and A. Polkovnikov, Geometric speed limit of accessible many-body state preparation, *Phys. Rev. X* **9**, 011034 (2019).
- [12] S. Deffner and S. Campbell, Quantum speed limits: From Heisenberg’s uncertainty principle to optimal quantum control, *J. Phys. A: Math. Theor.* **50**, 453001 (2017).
- [13] 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).
- [14] D. Stefanatos and E. Paspalakis, A shortcut tour of quantum control methods for modern quantum technologies, *Europhys. Lett.* **132**, 60001 (2020).
- [15] L. Giannelli, S. Sgroi, J. Brown, G. S. Paraoanu, M. Paternostro, E. Paladino, and G. Falci, A tutorial on optimal control and reinforcement learning methods for quantum technologies, *Phys. Lett. A* **434**, 128054 (2022).
- [16] T. Hatomura, Shortcuts to adiabaticity: Theoretical framework, relations between different methods, and versatile approximations, *J. Phys. B: At. Mol. Opt. Phys.* **57**, 102001 (2024).
- [17] Q. Ansel, E. Dionis, F. Arrouas, B. Peaudecerf, S. Guérin, D. Guéry-Odelin, and D. Sugny, Introduction to theoretical and experimental aspects of quantum optimal control, *J. Phys. B: At. Mol. Opt. Phys.* **57**, 133001 (2024).
- [18] S. Deffner and M. V. S. Bonança, Thermodynamic control—an old paradigm with new applications, *Europhys. Lett.* **131**, 20001 (2020).
- [19] Jupyter notebooks can be found on GitHub: <https://github.com/nqd-lab/quctrl-tutorial>.
- [20] C. P. Koch, Controlling open quantum systems: Tools, achievements, and limitations, *J. Phys.: Condens. Matter* **28**, 213001 (2016).
- [21] O. V. Ivakhnenko, S. N. Shevchenko, and F. Nori, Nonadiabatic Landau–Zener–Stückelberg–Majorana transitions, dynamics, and interference, *Phys. Rep.* **995**, 1 (2023).
- [22] 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).
- [23] L. D’Alessio, Y. Kafri, A. Polkovnikov, and M. Rigol, From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics, *Adv. Phys.* **65**, 239 (2016).
- [24] M. Born and V. Fock, Beweis des adiabatsatzes, *Z. Phys.* **51**, 165 (1928).
- [25] T. Kato, On the adiabatic theorem of quantum mechanics, *J. Phys. Soc. Jpn.* **5**, 435 (1950).
- [26] J. C. Budich and B. Trauzettel, From the adiabatic theorem of quantum mechanics to topological states of matter, *Phys. Status Solidi RRL* **7**, 109 (2013).
- [27] M. V. Berry, Quantal phase factors accompanying adiabatic changes, *Proc. R. Soc. Lond. A* **392**, 45 (1984).
- [28] T. Albash and D. A. Lidar, Adiabatic quantum computation, *Rev. Mod. Phys.* **90**, 015002 (2018).
- [29] A. Polkovnikov, Universal adiabatic dynamics in the vicinity of a quantum critical point, *Phys. Rev. B* **72**, 161201(R) (2005).
- [30] R. Schützhold, M. Uhlmann, Y. Xu, and U. R. Fischer, Sweeping from the superfluid to the Mott phase in the Bose-Hubbard model, *Phys. Rev. Lett.* **97**, 200601 (2006).
- [31] M. Uhlmann, R. Schützhold, and U. R. Fischer, Vortex quantum creation and winding number scaling in a quenched spinor Bose gas, *Phys. Rev. Lett.* **99**, 120407 (2007).
- [32] J. Dziarmaga, Dynamics of a quantum phase transition: Exact solution of the quantum Ising model, *Phys. Rev. Lett.* **95**, 245701 (2005).
- [33] B. Damski, The simplest quantum model supporting the Kibble-Zurek mechanism of topological defect production: Landau-Zener transitions from a new perspective, *Phys. Rev. Lett.* **95**, 035701 (2005).
- [34] A. del Campo and W. H. Zurek, Universality of phase transition dynamics: Topological defects from symmetry breaking, *Int. J. Mod. Phys. A* **29**, 1 (2014).
- [35] R. Puebla, S. Deffner, and S. Campbell, Kibble-Zurek scaling in quantum speed limits for shortcuts to adiabaticity, *Phys. Rev. Res.* **2**, 032020(R) (2020).
- [36] P. M. Schindler and M. Bukov, Counterdiabatic driving for periodically driven systems, *Phys. Rev. Lett.* **133**, 123402 (2024).
- [37] A. Eckardt and M. Holthaus, Avoided-level-crossing spectroscopy with dressed matter waves, *Phys. Rev. Lett.* **101**, 245302 (2008).
- [38] P. Weinberg, M. Bukov, L. D’Alessio, A. Polkovnikov, S. Vajna, and M. Kolodrubetz, Adiabatic perturbation theory and geometry of periodically-driven systems, *Phys. Rep.* **688**, 1 (2017).
- [39] M. Demirplak and S. A. Rice, Adiabatic population transfer with control fields, *J. Chem. Phys. A* **107**, 9937 (2003).
- [40] M. Demirplak and S. A. Rice, On the consistency, extremal, and global properties of counterdiabatic fields, *J. Chem. Phys.* **129**, 154111 (2008).

- [41] M. Berry, Transitionless quantum driving, *J. Phys. A: Math. Theor.* **42**, 365303 (2009).
- [42] A. del Campo, M. M. Rams, and W. H. Zurek, Assisted finite-rate adiabatic passage across a quantum critical point: Exact solution for the quantum Ising model, *Phys. Rev. Lett.* **109**, 115703 (2012).
- [43] S. Campbell, G. De Chiara, M. Paternostro, G. M. Palma, and R. Fazio, Shortcut to adiabaticity in the Lipkin-Meshkov-Glick model, *Phys. Rev. Lett.* **114**, 177206 (2015).
- [44] D. Sels and A. Polkovnikov, Minimizing irreversible losses in quantum systems by local counterdiabatic driving, *Proc. Natl. Acad. Sci. U.S.A.* **114**, E3909 (2017).
- [45] M. Kolodrubetz, D. Sels, P. Mehta, and A. Polkovnikov, Geometry and non-adiabatic response in quantum and classical systems, *Phys. Rep.* **697**, 1 (2017).
- [46] I. Čepaitė, A. Polkovnikov, A. J. Daley, and C. W. Duncan, Counterdiabatic optimized local driving, *PRX Quantum* **4**, 010312 (2023).
- [47] B. Bradlyn and M. Iraola, Lecture notes on Berry phases and topology, *SciPost Phys. Lect. Notes* **51** (2022).
- [48] C. Jarzynski, Generating shortcuts to adiabaticity in quantum and classical dynamics, *Phys. Rev. A* **88**, 040101(R) (2013).
- [49] E. Carolan, A. Kiely, and S. Campbell, Counterdiabatic control in the impulse regime, *Phys. Rev. A* **105**, 012605 (2022).
- [50] P. W. Claeys, M. Pandey, D. Sels, and A. Polkovnikov, Floquet-engineering counterdiabatic protocols in quantum many-body systems, *Phys. Rev. Lett.* **123**, 090602 (2019).
- [51] E. D. Lawrence, S. F. Schmid, I. Čepaitė, P. Kirton, and C. W. Duncan, A numerical approach for calculating exact non-adiabatic terms in quantum dynamics, *SciPost Phys.* **18**, 014 (2025).
- [52] T. Hatomura and K. Takahashi, Controlling and exploring quantum systems by algebraic expression of adiabatic gauge potential, *Phys. Rev. A* **103**, 012220 (2021).
- [53] S. Morawetz and A. Polkovnikov, Efficient paths for local counterdiabatic driving, *Phys. Rev. B* **110**, 024304 (2024).
- [54] O. Abah, R. Puebla, A. Kiely, G. De Chiara, M. Paternostro, and S. Campbell, Energetic cost of quantum control protocols, *New J. Phys.* **21**, 103048 (2019).
- [55] Y. Zheng, S. Campbell, G. De Chiara, and D. Poletti, Cost of counterdiabatic driving and work output, *Phys. Rev. A* **94**, 042132 (2016).
- [56] A. C. Santos and M. S. Sarandy, Superadiabatic controlled evolutions and universal quantum computation, *Sci. Rep.* **5**, 15775 (2015).
- [57] E. Calzetta, Not-quite-free shortcuts to adiabaticity, *Phys. Rev. A* **98**, 032107 (2018).
- [58] K. Funo, J.-N. Zhang, C. Chatou, K. Kim, M. Ueda, and A. del Campo, Universal work fluctuations during shortcuts to adiabaticity by counterdiabatic driving, *Phys. Rev. Lett.* **118**, 100602 (2017).
- [59] O. Abah and E. Lutz, Energy efficient quantum machines, *Europhys. Lett.* **118**, 40005 (2017).
- [60] S. Campbell and S. Deffner, Trade-off between speed and cost in shortcuts to adiabaticity, *Phys. Rev. Lett.* **118**, 100601 (2017).
- [61] S. Campbell, Quantum work statistics of controlled evolutions, *Europhys. Lett.* **143**, 68001 (2023).
- [62] C. W. Duncan, Exact counterdiabatic driving in finite topological lattice models, *Phys. Rev. B* **109**, 245421 (2024).
- [63] A. Hartmann and W. Lechner, Rapid counter-diabatic sweeps in lattice gauge adiabatic quantum computing, *New J. Phys.* **21**, 043025 (2019).
- [64] S. Sachdev, Quantum phase transitions, *Phys. World* **12**, 33 (1999).
- [65] B. Damski and W. H. Zurek, Adiabatic-impulse approximation for avoided level crossings: From phase-transition dynamics to Landau-Zener evolutions and back again, *Phys. Rev. A* **73**, 063405 (2006).
- [66] B. Damski and M. M. Rams, Exact results for fidelity susceptibility of the quantum Ising model: The interplay between parity, system size, and magnetic field, *J. Phys. A: Math. Theor.* **47**, 025303 (2013).
- [67] B. Damski, Counterdiabatic driving of the quantum Ising model, *J. Stat. Mech.* **2014**, P12019 (2014).
- [68] G. B. Mbeng, A. Russomanno, and G. E. Santoro, The quantum Ising chain for beginners, *SciPost Phys. Lect. Notes* **82** (2024).
- [69] S. Dusuel and J. Vidal, Finite-size scaling exponents of the Lipkin-Meshkov-Glick model, *Phys. Rev. Lett.* **93**, 237204 (2004).
- [70] S. Dusuel and J. Vidal, Continuous unitary transformations and finite-size scaling exponents in the Lipkin-Meshkov-Glick model, *Phys. Rev. B* **71**, 224420 (2005).
- [71] P. Cejnar, P. Stránský, M. Macek, and M. Kloc, Excited-state quantum phase transitions, *J. Phys. A: Math. Theor.* **54**, 133001 (2021).
- [72] T. Caneva, R. Fazio, and G. E. Santoro, Adiabatic quantum dynamics of the Lipkin-Meshkov-Glick model, *Phys. Rev. B* **78**, 104426 (2008).
- [73] J. G. Muga, X. Chen, S. Ibáñez, I. Lizuain, and A. Ruschhaupt, Transitionless quantum drivings for the harmonic oscillator, *J. Phys. B: At. Mol. Opt. Phys.* **43**, 085509 (2010).
- [74] K. Takahashi, Transitionless quantum driving for spin systems, *Phys. Rev. E* **87**, 062117 (2013).
- [75] M. H. Goerz, K. B. Whaley, and C. P. Koch, Hybrid optimization schemes for quantum control, *EPJ Quantum Technol.* **2**, 21 (2015).
- [76] V. D. Vaidya, Y. Guo, R. M. Kroeze, K. E. Ballantine, A. J. Kollár, J. Keeling, and B. L. Lev, Tunable-range, photon-mediated atomic interactions in multimode cavity QED, *Phys. Rev. X* **8**, 011002 (2018).
- [77] H. Saberi, T. Opatrný, K. Mølmer, and A. del Campo, Adiabatic tracking of quantum many-body dynamics, *Phys. Rev. A* **90**, 060301(R) (2014).
- [78] J.-F. Schaff, X.-L. Song, P. Capuzzi, P. Vignolo, and G. Labeyrie, Shortcut to adiabaticity for an interacting Bose-Einstein condensate, *Europhys. Lett.* **93**, 23001 (2011).
- [79] M. G. Bason, M. Viteau, N. Malossi, P. Huillery, E. Arimondo, D. Ciampini, R. Fazio, V. Giovannetti, R. Manfellina, and O. Morsch, High-fidelity quantum driving, *Nat. Phys.* **8**, 147 (2012).
- [80] Y.-X. Du, Z.-T. Liang, Y.-C. Li, X.-X. Yue, Q.-X. Lv, W. Huang, X. Chen, H. Yan, and S.-L. Zhu, Experimental realization of stimulated Raman shortcut-to-adiabatic passage with cold atoms, *Nat. Commun.* **7**, 12479 (2016).

- [81] S. An, D. Lv, A. del Campo, and K. Kim, Shortcuts to adiabaticity by counterdiabatic driving for trapped-ion displacement in phase space, *Nat. Commun.* **7**, 12999 (2016).
- [82] B. B. Zhou, A. Baksic, H. Ribeiro, C. G. Yale, F. J. Heremans, P. C. Jerger, A. Auer, G. Burkard, A. A. Clerk, and D. D. Awschalom, Accelerated quantum control using superadiabatic dynamics in a solid-state lambda system, *Nat. Phys.* **13**, 330 (2016).
- [83] J. Zhang, J. H. Shim, I. Niemeyer, T. Taniguchi, T. Teraji, H. Abe, S. Onoda, T. Yamamoto, T. Ohshima, J. Isoya, and D. Suter, Experimental implementation of assisted quantum adiabatic passage in a single spin, *Phys. Rev. Lett.* **110**, 240501 (2013).
- [84] S. Dogra, A. Vepsäläinen, and G. S. Paraoanu, Experimental demonstration of robustness under scaling errors for superadiabatic population transfer in a superconducting circuit, *Phil. Trans. R. Soc. A* **380**, 20210274 (2022).
- [85] T. Wang, Z. Zhang, L. Xiang, Z. Jia, P. Duan, W. Cai, Z. Gong, Z. Zong, M. Wu, J. Wu, L. Sun, Y. Yin, and G. Guo, The experimental realization of high-fidelity “shortcut-to-adiabaticity” quantum gates in a superconducting Xmon qubit, *New J. Phys.* **20**, 065003 (2018).
- [86] Z. Zhang, T. Wang, L. Xiang, Z. Jia, P. Duan, W. Cai, Z. Zhan, Z. Zong, J. Wu, L. Sun, Y. Yin, and G. Guo, Experimental demonstration of work fluctuations along a shortcut to adiabaticity with a superconducting Xmon qubit, *New J. Phys.* **20**, 085001 (2018).
- [87] J.-W. Zhang, J.-T. Bu, J. C. Li, W. Meng, W.-Q. Ding, B. Wang, W.-F. Yuan, H.-J. Du, G.-Y. Ding, W.-J. Chen, L. Chen, F. Zhou, Z. Xu, and M. Feng, Single-atom verification of the optimal trade-off between speed and cost in shortcuts to adiabaticity, *Phys. Rev. Lett.* **132**, 213602 (2024).
- [88] E. J. Meier, K. Ngan, D. Sels, and B. Gadway, Counterdiabatic control of transport in a synthetic tight-binding lattice, *Phys. Rev. Res.* **2**, 043201 (2020).
- [89] H. Zhou, Y. Ji, X. Nie, X. Yang, X. Chen, J. Bian, and X. Peng, Experimental realization of shortcuts to adiabaticity in a nonintegrable spin chain by local counterdiabatic driving, *Phys. Rev. Appl.* **13**, 044059 (2020).
- [90] J. Wurtz and P. J. Love, Counterdiabaticity and the quantum approximate optimization algorithm, *Quantum* **6**, 635 (2022).
- [91] P. Chandarana, N. N. Hegade, K. Paul, F. Albarrán-Arriagada, E. Solano, A. del Campo, and X. Chen, Digitized-counterdiabatic quantum approximate optimization algorithm, *Phys. Rev. Res.* **4**, 013141 (2022).
- [92] P. Chandarana, N. N. Hegade, I. Montalban, E. Solano, and X. Chen, Digitized counterdiabatic quantum algorithm for protein folding, *Phys. Rev. Appl.* **20**, 014024 (2023).
- [93] N. N. Hegade, P. Chandarana, K. Paul, X. Chen, F. Albarrán-Arriagada, and E. Solano, Portfolio optimization with digitized counterdiabatic quantum algorithms, *Phys. Rev. Res.* **4**, 043204 (2022).
- [94] D. d’Alessandro, *Introduction to Quantum Control and Dynamics* (Chapman and Hall/CRC, New York, 2021).
- [95] P. M. Poggi, Geometric quantum speed limits and short-time accessibility to unitary operations, *Phys. Rev. A* **99**, 042116 (2019).
- [96] A. I. Konnov and V. F. Krotov, On global methods for the successive improvement of control processes, *Avtomat. i Telemekh.* 77 (1999); *Autom. Remote Control* **60**, 1427 (1999).
- [97] D. M. Reich, M. Ndong, and C. P. Koch, Monotonically convergent optimization in quantum control using Krotov’s method, *J. Chem. Phys.* **136**, 104103 (2012).
- [98] P. Doria, T. Calarco, and S. Montangero, Optimal control technique for many-body quantum dynamics, *Phys. Rev. Lett.* **106**, 190501 (2011).
- [99] H. A. Rabitz, M. M. Hsieh, and C. M. Rosenthal, Quantum optimally controlled transition landscapes, *Science* **303**, 1998 (2004).
- [100] D. V. Zhdanov and T. Seideman, Role of control constraints in quantum optimal control, *Phys. Rev. A* **92**, 052109 (2015).
- [101] R. L. Kosut, G. Bhole, and H. Rabitz, Robust quantum control: Analysis & synthesis via averaging, [arXiv:2208.14193](https://arxiv.org/abs/2208.14193).
- [102] P. M. Poggi, G. De Chiara, S. Campbell, and A. Kiely, Universally robust quantum control, *Phys. Rev. Lett.* **132**, 193801 (2024).
- [103] M. M. Taddei, B. M. Escher, L. Davidovich, and R. L. de Matos Filho, Quantum speed limit for physical processes, *Phys. Rev. Lett.* **110**, 050402 (2013).
- [104] D. P. Pires, M. Cianciaruso, L. C. Céleri, G. Adesso, and D. O. Soares-Pinto, Generalized geometric quantum speed limits, *Phys. Rev. X* **6**, 021031 (2016).
- [105] D. Mondal and A. K. Pati, Quantum speed limit for mixed states using an experimentally realizable metric, *Phys. Rev. Lett.* **110**, 1395 (2013).
- [106] I. Marvian, R. W. Spekkens, and P. Zanardi, Quantum speed limits, coherence, and asymmetry, *Phys. Rev. A* **93**, 052331 (2016).
- [107] L. Mandelstam and I. Tamm, The uncertainty relation between energy and time in non-relativistic quantum mechanics, *J. Phys. (USSR)* **9**, 249 (1945).
- [108] K. Bhattacharyya, Quantum decay and the Mandelstam-Tamm-energy inequality, *J. Phys. A: Math. Gen.* **16**, 2993 (1983).
- [109] T. Caneva, M. Murphy, T. Calarco, R. Fazio, S. Montangero, V. Giovannetti, and G. E. Santoro, Optimal control at the quantum speed limit, *Phys. Rev. Lett.* **103**, 240501 (2009).
- [110] G. C. Hegerfeldt, Driving at the quantum speed limit: Optimal control of a two-level system, *Phys. Rev. Lett.* **111**, 260501 (2013).
- [111] 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. Res.* **172**, 296 (2005).
- [112] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework, *Phys. Rev. A* **84**, 022305 (2011).
- [113] 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).
- [114] M. Larocca, P. M. Poggi, and D. A. Wisniacki, Quantum control landscape for a two-level system near the quantum speed limit, *J. Phys. A* **51**, 385305 (2018).

- [115] H. Rabitz, T.-S. Ho, M. Hsieh, R. Kosut, and M. Demiralp, Topology of optimally controlled quantum mechanical transition probability landscapes, *Phys. Rev. A* **74**, 012721 (2006).
- [116] P. M. Poggi, F. C. Lombardo, and D. A. Wisniacki, Time-optimal control fields for quantum systems with multiple avoided crossings, *Phys. Rev. A* **92**, 053411 (2015).
- [117] M. Bukov, A. G. R. Day, P. Weinberg, A. Polkovnikov, P. Mehta, and D. Sels, Broken symmetry in a two-qubit quantum control landscape, *Phys. Rev. A* **97**, 052114 (2018).
- [118] M. Murphy, S. Montangero, V. Giovannetti, and T. Calarco, Communication at the quantum speed limit along a spin chain, *Phys. Rev. A* **82**, 022318 (2010).
- [119] P. M. Poggi, F. C. Lombardo, and D. A. Wisniacki, Enhancement of quantum speed limit time due to cooperative effects in multilevel systems, *J. Phys. A* **48**, 35FT02 (2015).
- [120] M. H. Goerz, F. Motzoi, K. B. Whaley, and C. P. Koch, Charting the circuit QED design landscape using optimal control theory, *npj Quantum Inf.* **3**, 37 (2017).
- [121] N. Khaneja, R. Brockett, and S. J. Glaser, Time optimal control in spin systems, *Phys. Rev. A* **63**, 032308 (2001).
- [122] U. Boscain and P. Mason, Time minimal trajectories for a spin 1/2 particle in a magnetic field, *J. Math. Phys.* **47**, 062101 (2006).
- [123] V. Evangelakos, E. Paspalakis, and D. Stefanatos, Minimum-time generation of a uniform superposition in a qubit with only transverse field control, *Phys. Rev. A* **108**, 062425 (2023).
- [124] P. M. Poggi, F. C. Lombardo, and D. A. Wisniacki, Quantum speed limit and optimal evolution time in a two-level system, *Europhys. Lett.* **104**, 40005 (2013).
- [125] A. D. Boozer, Time-optimal synthesis of SU(2) transformations for a spin-1/2 system, *Phys. Rev. A* **85**, 012317 (2012).
- [126] G. Tóth, Multipartite entanglement and high-precision metrology, *Phys. Rev. A* **85**, 022322 (2012).
- [127] L. Pezze, A. Smerzi, M. K. Oberthaler, R. Schmied, and P. Treutlein, Quantum metrology with nonclassical states of atomic ensembles, *Rev. Mod. Phys.* **90**, 035005 (2018).
- [128] R. H. Dicke, Coherence in spontaneous radiation processes, *Phys. Rev.* **93**, 99 (1954).
- [129] I. H. Deutsch and P. S. Jessen, Quantum control and measurement of atomic spins in polarization spectroscopy, *Opt. Commun.* **283**, 681 (2010).
- [130] S. Omanakuttan, A. Mitra, M. J. Martin, and I. H. Deutsch, Quantum optimal control of ten-level nuclear spin qubits in Sr-87, *Phys. Rev. A* **104**, L060401 (2021).
- [131] 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).
- [132] J. Ruths and J.-S. Li, Optimal control of inhomogeneous ensembles, *IEEE Trans. Autom. Control* **57**, 2021 (2012).
- [133] A. Smith, B. E. Anderson, H. Sosa-Martinez, C. A. Riofrío, I. H. Deutsch, and P. S. Jessen, Quantum control in the Cs  $6S_{1/2}$  ground manifold using radio-frequency and microwave magnetic fields, *Phys. Rev. Lett.* **111**, 170502 (2013).
- [134] B. E. Anderson, H. Sosa-Martinez, C. A. Riofrío, I. H. Deutsch, and P. S. Jessen, Accurate and robust unitary transformations of a high-dimensional quantum system, *Phys. Rev. Lett.* **114**, 240401 (2015).
- [135] N. K. Lysne, K. W. Kuper, P. M. Poggi, I. H. Deutsch, and P. S. Jessen, Small, highly accurate quantum processor for intermediate-depth quantum simulations, *Phys. Rev. Lett.* **124**, 230501 (2020).
- [136] A. Larrouy, S. Patsch, R. Richaud, J.-M. Raimond, M. Brune, C. P. Koch, and S. Gleyzes, Fast navigation in a large Hilbert space using quantum optimal control, *Phys. Rev. X* **10**, 021058 (2020).
- [137] 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. Vuletic, and M. D. Lukin, Generation and manipulation of Schrödinger cat states in Rydberg atom arrays, *Science* **365**, 570 (2019).
- [138] S. J. Evered, D. Bluvstein, M. Kalinowski, S. Ebadi, T. Manovitz, H. Zhou, S. H. Li, A. A. Geim, T. T. Wang, N. Maskara *et al.*, High-fidelity parallel entangling gates on a neutral-atom quantum computer, *Nature* **622**, 268 (2023).
- [139] A. Cao, W. J. Eckner, T. Lukin Yelin, A. W. Young, S. Jandura, L. Yan, K. Kim, G. Pupillo, J. Ye, N. D. O'Connell *et al.*, Multi-qubit gates and Schrödinger cat states in an optical clock, *Nature* **634**, 315 (2024).
- [140] 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).
- [141] J. D. Sterk, H. Coakley, J. Goldberg, V. Hietala, J. Lechtenberg, H. McGuinness, D. McMurtrey, L. P. Parazzoli, J. Van Der Wall, and D. Stick, Closed-loop optimization of fast trapped-ion shuttling with sub-quanta excitation, *npj Quantum Inf.* **8**, 68 (2022).
- [142] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, Implementing a universal gate set on a logical qubit encoded in an oscillator, *Nat. Commun.* **8**, 94 (2017).
- [143] L. M. Seifert, Z. Li, T. Roy, D. I. Schuster, F. T. Chong, and J. M. Baker, Exploring ququart computation on a transmon using optimal control, *Phys. Rev. A* **108**, 062609 (2023).
- [144] 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).
- [145] B. Bonnard, S. J. Glaser, and D. Sugny, A review of geometric optimal control for quantum systems in nuclear magnetic resonance, *Adv. Math. Phys.* **2012**, 857493 (2012).
- [146] D. J. Tannor and S. A. Rice, Control of selectivity of chemical reaction via control of wave packet evolution, *J. Chem. Phys.* **83**, 5013 (1985).
- [147] P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, A high-bias, low-variance introduction to machine learning for physicists, *Phys. Rep.* **810**, 1 (2019).
- [148] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).

- [149] J. Carrasquilla, Machine learning for quantum matter, *Adv. Phys. X* **5**, 1797528 (2020).
- [150] M. Krenn, J. Landgraf, T. Foessel, and F. Marquardt, Artificial intelligence and machine learning for quantum technologies, *Phys. Rev. A* **107**, 010101 (2023).
- [151] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, Cambridge, MA, 2018).
- [152] E. Todorov, Optimal control theory, in *Bayesian Brain: Probabilistic Approaches to Neural Coding* (2006).
- [153] J. M. Rodríguez-Borbón, X. Wang, A. P. Diéguez, K. Z. Ibrahim, and B. M. Wong, Van-damme: GPU-accelerated and symmetry-assisted quantum optimal control of multi-qubit systems, *Comput. Phys. Commun.* **307**, 109403 (2025).
- [154] D. Koutromanos, D. Stefanatos, and E. Paspalakis, TorchQC - a framework for efficiently integrating machine and deep learning methods in quantum dynamics and control, *Comput. Phys. Commun.* **310**, 109505 (2025).
- [155] F. Metz and M. Bukov, Self-correcting quantum many-body control using reinforcement learning with tensor networks, *Nat. Mach. Intell.* **5**, 780 (2023).
- [156] E. Gillman, D. C. Rose, and J. P. Garrahan, Combining reinforcement learning and tensor networks, with an application to dynamical large deviations, *Phys. Rev. Lett.* **132**, 197301 (2024).
- [157] R. J. Williams, Simple statistical gradient-following algorithms for connectionist reinforcement learning, *Mach. Learn.* **8**, 229 (1992).
- [158] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, Proximal policy optimization algorithms, [arXiv:1707.06347](https://arxiv.org/abs/1707.06347).
- [159] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, Trust region policy optimization, *Proc. Mach. Learn. Res.* **37**, 1889 (2015).
- [160] V. Konda and J. Tsitsiklis, Actor-critic algorithms, *Adv. Neural Inf. Process. Syst.* **12**, 1008 (1999).
- [161] C. J. Watkins and P. Dayan, Q-learning, *Mach. Learn.* **8**, 279 (1992).
- [162] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, Human-level control through deep reinforcement learning, *Nature* **518**, 529 (2015).
- [163] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, Continuous control with deep reinforcement learning, [arXiv:1509.02971](https://arxiv.org/abs/1509.02971).
- [164] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, Reinforcement learning in different phases of quantum control, *Phys. Rev. X* **8**, 031086 (2018).
- [165] 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).
- [166] V. V. Sivak, A. Eickbusch, H. Liu, B. Royer, I. Tsoutsios, and M. H. Devoret, Model-free quantum control with reinforcement learning, *Phys. Rev. X* **12**, 011059 (2022).
- [167] R. Porotti, V. Peano, and F. Marquardt, Gradient-ascent pulse engineering with feedback, *PRX Quantum* **4**, 030305 (2023).
- [168] K. Reuer, J. Landgraf, T. Fösel, J. O'Sullivan, L. Beltrán, A. Akin, G. J. Norris, A. Remm, M. Kerschbaum, J.-C. Besse *et al.*, Realizing a deep reinforcement learning agent for real-time quantum feedback, *Nat. Commun.* **14**, 7138 (2023).
- [169] Y. Ding, Y. Ban, J. D. Martín-Guerrero, E. Solano, J. Casanova, and X. Chen, Breaking adiabatic quantum control with deep learning, *Phys. Rev. A* **103**, L040401 (2021).
- [170] Y. Ding, X. Chen, R. Magdalena-Benedito, and J. D. Martín-Guerrero, Closed-loop control of a noisy qubit with reinforcement learning, *Mach. Learn.: Sci. Technol.* **4**, 025020 (2023).
- [171] M.-Z. Ai, Y. Ding, Y. Ban, J. D. Martín-Guerrero, J. Casanova, J.-M. Cui, Y.-F. Huang, X. Chen, C.-F. Li, and G.-C. Guo, Experimentally realizing efficient quantum control with reinforcement learning, *Sci. China Phys. Mech. Astron.* **65**, 250312 (2022).
- [172] A. Fallani, M. A. C. Rossi, D. Tamascelli, and M. G. Genoni, Learning feedback control strategies for quantum metrology, *PRX Quantum* **3**, 020310 (2022).
- [173] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven, Universal quantum control through deep reinforcement learning, *npj Quantum Inf.* **5**, 33 (2019).
- [174] M. Dalgaard, F. Motzoi, J. J. Sørensen, and J. Sherson, Global optimization of quantum dynamics with alphazero deep exploration, *npj Quantum Inf.* **6**, 6 (2020).
- [175] Y. Baum, M. Amico, S. Howell, M. Hush, M. Liuzzi, P. Mundada, T. Merkh, A. R. R. Carvalho, and M. J. Biercuk, Experimental deep reinforcement learning for error-robust gate-set design on a superconducting quantum computer, *PRX Quantum* **2**, 040324 (2021).
- [176] L. Moro, M. G. Paris, M. Restelli, and E. Prati, Quantum compiling by deep reinforcement learning, *Commun. Phys.* **4**, 178 (2021).
- [177] H.-N. Nguyen, F. Motzoi, M. Metcalf, K. B. Whaley, M. Bukov, and M. Schmitt, Reinforcement learning pulses for transmon qubit entangling gates, *Mach. Learn.: Sci. Technol.* **5**, 025066 (2024).
- [178] T. Fösel, M. Y. Niu, F. Marquardt, and L. Li, Quantum circuit optimization with deep reinforcement learning, [arXiv:2103.07585](https://arxiv.org/abs/2103.07585).
- [179] A. Bolens and M. Heyl, Reinforcement learning for digital quantum simulation, *Phys. Rev. Lett.* **127**, 110502 (2021).
- [180] Z. He, L. Li, S. Zheng, Y. Li, and H. Situ, Variational quantum compiling with double q-learning, *New J. Phys.* **23**, 033002 (2021).
- [181] D. A. Herrera-Martí, Policy gradient approach to compilation of variational quantum circuits, *Quantum* **6**, 797 (2022).
- [182] Y. J. Patel, A. Kundu, M. Ostaszewski, X. Bonet-Monroig, V. Dunjko, and O. Danaci, Curriculum reinforcement learning for quantum architecture search under hardware errors, [arXiv:2402.03500](https://arxiv.org/abs/2402.03500).
- [183] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, Reinforcement learning with neural networks for quantum feedback, *Phys. Rev. X* **8**, 031084 (2018).
- [184] R. Sweke, M. S. Kesselring, E. P. van Nieuwenburg, and J. Eisert, Reinforcement learning decoders for fault-tolerant quantum computation, *Mach. Learn.: Sci. Technol.* **2**, 025005 (2020).

- [185] P. Andreasson, J. Johansson, S. Liljestrand, and M. Granath, Quantum error correction for the toric code using deep reinforcement learning, *Quantum* **3**, 183 (2019).
- [186] D. Fitzek, M. Eliasson, A. F. Kockum, and M. Granath, Deep q-learning decoder for depolarizing noise on the toric code, *Phys. Rev. Res.* **2**, 023230 (2020).
- [187] V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. Brock, A. Ding, L. Frunzio *et al.*, Real-time quantum error correction beyond break-even, *Nature* **616**, 50 (2023).
- [188] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/j8c7-v2hd> for movies show the time-dependent dynamics corresponding to the cases of Fig. 12(b).
- [189] A. Ferrer-Sánchez, C. Flores-Garrigos, C. Hernani-Morales, J. J. Orquín-Marqués, N. N. Hegade, A. G. Cadavid, I. Montalban, E. Solano, Y. Vives-Gilbert, and J. D. Martín-Guerrero, Physics-informed neural networks for an optimal counterdiabatic quantum computation, [arXiv:2309.04434](https://arxiv.org/abs/2309.04434).
- [190] A. Kiely and S. Campbell, Fast and robust magnon transport in a spin chain, *New J. Phys.* **23**, 033033 (2021).
- [191] L. Coopmans, S. Campbell, G. De Chiara, and A. Kiely, Optimal control in disordered quantum systems, *Phys. Rev. Res.* **4**, 043138 (2022).
- [192] R. Chakrabarti and H. Rabitz, Quantum control landscapes, *Int. Rev. Phys. Chem.* **26**, 671 (2007).
- [193] N. Beato, P. Patil, and M. Bukov, Towards a theory of phase transitions in quantum control landscapes, [arXiv:2408.11110v1](https://arxiv.org/abs/2408.11110v1).
- [194] N. Beato, P. Patil, and M. Bukov, Topological phase transitions in a constrained two-qubit quantum control landscape, [arXiv:2411.08736v1](https://arxiv.org/abs/2411.08736v1).
- [195] H. W. Fentaw, S. Campbell, and S. Caton, Exploring quantum control landscape and solution space complexity through optimization algorithms and dimensionality reduction, *Sci. Rep.* **15**, 14605 (2025).
- [196] D. Stefanatos, J. Ruths, and J.-S. Li, Frictionless atom cooling in harmonic traps: A time-optimal approach, *Phys. Rev. A* **82**, 063422 (2010).
- [197] C. Whitty, A. Kiely, and A. Ruschhaupt, Quantum control via enhanced shortcuts to adiabaticity, *Phys. Rev. Res.* **2**, 023360 (2020).
- [198] C. Whitty, A. Kiely, and A. Ruschhaupt, Robustness of enhanced shortcuts to adiabaticity in lattice transport, *Phys. Rev. A* **105**, 013311 (2022).
- [199] S. H. Hauck, G. Alber, and V. M. Stojanović, Single-atom transport in optical conveyor belts: Enhanced shortcuts-to-adiabaticity approach, *Phys. Rev. A* **104**, 053110 (2021).
- [200] C. Whitty, A. Kiely, and A. Ruschhaupt, Improved anharmonic trap expansion through enhanced shortcuts to adiabaticity, *J. Phys. B: At. Mol. Opt. Phys.* **55**, 194003 (2022).
- [201] S. H. Hauck and V. M. Stojanović, Coherent atom transport via enhanced shortcuts to adiabaticity: Double-well optical lattice, *Phys. Rev. Appl.* **18**, 014016 (2022).
- [202] M. Odelli, A. Ruschhaupt, and V. M. Stojanović, Twist-and-turn dynamics of spin squeezing in bosonic Josephson junctions: Enhanced shortcuts-to-adiabaticity approach, *Phys. Rev. A* **110**, 022610 (2024).
- [203] J. Yao, H. Li, M. Bukov, L. Lin, and L. Ying, in *Mathematical and Scientific Machine Learning* (Proceedings of Machine Learning Research, 2022), Vol. 190, pp. 49–64.
- [204] G. N. Fleming, A unitarity bound on the evolution of non-stationary states, *Il Nuovo Cimento A* (1971-1996) **16**, 232 (1973).
- [205] J. Anandan and Y. Aharonov, Geometry of quantum evolution, *Phys. Rev. Lett.* **65**, 1697 (1990).
- [206] N. Margolus and L. B. Levitin, The maximum speed of dynamical evolution, *Physica D* **120**, 188 (1998).
- [207] A. Grabarits, F. Balducci, B. C. Sanders, and A. del Campo, Nonadiabatic quantum optimization for crossing quantum phase transitions, *Phys. Rev. A* **111**, 012215 (2025).
- [208] S. Pang and A. N. Jordan, Optimal adaptive control for quantum metrology with time-dependent hamiltonians, *Nat. Commun.* **8**, 14695 (2017).
- [209] J. Liu and H. Yuan, Quantum parameter estimation with optimal control, *Phys. Rev. A* **96**, 012117 (2017).
- [210] C. Lin, Y. Ma, and D. Sels, Optimal control for quantum metrology via Pontryagin’s principle, *Phys. Rev. A* **103**, 052607 (2021).
- [211] K. Gietka, F. Metz, T. Keller, and J. Li, Adiabatic critical quantum metrology cannot reach the Heisenberg limit even when shortcuts to adiabaticity are applied, *Quantum* **5**, 489 (2021).
- [212] R. Kaubruegger, D. V. Vasilyev, M. Schulte, K. Hammerer, and P. Zoller, Quantum variational optimization of Ramsey interferometry and atomic clocks, *Phys. Rev. X* **11**, 041045 (2021).
- [213] F. Mazzoncini, V. Cavina, G. M. Andolina, P. A. Erdman, and V. Giovannetti, Optimal control methods for quantum batteries, *Phys. Rev. A* **107**, 032218 (2023).
- [214] R. R. Rodríguez, B. Ahmadi, G. Suárez, P. Mazurek, S. Barzanjeh, and P. Horodecki, Optimal quantum control of charging quantum batteries, *New J. Phys.* **26**, 043004 (2024).
- [215] V. Evangelakos, E. Paspalakis, and D. Stefanatos, Fast charging of an Ising-spin-pair quantum battery using optimal control, *Phys. Rev. A* **110**, 052601 (2024).
- [216] V. Evangelakos, E. Paspalakis, and D. Stefanatos, Rapid charging of a two-qubit quantum battery by transverse field amplitude and phase control, *Quantum Sci. Technol.* **10**, 035024 (2025).
- [217] P. A. Erdman, G. M. Andolina, V. Giovannetti, and F. Noé, Reinforcement learning optimization of the charging of a Dicke quantum battery, *Phys. Rev. Lett.* **133**, 243602 (2024).
- [218] P.-Y. Sun, H. Zhou, and F.-Q. Dou, Cavity-Heisenberg spin- $j$  chain quantum battery and reinforcement learning optimization, [arXiv:2412.01442](https://arxiv.org/abs/2412.01442).
- [219] C.-F. A. Chen, A. Lucas, and C. Yin, Speed limits and locality in many-body quantum dynamics, *Rep. Prog. Phys.* **86**, 116001 (2023).
- [220] N. Lashkari, D. Stanford, M. Hastings, T. Osborne, and P. Hayden, Towards the fast scrambling conjecture, *J. High Energy Phys.* **2013**, 22 (2013).
- [221] L. Campos Venuti and P. Zanardi, Quantum critical scaling of the geometric tensors, *Phys. Rev. Lett.* **99**, 095701 (2007).

- [222] S.-J. Gu, Fidelity approach to quantum phase transitions, *Int. J. Mod. Phys. B* **24**, 4371 (2010).
- [223] M. Pandey, P. W. Claeys, D. K. Campbell, A. Polkovnikov, and D. Sels, Adiabatic eigenstate deformations as a sensitive probe for quantum chaos, *Phys. Rev. X* **10**, 041017 (2020).
- [224] C. Lim, K. Matirko, A. Polkovnikov, and M. O. Flynn, Defining classical and quantum chaos through adiabatic transformations, [arXiv:2401.01927](https://arxiv.org/abs/2401.01927).
- [225] H. Kim and A. Polkovnikov, Integrability as an attractor of adiabatic flows, *Phys. Rev. B* **109**, 195162 (2024).
- [226] N. Goldman and J. Dalibard, Periodically driven quantum systems: Effective Hamiltonians and engineered gauge fields, *Phys. Rev. X* **4**, 031027 (2014).
- [227] M. Bukov, L. D'Alessio, and A. Polkovnikov, Universal high-frequency behavior of periodically driven systems: from dynamical stabilization to Floquet engineering, *Adv. Phys.* **64**, 139 (2015).
- [228] A. Eckardt, Colloquium: Atomic quantum gases in periodically driven optical lattices, *Rev. Mod. Phys.* **89**, 011004 (2017).
- [229] M. Aidelsburger, S. Nascimbene, and N. Goldman, Artificial gauge fields in materials and engineered systems, *C. R. Phys.* **19**, 394 (2018).
- [230] T. Oka and S. Kitamura, Floquet engineering of quantum materials, *Annu. Rev. Condens. Matter Phys.* **10**, 387 (2019).
- [231] C. Weitenberg and J. Simonet, Tailoring quantum gases by Floquet engineering, *Nat. Phys.* **17**, 1342 (2021).
- [232] E. Boyers, M. Pandey, D. K. Campbell, A. Polkovnikov, D. Sels, and A. O. Sushkov, Floquet-engineered quantum state manipulation in a noisy qubit, *Phys. Rev. A* **100**, 012341 (2019).
- [233] P. M. Schindler and M. Bukov, Geometric Floquet theory, *Phys. Rev. X* **15**, 031037 (2025).
- [234] K. Takahashi and A. del Campo, Shortcuts to adiabaticity in Krylov space, *Phys. Rev. X* **14**, 011032 (2024).
- [235] H. Kim, M. Fishman, and D. Sels, Variational adiabatic transport of tensor networks, *PRX Quantum* **5**, 020361 (2024).
- [236] C. Mc Keever and M. Lubasch, Towards adiabatic quantum computing using compressed quantum circuits, *PRX Quantum* **5**, 020362 (2024).
- [237] D. A. Lidar, Review of decoherence free subspaces, noiseless subsystems, and dynamical decoupling, *Adv. Chem. Phys.* **154**, 295 (2014).
- [238] E. Kapit, The upside of noise: Engineered dissipation as a resource in superconducting circuits, *Quantum Sci. Technol.* **2**, 033002 (2017).
- [239] L. C. Venuti, T. Albash, D. A. Lidar, and P. Zanardi, Adiabaticity in open quantum systems, *Phys. Rev. A* **93**, 032118 (2016).
- [240] A. C. Santos and M. S. Sarandy, Sufficient conditions for adiabaticity in open quantum systems, *Phys. Rev. A* **102**, 052215 (2020).
- [241] G. Vacanti, R. Fazio, S. Montangero, G. M. Palma, M. Paternostro, and V. Vedral, Transitionless quantum driving in open quantum systems, *New J. Phys.* **16**, 053017 (2014).
- [242] J. Jing, M. S. Sarandy, D. A. Lidar, D.-W. Luo, and L.-A. Wu, Eigenstate tracking in open quantum systems, *Phys. Rev. A* **94**, 042131 (2016).
- [243] R. Verresen, N. Tantivasadakarn, and A. Vishwanath, Efficiently preparing Schrödinger's cat, fractons and non-Abelian topological order in quantum devices, [arXiv:2112.03061](https://arxiv.org/abs/2112.03061).
- [244] L. Piroli, G. Styliaris, and J. I. Cirac, Quantum circuits assisted by local operations and classical communication: Transformations and phases of matter, *Phys. Rev. Lett.* **127**, 220503 (2021).
- [245] K. C. Smith, E. Crane, N. Wiebe, and S. M. Girvin, Deterministic constant-depth preparation of the AKLT state on a quantum processor using fusion measurements, *PRX Quantum* **4**, 020315 (2023).
- [246] G.-Y. Zhu, N. Tantivasadakarn, A. Vishwanath, S. Trebst, and R. Verresen, Nishimori's cat: Stable long-range entanglement from finite-depth unitaries and weak measurements, *Phys. Rev. Lett.* **131**, 200201 (2023).
- [247] K. C. Smith, A. Khan, B. K. Clark, S. M. Girvin, and T.-C. Wei, Constant-depth preparation of matrix product states with adaptive quantum circuits, *PRX Quantum* **5**, 030344 (2024).
- [248] M. Iqbal, N. Tantivasadakarn, T. M. Gatterman, J. A. Gerber, K. Gilmore, D. Gresh, A. Hankin, N. Hewitt, C. V. Horst, M. Matheny *et al.*, Topological order from measurements and feed-forward on a trapped ion quantum computer, *Commun. Phys.* **7**, 205 (2024).
- [249] T. Jörg, F. Krzakala, J. Kurchan, A. C. Maggs, and J. Pujos, Energy gaps in quantum first-order mean-field-like transitions: The problems that quantum annealing cannot solve, *Europhys. Lett.* **89**, 40004 (2010).
- [250] R. Ghosh, L. A. Nutricati, N. Feinstein, P. Warburton, and S. Bose, Exponential speed-up of quantum annealing via n-local catalysts, [arXiv:2409.13029](https://arxiv.org/abs/2409.13029).
- [251] F. Balducci, A. Grabarits, and A. del Campo, Fighting exponentially small gaps by counterdiabatic driving, [arXiv:2410.02520](https://arxiv.org/abs/2410.02520).
- [252] L. Prielinger, A. Hartmann, Y. Yamashiro, K. Nishimura, W. Lechner, and H. Nishimori, Two-parameter counterdiabatic driving in quantum annealing, *Phys. Rev. Res.* **3**, 013227 (2021).
- [253] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, Variational quantum algorithms, *Nat. Rev. Phys.* **3**, 625 (2021).
- [254] A. B. Magann, C. Arenz, M. D. Grace, T.-S. Ho, R. L. Kosut, J. R. McClean, H. A. Rabitz, and M. Sarovar, From pulses to circuits and back again: A quantum optimal control perspective on variational quantum algorithms, *PRX Quantum* **2**, 010101 (2021).
- [255] M. Larocca, P. Czarnik, K. Sharma, G. Muraleedharan, P. J. Coles, and M. Cerezo, Diagnosing barren plateaus with tools from quantum optimal control, *Quantum* **6**, 824 (2022).
- [256] M. Ragone, B. N. Bakalov, F. Sauvage, A. F. Kemper, C. Ortiz Marrero, M. Larocca, and M. Cerezo, A Lie algebraic theory of barren plateaus for deep parameterized quantum circuits, *Nat. Commun.* **15**, 7172 (2024).

- [257] 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).
- [258] P. Groszkowski, A. G. Fowler, F. Motzoi, and F. K. Wilhelm, Tunable coupling between three qubits as a building block for a superconducting quantum computer, *Phys. Rev. B* **84**, 144516 (2011).
- [259] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, Procedure for systematically tuning up cross-talk in the cross-resonance gate, *Phys. Rev. A* **93**, 060302(R) (2016).
- [260] 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).
- [261] K. X. Wei, E. Magesan, I. Lauer, S. Srinivasan, D. F. Bogorin, S. Carnevale, G. A. Keefe, Y. Kim, D. Klaus, W. Landers *et al.*, Hamiltonian engineering with multicolor drives for fast entangling gates and quantum crosstalk cancellation, *Phys. Rev. Lett.* **129**, 060501 (2022).
- [262] K. Heya and N. Kanazawa, Cross-cross resonance gate, *PRX Quantum* **2**, 040336 (2021).
- [263] M. Malekakhlagh and E. Magesan, Mitigating off-resonant error in the cross-resonance gate, *Phys. Rev. A* **105**, 012602 (2022).
- [264] B. Li, T. Calarco, and F. Motzoi, Experimental error suppression in cross-resonance gates via multi-derivative pulse shaping, *npj Quantum Inf.* **10**, 66 (2024).
- [265] E. Hyypä, A. Vepsäläinen, M. Papič, C. F. Chan, S. Inel, A. Landra, W. Liu, J. Luus, F. Marxer, C. Ockeloen-Korppi *et al.*, Reducing leakage of single-qubit gates for superconducting quantum processors using analytical control pulse envelopes, *PRX Quantum* **5**, 030353 (2024).
- [266] R. Gautier, É. Genois, and A. Blais, Optimal control in large open quantum systems: The case of transmon readout and reset, *Phys. Rev. Lett.* **134**, 070802 (2025).
- [267] D. A. Rower, L. Ding, H. Zhang, M. Hays, J. An, P. M. Harrington, I. T. Rosen, J. M. Gertler, T. M. Hazard, B. M. Niedzielski *et al.*, Suppressing counter-rotating errors for fast single-qubit gates with fluxonium, *PRX Quantum* **5**, 040342 (2024).
- [268] M. F. Zwanenburg, S. Singh, E. Y. Huang, F. Yilmaz, T. V. Stefanski, J. Hu, P. Kumaravadivel, and C. K. Andersen, Single-qubit gates beyond the rotating-wave approximation for strongly anharmonic low-frequency qubits, [arXiv:2503.08238](https://arxiv.org/abs/2503.08238).
- [269] A. Sørensen and K. Mølmer, Entanglement and quantum computation with ions in thermal motion, *Phys. Rev. A* **62**, 022311 (2000).
- [270] V. Schäfer, C. Ballance, K. Thirumalai, L. Stephenson, T. Ballance, A. Steane, and D. Lucas, Fast quantum logic gates with trapped-ion qubits, *Nature* **555**, 75 (2018).
- [271] L. Gerster, F. Martínez-García, P. Hrmo, M. W. van Mourik, B. Wilhelm, D. Vodola, M. Müller, R. Blatt, P. Schindler, and T. Monz, Experimental Bayesian calibration of trapped-ion entangling operations, *PRX Quantum* **3**, 020350 (2022).
- [272] S. A. Moses, C. H. Baldwin, M. S. Allman, R. Ancona, L. Ascarrunz, C. Barnes, J. Bartolotta, B. Bjork, P. Blanchard, M. Bohn *et al.*, A race-track trapped-ion quantum processor, *Phys. Rev. X* **13**, 041052 (2023).
- [273] S. Kirchhoff, F. K. Wilhelm, and F. Motzoi, Correction formulas for the Mølmer-Sørensen gate under strong driving, *PRX Quantum* **6**, 010328 (2025).
- [274] R. Barends, L. Lamata, J. Kelly, L. García-Álvarez, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus *et al.*, Digital quantum simulation of fermionic models with a superconducting circuit, *Nat. Commun.* **6**, 7654 (2015).
- [275] J. M. Martinis and M. R. Geller, Fast adiabatic qubit gates using only  $\sigma$  z control, *Phys. Rev. A* **90**, 022307 (2014).
- [276] M. Ganzhorn, D. J. Egger, P. Barkoutsos, P. Ollitrault, G. Salis, N. Moll, M. Roth, A. Fuhrer, P. Mueller, S. Woerner *et al.*, Gate-efficient simulation of molecular eigenstates on a quantum computer, *Phys. Rev. Appl.* **11**, 044092 (2019).
- [277] B. Foxen, C. Neill, A. Dunsworth, P. Roushan, B. Chiaro, A. Megrant, J. Kelly, Z. Chen, K. Satzinger, R. Barends *et al.*, Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms, *Phys. Rev. Lett.* **125**, 120504 (2020).
- [278] N. Lacroix, C. Hellings, C. K. Andersen, A. Di Paolo, A. Remm, S. Lazar, S. Krinner, G. J. Norris, M. Gabureac, J. Heinsoo *et al.*, Improving the performance of deep quantum optimization algorithms with continuous gate sets, *PRX Quantum* **1**, 020304 (2020).
- [279] S. E. Rasmussen and N. T. Zinner, Parameterized two-qubit gates for enhanced variational quantum eigensolver, *Ann. Phys.* **534**, 2200338 (2022).
- [280] Y. Shi, P. Gokhale, P. Murali, J. M. Baker, C. Duckering, Y. Ding, N. C. Brown, C. Chamberland, A. Javadi-Abhari, A. W. Cross *et al.*, Resource-efficient quantum computing by breaking abstractions, *Proc. IEEE* **108**, 1353 (2020).
- [281] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2010).
- [282] E. Zahedinejad, J. Ghosh, and B. C. Sanders, High-fidelity single-shot Toffoli gate via quantum control, *Phys. Rev. Lett.* **114**, 200502 (2015).
- [283] M. Khazali and K. Mølmer, Fast multiqubit gates by adiabatic evolution in interacting excited-state manifolds of Rydberg atoms and superconducting circuits, *Phys. Rev. X* **10**, 021054 (2020).
- [284] S. E. Rasmussen, K. Groenland, R. Gerritsma, K. Schoutens, and N. T. Zinner, Single-step implementation of high-fidelity n-bit Toffoli gates, *Phys. Rev. A* **101**, 022308 (2020).
- [285] C. W. Warren, J. Fernández-Pendás, S. Ahmed, T. Abad, A. Bengtsson, J. Biznárová, K. Debnath, X. Gu, C. Križan, A. Osman *et al.*, Extensive characterization and implementation of a family of three-qubit gates at the coherence limit, *npj Quantum Inf.* **9**, 44 (2023).
- [286] S. Jandura and G. Pupillo, Time-optimal two- and three-qubit gates for Rydberg atoms, *Quantum* **6**, 712 (2022).
- [287] L. Henriët, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak, Quantum computing with neutral atoms, *Quantum* **4**, 327 (2020).

- [288] M. Morgado and S. Whitlock, Quantum simulation and computing with Rydberg-interacting qubits, *AVS Quantum Sci.* **3**, 023501 (2021).
- [289] C. Figgatt, D. Maslov, K. A. Landsman, N. M. Linke, S. Debnath, and C. Monroe, Complete 3-qubit grover search on a programmable quantum computer, *Nat. Commun.* **8**, 1918 (2017).
- [290] O. Katz, L. Feng, A. Risinger, C. Monroe, and M. Cetina, Demonstration of three-and four-body interactions between trapped-ion spins, *Nat. Phys.* **19**, 1452 (2023).
- [291] Y. Wang, Z. Hu, B. C. Sanders, and S. Kais, Qudits and high-dimensional quantum computing, *Front. Phys.* **8**, 589504 (2020).
- [292] S. E. Rasmussen, N. J. S. Loft, T. Bækkegaard, M. Kues, and N. T. Zinner, Reducing the amount of single-qubit rotations in VQE and related algorithms, *Adv. Quantum Technol.* **3**, 2000063 (2020).
- [293] T. Bækkegaard, L. Kristensen, N. J. Loft, C. K. Andersen, D. Petrosyan, and N. T. Zinner, Realization of efficient quantum gates with a superconducting qubit-qutrit circuit, *Sci. Rep.* **9**, 13389 (2019).
- [294] E. J. Davis, G. Bentsen, L. Homeier, T. Li, and M. H. Schleier-Smith, Photon-mediated spin-exchange dynamics of spin-1 atoms, *Phys. Rev. Lett.* **122**, 010405 (2019).
- [295] M. S. Blok, V. V. Ramasesh, T. Schuster, K. O'Brien, J.-M. Kreikebaum, D. Dahlen, A. Morvan, B. Yoshida, N. Y. Yao, and I. Siddiqi, Quantum information scrambling on a superconducting qutrit processor, *Phys. Rev. X* **11**, 021010 (2021).
- [296] M. Ringbauer, M. Meth, L. Postler, R. Stricker, R. Blatt, P. Schindler, and T. Monz, A universal qudit quantum processor with trapped ions, *Nat. Phys.* **18**, 1053 (2022).
- [297] C. Reimer, S. Sciara, P. Roztock, M. Islam, L. Romero Cortés, Y. Zhang, B. Fischer, S. Loranger, R. Kashyap, A. Cino *et al.*, High-dimensional one-way quantum processing implemented on  $d$ -level cluster states, *Nat. Phys.* **15**, 148 (2019).
- [298] Y. Chi *et al.*, A programmable qudit-based quantum processor, *Nat. Commun.* **13**, 1166 (2022).
- [299] S. Omanakuttan, V. Buchemavari, J. A. Gross, I. H. Deutsch, and M. Marvian, Fault-tolerant quantum computation using large spin-cat codes, *PRX Quantum* **5**, 020355 (2024).
- [300] D. Lewis, R. Wiersema, J. Carrasquilla, and S. Bose, Geodesic algorithm for unitary gate design with time-independent Hamiltonians, *Phys. Rev. A* **111**, 052618 (2025).
- [301] E. Carolan, B. Çakmak, and S. Campbell, Robustness of controlled Hamiltonian approaches to unitary quantum gates, *Phys. Rev. A* **108**, 022423 (2023).
- [302] O. Raii, A. Dey, F. Mintert, and D. Burgarth, Creation and manipulation of surface code defects with quantum optimal control, *Phys. Rev. A* **111**, 012423 (2025).
- [303] P. Kairys and T. S. Humble, Parametrized Hamiltonian simulation using quantum optimal control, *Phys. Rev. A* **104**, 042602 (2021).
- [304] C. W. Duncan, P. M. Poggi, M. Bukov, N. T. Zinner, and S. Campbell, Taming quantum systems: A tutorial for using shortcuts-to-adiabaticity, quantum optimal control, and reinforcement learning, Zenodo (2025), <https://doi.org/10.5281/zenodo.17169846>.