

Protocol Discovery for the Quantum Control of Majoranas by Differentiable Programming and Natural Evolution Strategies

Luuk Coopmans^{1,2,*}, Di Luo^{3,†}, Graham Kells¹, Bryan K. Clark³ and Juan Carrasquilla^{4,5}

¹*Dublin Institute for Advanced Studies, School of Theoretical Physics, 10 Burlington Road, Dublin, Ireland*

²*School of Physics, Trinity College Dublin, College Green, Dublin 2, Ireland*

³*Department of Physics and IQUIST and Institute for Condensed Matter Theory, University of Illinois at Urbana-Champaign, Illinois 61801, USA*

⁴*Vector Institute for Artificial Intelligence, MaRS Centre, Toronto, Ontario, Canada*

⁵*Department of Physics and Astronomy, University of Waterloo, Ontario N2L 3G1, Canada*



(Received 28 August 2020; revised 8 April 2021; accepted 10 May 2021; published 2 June 2021)

Quantum control, which refers to the active manipulation of physical systems described by the laws of quantum mechanics, constitutes an essential ingredient for the development of quantum technology. Here we apply differentiable programming (DP) and natural evolution strategies (NES) to the optimal transport of Majorana zero modes in superconducting nanowires, a key element to the success of Majorana-based topological quantum computation. We formulate the motion control of Majorana zero modes as an optimization problem for which we propose a new categorization of four different regimes with respect to the critical velocity of the system and the total transport time. In addition to correctly recovering the anticipated smooth protocols in the adiabatic regime, our algorithms uncover efficient but strikingly counterintuitive motion strategies in the nonadiabatic regime. The emergent picture reveals a simple but high-fidelity strategy that makes use of pulselike jumps at the beginning and the end of the protocol with a period of constant velocity in between the jumps, which we dub the jump-move-jump protocol. We provide a transparent semianalytical picture, which uses the sudden approximation and a reformulation of the Majorana motion in a moving frame, to illuminate the key characteristics of the jump-move-jump control strategy. We verify that the jump-move-jump protocol remains robust against the presence of interactions or disorder, and corroborate its high efficacy on a realistic proximity-coupled nanowire model. Our results demonstrate that machine learning for quantum control can be applied efficiently to quantum many-body dynamical systems with performance levels that make it relevant to the realization of large-scale quantum technology.

DOI: [10.1103/PRXQuantum.2.020332](https://doi.org/10.1103/PRXQuantum.2.020332)

I. INTRODUCTION

Through the use of a wide array of promising experimental platforms ranging from superconducting qubits [1] to trapped ions [2–4], optical lattices [5], and nitrogen-vacancy centers [6,7], scientists are exploring groundbreaking ways to build quantum technology with an eye on deepening our understanding of complex natural systems, improving artificial intelligence, and impacting

industry more broadly. While promising, several fundamental and practical difficulties must be overcome for quantum machines to become practical [8]. Quantum control, which studies the manipulation of physical systems whose behavior is dominated by the laws of quantum mechanics, remains a fundamental ingredient in the quest for practical quantum technology. Thus, the development of principles and strategies for quantum control [9–13] remains an essential task for future quantum technologies [14], e.g., in the preparation of complex quantum states in quantum computing and quantum simulators.

The practical encoding and manipulation of quantum information has, however, been hampered by the presence of noise and decoherence inherent to brittle quantum devices. Majorana zero modes, which are special zero-energy quasiparticles with non-Abelian braiding statistics that can potentially be realized in proximity-coupled superconductors [15–24], represent a promising alternative

*coopmanl@tcd.ie

†diluo2@illinois.edu

‡Co-first authors.

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.

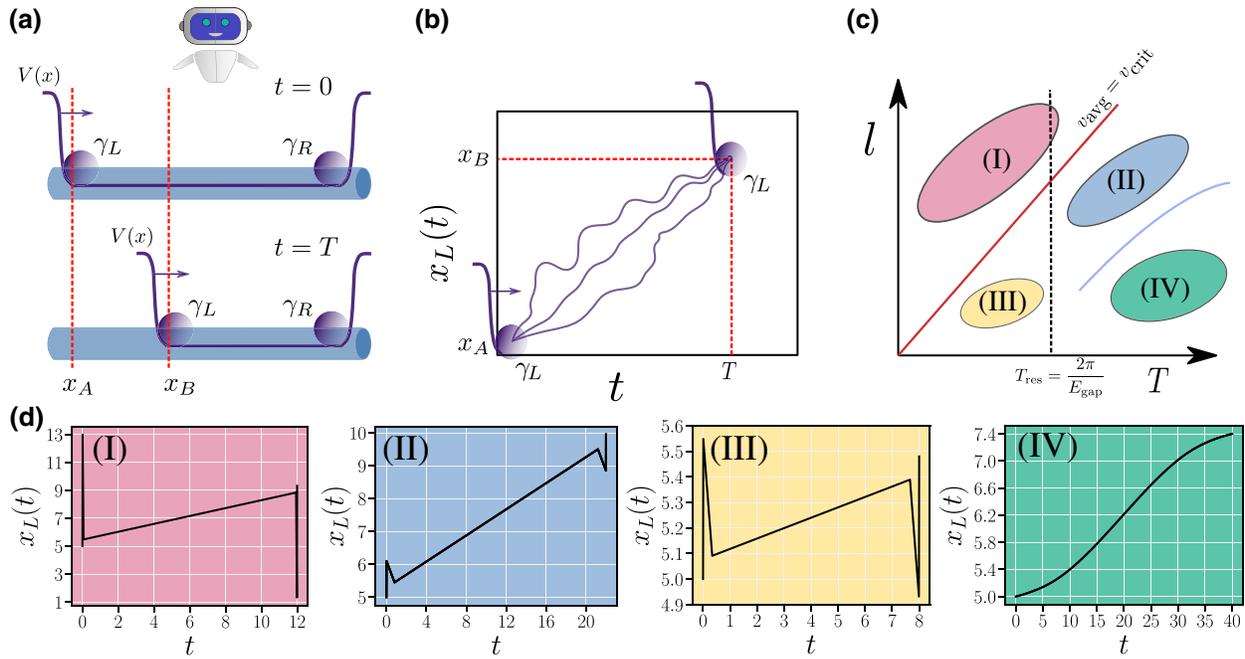


FIG. 1. (a) In the Majorana game, the agent moves the left Majorana γ_L from a position x_A to a position x_B by tuning the external potential profile $V(x, t)$. (b) The agent attempts different movement paths and tries to find the optimal one that minimizes the infidelity $\mathcal{I}(T)$. (c) Different Majorana control regimes and their corresponding strategies [boxes in (d)]: (I) critical regime ($v_{\text{avg}} > v_{\text{crit}}$), (II) subcritical regime ($v_{\text{crit}}/2 < v_{\text{avg}} < v_{\text{crit}}$), (III) short time T , low-velocity regime, and (IV) (super)adiabatic regime (long time T and low-velocity v_{avg}). (d) In regimes (I)-(II)-(III) we find the jump-move-jump (JMJ) strategy whereas in regime (IV) we recover a smooth superadiabatic protocol. The infidelities are $\mathcal{I}(T) = 0.3575$ in regime (I), $\mathcal{I}(T) = 0.1589$ in (II), $\mathcal{I}(T) = 0.0057$ in (III) and $\mathcal{I}(T) = 0.0005$ in regime (IV); the parameters for the JMJ strategies displayed in regimes (I), (II), (III) can be found in Appendix G and we set $x_L(0) = x_A = 5.0$ to ensure a smooth potential profile $V(x, t)$.

approach due to their nonlocal nature and anticipated topological protection against local errors [25–30]. The key to manipulating this topologically encoded quantum information is the development of protocols to transport the Majoranas, as will be necessary for their braiding and measurement. While the simplest approaches proposed to move the Majoranas adiabatically, decoherence processes such as quasiparticle poisoning [31], motivate the need for faster and more efficient control protocols.

In this context, machine learning (ML) offers a powerful and unifying approach to the study, design, benchmarking, and control of quantum systems and devices. Motivated by a series of remarkable technological breakthroughs in research areas as diverse as computer vision [32], natural language processing [33], and genomics [34], physicists have started to explore the potential of ML for fundamental research [35]. In particular, researchers interested in quantum many-body physics have initiated the development of a machine-learning perspective on the many-body problem [36] with recent advances such as neural-network representation of quantum states [37–48], neural-network tomography [49–51], machine-learning classification of phases [38], as well as advances in quantum chemistry [52–54], density-functional theory

[55,56], and the acceleration of Monte Carlo simulations [57–59], preparation of quantum states [60,61], quantum feedback [62], among many others [36,63]. ML applications in the context of quantum control [64–67] are largely based on reinforcement-learning (RL) techniques, which are the key ingredient behind major artificial intelligence breakthroughs in game play [68,69]. A RL perspective on quantum control has been developed in the context of quantum-state preparation [64], where the authors re-examine the quantum-state preparation in terms of a modified version of the Watkins Q -learning algorithm [70]. Likewise, Ref. [65] considers the problem of gate synthesis from a RL viewpoint.

In this work we exploit differentiable programming (DP) [71] and natural evolutionary strategies (NES) [72] to find optimal strategies to transport Majorana zero modes. We reformulate the complex task of transporting a Majorana zero mode as a quantum control optimization problem amenable to advanced ML techniques. This can be most easily understood from a reinforcement-learning perspective, where the movement of the Majoranas can be recast as a “game”, see Fig. 1. In this game the player (agent) has to move a Majorana from a position x_A to a position x_B in a fixed amount of time T . At the end of the game the

agent gets rewarded depending on how well the Majorana reached its target state at position x_B . The objective of the agent is to learn the best strategy (path of the Majorana) that maximizes the reward.

Our machine-learning techniques discover a novel and counterintuitive approach to transporting Majorana zero modes, here dubbed the *jump-move-jump* approach and exemplified in Fig. 1, which yields a high-fidelity control significantly superior in terms of speed and quality to previous strategies applied to this problem [73–75]. The jump-move-jump protocol, which is relevant for transporting Majoranas over a large distance l in a short time T , makes use of pulselike jumps at the beginning and end of the protocol [see regions (I)–(III) of Fig. 1], with a period of nearly constant velocity between the jumps. This is particularly interesting given the previous results for other models showed that bang-bang protocols [74] were the most efficient approach to this form of quantum control.

Using the insight gleaned from our ML-inspired control protocols, we construct the core strategy for these paths and provide a theoretical understanding for the protocol by analyzing the motion of the Majoranas in a moving frame. We find that these protocols simultaneously use the stability of the system at finite velocities together with the fact that small sudden jumps in wall position do not significantly reduce the ground-state fidelity. In contrast, in the limit where there is a significant total time T to move a relatively small distance l , the best approach is to follow smooth protocols that follow an adiabatic path. Our ML technology recovers these protocols without any prior knowledge [see region (IV) of Fig. 1]. In addition, we verify that the jump-move-jump protocol retains its usefulness in the presence of interactions or disorder, and also corroborate its high efficacy using the more experimentally relevant proximity-coupled semiconductor nanowire model.

We structure this paper as follows: in Sec. II we introduce the setup of the Majorana control problem. In Sec. III we introduce the optimization methods, DP and NES, for which we provide the optimization results in Sec. IV, which we benchmark against a standard simulated annealing method. In Sec. IV we also highlight some of the advantages of the DP and NES methods compared to others. We describe and analyze the physical mechanisms behind the jump-move-jump strategy in Sec. V. Then we provide some numerical results of the robustness of the jump-move-jump strategy with respect to interactions or disorder in Sec. VI before concluding our work.

We also include several appendices where we provide extra details regarding the moving frame in Appendix A, ML methods in Appendices B, C, and E, extra results for the proximity-coupled semiconductor nanowire in Appendix F, and the robustness of the JMJ strategy with respect to interactions and disorder in Appendix H. Finally

in Appendices G and I we discuss extra details for the analysis of the JMJ strategy.

II. SETUP OF THE MAJORANA CONTROL PROBLEM

To model the movement of the Majorana zero modes in superconducting nanowires we use a one-dimensional (1D) Kitaev chain [25] described by the Hamiltonian

$$\begin{aligned} \mathcal{H}(t) = & - \sum_{x=1}^N [\mu(x) - V(x, t)] (c_x^\dagger c_x - 1/2) \\ & - w \sum_{x=1}^{N-1} (c_x^\dagger c_{x+1} + \text{h.c.}) + \Delta \sum_{x=1}^{N-1} (c_x^\dagger c_{x+1}^\dagger + \text{h.c.}), \end{aligned} \quad (1)$$

where $c_x^{(\dagger)}$ are fermionic annihilation (creation) operators, $\mu(x)$ the onsite chemical potential, $V(x, t)$ a time-dependent external potential, w the hopping amplitude, and Δ the p -wave superconducting gap. This model can be realized effectively in variety of setups by proximity coupling to a conventional s -wave superconductor [15–17, 19, 20, 76–78]. When $|\mu| \geq |2w|$ the gap in the energy spectrum closes inducing a transition to a topological trivial phase [25, 79]. For an open chain, a pair of Majorana zero modes are found to reside on the edges of the wire if the bulk is topologically nontrivial. The existence of such Majorana modes implies the existence of a degenerate space of ground states. The ground state of the system is protected by a robust topological energy gap, which in the continuum limit of the model is given by $\min[\mu_c, \Delta_c k_F]$. As a consequence of this protection, two pairs of Majorana zero modes can be used to produce a qubit whose information content is protected nonlocally [27].

Manipulation of the information encoded in the ground state requires the braiding of the Majorana quasiparticles [18, 26, 28, 29, 80–83], while staying as much as possible within the degenerate ground-state space. To achieve this, a generic strategy is to try to perform the braiding adiabatically. An adiabatic path generically must have a total time T larger than the inverse gap; in practice, the size of the topological gap makes such a strategy prohibitively slow in view of numerous sources of decoherence in proximity-coupled setups [31, 84–91].

In light of this, our strategy in what follows (see also Refs. [73, 75, 92]), is to attempt to move Majoranas as quickly and efficiently as possible. To this end, and to have smooth control over the position of the Majorana bound states on the lattice, we encode the external potential profile as

$$V(x, t) = V_{\text{height}}\{f[x - x_L(t)] + f[x_R - x]\}, \quad (2)$$

where V_{height} is the maximum height of the outer potential walls and $f(x) = 1/[1 + \exp(x/\sigma)]$ is a sigmoid function shifted by the wall positions $x_{L(R)}$ [Fig. 1(a)]. The Majoranas are localized at the outer edges of the potential profile, which can be seen as domain walls between topological and nontopological phases since $V_{\text{height}} \gg 2w$. The position of these domain walls can be controlled through $x_{L(R)}$; to move the left Majorana we give the agent control over $x_L(t)$ as a function of time.

The motivation for this specific form of potential profile comes from experimental proposals for moving the Majorana by the so-called piano-key potentials [18,73,93]. In this proposal the position of the domain wall in the wire is controlled by applying gate electrodes [94]. At variance with protocols found in Ref. [74], the presence of nonlinearities in the potential profile implies that bang-bang protocols, which are expected for linear control [95,96], may no longer be optimal. As a consequence of this, it is conceivable that the space of possible solutions is expanded in our setting.

The Majorana game starts at $t = 0$ with the system in the ground state of Eq. (1) with $x_L(0) = x_A$. The aim is to reach the ground state $|\Psi_B\rangle$ with $x_L(T) = x_B$ located at a distance $x_B - x_A = l$ spending a total time T . A natural choice to quantify the accuracy of the protocol is the infidelity

$$\mathcal{I}(T) = 1 - |\langle \psi_B | \mathcal{T} e^{-i \int_0^T \mathcal{H}(t) dt} | \psi_A \rangle|^2 \equiv 1 - \mathcal{F}(T). \quad (3)$$

Here, $\mathcal{H}(t)$ is the time-dependent Hamiltonian that describes the Majorana wire setup, Eq. (1), during the protocol. Whereas $\mathcal{I}(T) = 0$ means that we have fully preserved the encoded quantum information, $\mathcal{I}(T) = 1$ implies that information has been completely lost.

A key timescale related to the control problem is $T_{\text{res}} = \frac{2\pi}{\Delta k_F}$, which naturally arises from the response of the system to boundary-wall oscillations (see Ref. [75] and Appendix A). T_{res} corresponds to the time above which one can make changes slowly enough for genuinely superadiabatic motion [73,97,98]. Superadiabaticity in this scenario allows the static ground state to be accelerated into the ground state of a moving frame, provided the transition is done slowly enough, i.e., in times large compared to T_{res} . Another important additional physical constraint is the presence of a critical velocity $v_{\text{crit}} = \Delta$ [73,92] above which the moving-frame Hamiltonian becomes gapless (Appendix A) and the ground states lose their topological protection.

Based on this we divide up the Majorana control problem into four different regimes [(I)–(IV)] [see Fig. 1(c)].

- (a) Regime (I) corresponds to the critical regime in which the Majorana must move on average $v_{\text{avg}} = l/T$ above the critical velocity. In this regime the ground-state fidelity is expected to rapidly decrease to zero.

- (b) Regime (II) is the subcritical regime for which the velocity is on average close to but nonetheless below v_{crit} . This regime is open ended in both time T and length l . The key feature distinguishing this regime from regime (IV) below is the character of the found optimal protocols.
- (c) Regime (III) is again a subcritical regime, defined for the times that are short with respect to the resonance time T_{res} but also has low velocity. This region cannot be used to efficiently move Majorana states over long distances, but we expect it to be relevant for braiding protocols based on small relative movements that change the effective couplings between Majoranas.
- (d) Finally, regime (IV) is the adiabatic regime in which we are above T_{res} and we have sufficient time to expect that slow ramp-up and ramp-down protocols [73,75] from earlier studies to be optimal. Ideally one would always like to be in this regime, however a gradual build up of noise and decoherence may make it necessary to get things done more quickly.

A. Relevance to other braiding schemes

It is worth mentioning here that the aforementioned trade-off between adiabaticity and the need to perform operations quickly has led to the emergence of other approaches to braiding. These schemes seek to circumvent the need to manipulate and shuttle Majorana's over excessively long distances. This includes local [18,99,100] and nonlocal [101,102] Majorana coupling in wire networks, and the implementation of one-way computation schemes [103,104] via projective charge measurements [30,105–109].

All of these methods still require some level of adiabatic control, or some effective short cutting thereof (see, e.g., Refs. [110,111]) and there can also be some additional downsides. The interaction and coupling schemes, for example, require precise control of the couplings between neighboring Majoranas and the loss of some topological protection [87,112]. In implementations of wire measurement-only schemes one must also tune a coupling between wire and a quantum dot [30,107] and accept state manipulation that is inherently probabilistic [109].

The phase diagram that we outline above is directly applicable to the local coupling schemes, where the relative Majorana position can be seen as a proxy for the coupling strength. A similar proxy may also hold for measurement-only schemes via the distance between topological boundary and quantum dot, although the connection here is harder to make directly because one should also account for fermion number conservation [30,105–108]. That said, the DP and NES methods we apply below do have direct relevance for even more sophisticated numerical treatments of this problem. We

refer to our concluding remarks for further discussion of this point.

III. MACHINE-LEARNING METHODS

In our study we apply two machine-learning strategies, namely differentiable programming and natural evolution strategies, to minimize the infidelity, Eq. (3), with respect to the position of the domain wall $x_L(t)$ [113]. In this section we introduce these methods in a way that directly applies to our Majorana control setup. For the interested reader we provide the corresponding programming codes in Ref. [114].

A. Differentiable programming

DP is a programming paradigm to evaluate gradients of computer programs [115] that has been recently introduced as an optimization method in physics applications like tensor networks [116,117], Monte Carlo [40,118,119], and optimal control [67,120]. DP obtains numerically exact gradients with similar computational-time complexity as the forward evaluation of the computer program due to the use of backward propagation [121–123].

In our Majorana control optimization problem we are interested in the total derivative of the infidelity with respect to the control $\frac{d\mathcal{I}(T)}{dx_L(t)}$. As shown in Appendix B, the algorithm to compute $\mathcal{I}(T)$ consists of a sequence of elementary operations f_i , in DP language known as primitives, which maps the input control $x_L(t)$ to the value of the infidelity by $\mathcal{I}(T) = f_n \circ f_{n-1} \circ \dots \circ f_1[x_L(t)]$. Moreover, the derivatives of each individual operation f_i are known and the total derivative $d\mathcal{I}(T)/dx_L(t)$ can be assembled by recursively applying the chain rule with automatic differentiation (see Appendix B).

In practice, we write a code to evaluate the real-time dynamics of the Majorana induced by the time-dependent profile $V(x, t)$ such that all the individual operations are differentiable. To obtain the necessary gradients we use a language that supports automatic differentiation, which in our case is the JAX library [124]. The gradients are used to minimize the infidelity of the final state using gradient descent (GD) and ADAM [125]. For GD this is done iteratively, where at each iteration of the algorithm the protocol is updated as $x_L(t) \rightarrow x_L(t) - \nabla_{x_L(t)} \mathcal{I}(T)$. Moreover, since neural networks consist of differentiable elementary operations, we also explore representing and training the control $x_L(t)$ as the output of a neural network for which the weights are updated with the standard update scheme ADAM.

B. Natural evolution strategies

Evolution strategies are a family of black-box optimization algorithms inspired by natural evolution [72,126]. These methodologies have been recently revisited in the

context of machine learning [127], in particular, in reinforcement learning [128], where OpenAI demonstrated that a particular incarnation of the algorithm called natural evolution strategies offers a powerful alternative to popular Markov-decision process-based reinforcement-learning methods [129].

NES starts with an objective function $F(\phi)$ that acts on parameters ϕ from a population described by a distribution p_θ , where θ parameterizes the population distribution. In our work, we choose the objective function $F(\phi)$ to be the infidelity $\mathcal{I}(T, \phi)$ in Eq. (3). Depending on the setting, the parameter ϕ corresponds to either the position profile $x_L(t)$, the velocity $v_L(t) = \dot{x}_L(t)$ or the parameters of a neural network whose output is $x_L(t) = \text{NN}_\phi(t)$.

The NES algorithm proceeds to optimize the expectation value of the objective function $\mathbb{E}_{\phi \sim p_\theta} [\mathcal{I}(T, \phi)]$ over the population. We choose p_θ to be a Gaussian distribution with mean θ and diagonal covariance matrix $\Sigma = \sigma^2 I$ with $\sigma = 0.1$, i.e., $\mathcal{N}(\theta, \sigma^2 I)$. It follows that the gradient of the cost function is (see Appendix C)

$$\begin{aligned} \nabla_\theta \mathbb{E}_{\phi \sim \mathcal{N}(\theta, \sigma^2 I)} [\mathcal{I}(T, \phi)] &= \nabla_\theta \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} [\mathcal{I}(T, \theta + \sigma \epsilon)] \\ &= \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} [\mathcal{I}(T, \theta + \sigma \epsilon) \epsilon / \sigma]. \end{aligned} \quad (4)$$

This equation provides an efficient way for computing gradients without differentiation, but instead through the expectation of the objective function. Notice that $\mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} [\mathcal{I}(T, \theta) \epsilon / \sigma] = 0$, which implies the above equation is equivalent to $\mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} \{[\mathcal{I}(T, \theta + \sigma \epsilon) - \mathcal{I}(T, \theta)] \epsilon / \sigma\}$. This means that the estimation of the gradient can be seen as the finite difference of the objective function with random search [128]. To update the parameters θ , we apply gradient descent to θ with Eq. (4).

IV. MACHINE-LEARNED STRATEGIES FOR MAJORANA CONTROL

Our main objective is to investigate the use of DP and NES for the motion of the Majorana modes in the four different movement regimes (I)–(IV). We choose four different points (l, T) (which fixes $v_{\text{avg}} = l/T$) each of which belongs to a different regime. For each of these points we apply the optimization algorithms to minimize the infidelity $\mathcal{I}(T)$ in Eq. (3) with respect to the control of the domain wall. The control can be parameterized by either the wall position $x_L(t)$, the wall velocity $v_L(t)$, or a neural network $x_L(t) = \text{NN}_\theta(t)$, where θ represent the parameters in the neural network. We test all these different parameterizations (see Appendix D) and in the following we focus on the parameterization choices that give the best performance (the lowest infidelity values).

For the DP optimization, we first parameterize the wall position by a neural network $x_L(t) = \text{NN}_\theta(t)$ and use ADAM [125] to optimize the parameters. We then take

is dressed to various degrees with high-frequency oscillations. We observe that the velocity in the middle part of the protocols is typically lower than v_{crit} even for regime (I), where the total average velocity v_{avg} (including the sudden movements) is above v_{crit} . The size and character of the additional oscillations largely depends on the optimization strategy being used.

The infidelity values we find in regimes (I)-(II)-(III) are significantly better compared to strategies like linear motion $x_L(t) = v_{\text{avg}}t$ or bang-bang as shown in Appendix D. In the critical regime (I) we get an infidelity of about $\mathcal{I}(T) \approx 0.35$ compared to $\mathcal{I}(T) = 0.47$ for a linear protocol whereas in regime (II) we get $\mathcal{I}(T) \approx 0.15$ versus $\mathcal{I}(T) = 0.22$ for linear. In the low average velocity regime (III) we get infidelities as low as $\mathcal{I}(T) \approx 0.006$ while linear motion gives $\mathcal{I}(T) = 0.012$. The infidelity value improvement in regime (I) is rather surprising given that the Majorana moves on average above the critical velocity. Below, we explain why the jumps in this scenario are particularly beneficial.

In regime (IV) the results show a globally different strategy: in this regime the optimal protocol is to slowly accelerate the Majorana up to some finite velocity and then slowly decelerate back down to the target position. This type of protocol was discussed in earlier work [73,75] in regimes where there is enough time to accelerate to a moving frame, i.e., regime (IV). We note that due to the nearly adiabatic motion in regime (IV), the values of infidelity are the lowest $\mathcal{I}(T) \approx \mathcal{O}(10^{-4})$. These infidelity values make regime (IV) optimal for braiding of Majoranas in an ideal experimental setup. However, due to imperfection in the devices and their coupling to a noisy environment, balancing the infidelity associated with shorter protocols with respect to the infidelity induced by dephasing in a longer time protocol makes controlling Majorana movement in faster regimes relevant for near-term experiments.

Compared to previous studies, however, [73,75,89] we find that the starting velocity $v_L(t=0)$ of the protocols in regime (IV) can be finite if one allows for small but finite infidelity values $\mathcal{I}(T) \approx \mathcal{O}(10^{-4})$. In the obtained protocols, for example, we have $v_L(t=0) \neq 0$ followed by an approximately smooth gradual build up and down of the velocity. Moreover, the optimization techniques with which these protocols are obtained have an additional advantage that they can fine tune the protocols up to a finer level (bigger search space) than the previously studied parameterized adiabatic protocols [75], as we show in Appendix D. This advantage of our methods might be particularly beneficial for finding optimal protocols in the presence of disorder in the wires as discussed briefly in Sec. VIB and Appendix H.

Note that although we focus on the Kitaev-chain model in Eq. (1) the results and conclusions discussed here remain true on the more realistic proximity-coupled nanowire models, where the JMJ protocol is optimal in the

nonadiabatic regime (see details of the proximity-coupled nanowire simulations in Appendix F). Additionally, the results on the proximity-coupled nanowire in Appendix F are compatible with a smooth protocol in the adiabatic regime, as well as consistent with the results on the Kitaev chain in the large Zeeman field and large strong spin-orbit-coupling limits.

A. Comparison between simulated annealing and DP and NES optimization

We now benchmark the results of the DP and NES algorithms against the standard simulated-annealing (SA) method following Ref. [74] closely. In the SA method the wall velocity $v_L(t)$ of the domain wall is iteratively updated by choosing two random intervals of length Δt of which one interval is increased by Δv and the other is decreased by Δv . The new infidelity \mathcal{I}_i is calculated for the updated velocity profile and the move is accepted with a probability $e^{-\delta\mathcal{I}/T_{\text{SA}}}$, where $\delta\mathcal{I} = \mathcal{I}_i - \mathcal{I}_{i-1}$ is the difference in the infidelity with respect to previous iteration. The annealing temperature T_{SA} is slowly cooled down to zero. The results of this benchmark in each of the four regimes are shown in Figs. 2(i)–2(l) and are qualitatively similar to the results obtained with NES and DP [Figs. 2(a)–2(h)].

In practice, we find that SA is significantly more computationally demanding than NES or DP since a typical SA simulation entails the evaluation of the many-body infidelity for each SA update step. To obtain results with comparable infidelities, the SA simulation requires a total of 10^5 evaluations of the infidelity while for DP we need only 440 update steps each involving a single infidelity and gradient evaluation (the latter requiring similar computational complexity as the infidelity calculation); the NES algorithm requires 50 update steps each with 100 parallel evaluations to reach convergence.

This difference may be partially attributed to the fact that SA does not take advantage of any gradient signal. Beyond these practical observations, a careful scaling analysis of the convergence of these methods requires an analysis of the convexity properties of the infidelity landscape as a function of the control parameters. This can be either done analytically for simple infidelity landscapes [131], but may require a numerical investigation for control problems exhibiting a glassy landscape [132].

Overall, we highlight that DP offers a powerful tool for quantum control as long as an accurate and differentiable physical model is available. It is found that direct gradient-based methods are usually more stable and efficient than RL [133]. In contrast, NES can optimize nondifferentiable and discrete control protocols, both of which remain challenging for DP. NES can also be directly applied to experimental settings as long as a suitable objective function, such as the expectation value of a Hermitian observable or the fidelity, is available. For instance,

recent experiments demonstrate the teleportation of Majorana modes in a quantum simulation of a small Majorana chain [134], including access to the fidelity of the protocol. This possibility makes NES a viable tool for controlling and designing quantum simulations of Majorana modes.

Additional discussions detailing connections, comparisons, advantages, and disadvantages of the machine-learning approaches in relation to other advanced methods used in quantum control are presented in Appendix E.

V. JUMP-MOVE-JUMP [JMJ] MAJORANA CONTROL STRATEGY

The main features of the ML strategies in regimes (I), (II), and (III) can be encapsulated in a simple model strategy, the jump-move-jump strategy as shown in Fig. 3(a), amenable to semianalytical and numerical analysis. This strategy consists of two dressed jumps in the position of the domain wall at the beginning and end of the protocol interluded by a period of motion at a constant velocity v . The initial dressed jump comprises an instantaneous forward jump over a distance $\Delta x_{\text{forward}}$ followed by a rapid backward movement over a distance Δx_{back} in a time Δt_{back} . Similarly, the last pulse starts with a rapid backward movement Δx_{back} followed by a forward jump over a distance $\Delta x_{\text{forward}}$. In what follows, we assume that jumps at the beginning and end of the protocol are symmetric and are described by the same parameters.

The simplicity of this model allows us to develop an understanding for the key mechanisms behind the ML strategies and estimate the value of the infidelity of the JMJ strategy for a wide range of parameters l, T in regimes (I)–(III) [see Fig. 5 and Fig. 3(b)]. In Sec. VA below we first argue that when we disregard the backward movements, i.e., $\Delta x_{\text{back}} = 0$, the overall strategy can be explained via a trade-off (Fig. 4) between the amount of fidelity loss due to the boundary jumps and the loss due to sudden changes in velocity.

In Sec. VB we also explore how this mechanism can be better controlled with the backward movements $\Delta x_{\text{back}} \neq 0$, which, for a certain subset of parameters $(\Delta x_{\text{forward}}, \Delta x_{\text{back}})$, leads to a model strategy with infidelities that compare competitively with the best machine-learned protocols [see the discussion in Appendix G and Figs. 1(d) and 2]. Our analysis shows that this dressed protocol allows one to better target the ground state of the system in a moving (constant velocity) frame.

A. Bare jump-move-jump ($\Delta x_{\text{back}} = 0$)

For the case $\Delta x_{\text{back}} = 0$ we minimize the infidelity, Eq. (3), with respect to the velocity v of the bulk of the protocol, which fixes the instantaneous forward jump to $\Delta x_{\text{forward}} = (l - vT)/2 \equiv \delta$. To make analytic analysis simpler we focus on the case that both the left and

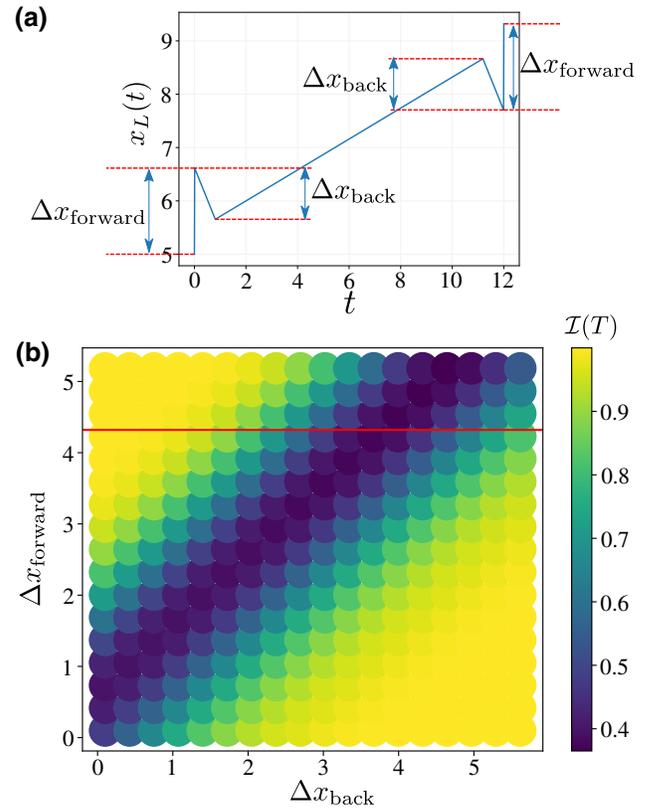


FIG. 3. (a) Setup of the ML-inspired simple control strategy consisting of two edge pulses with a forward jump $\Delta x_{\text{forward}}$ and backward jump Δx_{back} together with a constant linear motion in between the dressed jumps. (b) Infidelity $\mathcal{I}(T)$ surface plot for the scans over the free parameters $\Delta x_{\text{forward}}$ and Δx_{back} of the model strategy with $\Delta t_{\text{back}} = 0.07w^{-1}$ and $\Delta t_{\text{forward}} = 0.01w^{-1}$ in regime (I) ($l = 4.32, T = 12$). The red line indicates the line where the size of the forward jump $\Delta x_{\text{forward}}$ is equal to the total movement length l . The parameters of the Majorana control setup are the same as in Fig. 2. The optimal dressed jumps appear at a diagonal set of parameters $\Delta x_{\text{forward}} - \Delta x_{\text{back}} \sim C$, where $C \approx 0.96$. We note that the best protocol has a jump size bigger than the movement length $\Delta x_{\text{forward}} > l$, which means that it jumps over the target position $x_L(T) = x_B$ and then jumps back within the range $x_A < x_L < x_B$, as can be seen in Fig. 1 (d) [for regimes (I) and (III)].

the right domain-wall positions are being evolved simultaneously with the JMJ strategy. This means that the right wall position x_R in Eq. (2) becomes time dependent via $x_R(t) = x_L(t) + C_x$ with C_x a fixed constant. This two-wall scenario makes it possible to evaluate the strategy in the moving frame basis (Appendix A), which allows the key rationale behind the JMJ strategy to be revealed.

To evaluate this strategy we expand the infidelity after the first jump $x_L = x_A \mapsto x_L = x_A + \delta$ in terms of the eigenbasis $|\psi_{A+\delta}^i\rangle$ of the Hamiltonian, Eq. (1), with the

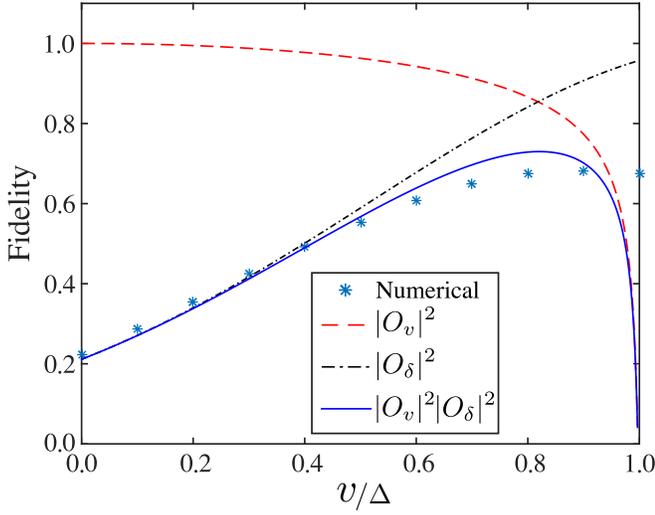


FIG. 4. The jump-move-jump protocol maximizes the product of the two functions O_δ [Eq. (8)] and O_v [Eq. (9)]. Here $l = 4.32$ and $T = 12$ [regime (I)] and we use $s = 2.45$ and $\beta = 0.065$. The ground-state projection analysis (solid blue) captures the key features of the ML-inspired (one-wall) JMJ strategy (blue stars) for low velocities v but tends to overestimate the penalty for moving at velocities near v_{crit} .

wall at position $x_L = x_A + \delta$ giving

$$\mathcal{I}(T) = 1 - \left| \sum_i \langle \psi_B | U(T) | \psi_{A+\delta}^i \rangle \langle \psi_{A+\delta}^i | \psi_A \rangle \right|^2 \quad (5)$$

in which $U(T) = \mathcal{T} e^{-i \int_0^T H(t) dt}$ is the time-ordered evolution operator with $H(t)$ following the time dependence of the strategy. We break this equation up into two separate terms

$$\mathcal{I}(T) = 1 - \left| \langle \psi_B | U(T) | \psi_{A+\delta}^0 \rangle \langle \psi_{A+\delta}^0 | \psi_A \rangle + \sum_{i>0} \langle \psi_B | U(T) | \psi_{A+\delta}^i \rangle \langle \psi_{A+\delta}^i | \psi_A \rangle \right|^2, \quad (6)$$

which allows for an approximate analysis of the separate contributions. Here we focus on the ground-state contribution (first line), making the assumption that the second line is small in comparison.

To approximate the ground-state contribution we insert projections onto the ground state $|\psi_v^0\rangle$ of a moving-frame Hamiltonian $H(v)$ with a velocity v equal to the bulk velocity of the strategy and finally a projection onto the ground state $|\psi_{B-\delta}^0\rangle$ of the Hamiltonian, Eq. (1), with the wall at position $x_L = x_B - \delta$ (just before the final jump) resulting in

$$\mathcal{I}(T) \approx 1 - \left| \langle \psi_B | \psi_{B-\delta}^0 \rangle \langle \psi_{B-\delta}^0 | \psi_v^0 \rangle \times \langle \psi_v^0 | U(T) | \psi_v^0 \rangle \langle \psi_v^0 | \psi_{A+\delta}^0 \rangle \langle \psi_{A+\delta}^0 | \psi_A \rangle \right|^2. \quad (7)$$

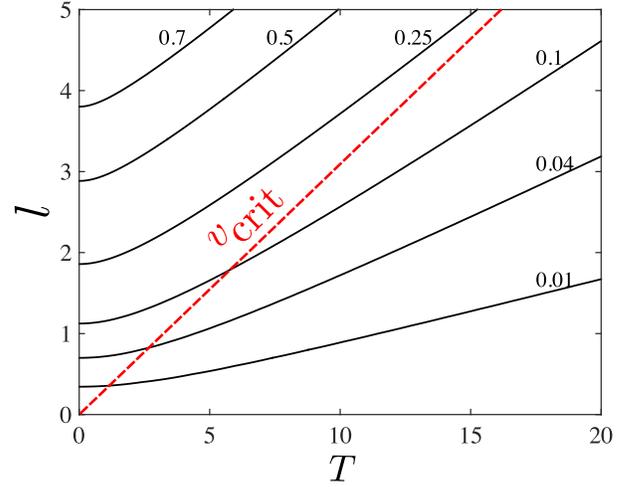


FIG. 5. Contours showing the infidelity \mathcal{I} for the bare (one-wall) jump-move-jump strategy [Eq. (10)] as a function of distance l and time T using the same fitted parameters as Fig. 4. The red dashed line indicates when $v_{\text{avg}} = l/T = v_{\text{crit}}$, which can be used to distinguish the different Majorana motion regimes as defined in Fig. 1 (c). The jump-move-jump protocol allows for low infidelities even in cases where the average velocity exceeds v_{crit} [regime (I)].

The $\langle \psi_{A(B)\pm\delta} | \psi_{A(B)}$ represent the initial and final jumps in position space of size δ , which by fitting to our numerical model can be characterized as

$$O_\delta = |\langle \psi_{x_L} | \psi_{x_L+\delta} \rangle|^2 \sim \exp(-\delta^2/s^2), \quad (8)$$

where $s \sim \epsilon + \alpha \lambda_F$ with λ_F the Fermi wavelength and the fitting parameters $(\epsilon, \alpha) = (-0.33, 0.44)$ for one wall and $(-0.12, 0.3)$ for two walls when $\Delta = 0.3$. In Appendix I we discuss this in more details, showing also how the fit s scales with the coherence length ξ of the Majorana mode.

The amplitude $\langle \psi_v^0 | \psi_{A+\delta}^0 \rangle$ represents the overlap between the static ground state and that of a moving frame with a constant velocity v . We estimate this as follows:

$$O_v = |\langle \psi_{x_L}^0 | \psi_v^0 \rangle|^2 \sim 1 - \beta \left(\frac{v}{\Delta} \right)^2 \gamma, \quad (9)$$

with $\gamma = 1/\sqrt{1 - v^2/\Delta^2}$. One way to argue for this form is to use the results of Ref. [75] that showed that the energy of the moving-frame ground state with respect to the static frame goes as $E \sim \gamma M v^2$, where $M \propto k_F/\Delta$. The overlap can be related to this energy via $O_v \sim 1 - E/k_F \Delta$. We arrive at a value of $\beta \sim 0.13$ by fitting Eq. (9) to moving-frame (two-wall) numerical simulations for $v \ll v_{\text{crit}}$. For single-wall motion we can calculate a value of $\beta \sim 0.065$ by slowly accelerating the wall up to $v \ll v_{\text{crit}}$ and comparing the evolved state with the instantaneous ground state. Note that this analysis also holds for the more general proximity-coupled semiconducting nanowire with different values for fitting parameters α and β .

If we assume a symmetric strategy (e.g., jumps of the same size at the beginning and end) then the fidelity function $\mathcal{F}(T) \approx |O_\delta O_v|^2$ [where $\mathcal{I}(T) = 1 - \mathcal{F}(T)$] to be maximized is approximated as

$$\mathcal{F}(T) \approx \exp\left(-\frac{(l-vT)^2}{2s^2}\right) \times \left[1 - \beta\gamma \left(\frac{v}{\Delta}\right)^2\right]^2, \quad (10)$$

where we substitute $\delta = (l-vT)/2$. The first function is maximized by making v as large as possible while the second prefers to have small v , with a very severe penalty kicking in as v approaches the critical velocity Δ . In Fig. 4 we show how this function behaves for the protocol in regime (I), which indicates that the analysis accurately captures the behavior of the fidelity at small v/Δ but also displays overall qualitative agreement with the numerical results for a wide range of values of v/Δ , including the presence of a maximum in the fidelity.

An important part of the strategy are the jumps at the beginning and end of the protocol that allow for a constant middle motion at lower velocities. In Fig. 5 we show how the infidelity behaves for a range of movement lengths l and times T obtained from maximizing Eq. (10) with respect to v . By comparing with Fig. 1(c) we see that in regime (I) the bare JMJ strategy gives infidelities of $\mathcal{I}(T) > 0.25$, in regime (II) (between the dashed lines) about $0.04 < \mathcal{I}(T) < 0.25$ and in the short l -short T regime (III) (below the blue dashed line) performs particularly well with infidelities $\mathcal{I}(T) \sim \mathcal{O}(10^{-2})$.

This method is only effective, however, if the jumps needed are not too large [$\delta \sim \mathcal{O}(\lambda_F)$]. Interestingly, as the overlap behavior of O_v is relatively unaffected by changes to the chemical potential, this suggests that a chemical potential nearer to the bottom of the band, $\mu \rightarrow -2w$ can help to extend the range of the protocol.

B. Dressed jump-move-jump ($\Delta x_{\text{back}} \neq 0$)

On top of the bare jump-move-jump structure, the ML algorithms dress the protocols with additional forward and backward movements. In this section we show that the primary purpose behind these additional motions is to better target the moving-frame ground state in the *move* part of the *jump-move-jump* protocol.

To proceed we focus on the simple upgraded JMJ protocol that also allows backward motions over distances $\Delta x_{\text{back}} \neq 0$, see Fig. 3(a). Our first key observation is that the optimal protocols here tend to choose a combined jump size $\Delta x_{\text{forward}} - \Delta x_{\text{back}} \sim C$, see Fig. 3 (b), where C is a fixed constant. This indicates the same primary goal as the JMJ strategy: to allow a period of optimal and roughly constant subcritical motion.

To understand why the backward motion is a further benefit we examine the instantaneous infidelity in both the static and moving frames for a series of protocols

with the fixed bulk velocity v , see Fig. 6(a). The instantaneous infidelity is defined as $I(t) = 1 - |\langle \psi(t) | \psi_{\text{ins}}^0(t) \rangle|^2$ with $|\psi(t)\rangle = \mathcal{T} e^{-i \int_0^t H(t') dt'} |\psi_A\rangle$ in which $H(t)$ follows the time dependence of the dressed JMJ strategy and $|\psi_{\text{ins}}^0(t)\rangle$ corresponds to the instantaneous ground state (at the fixed

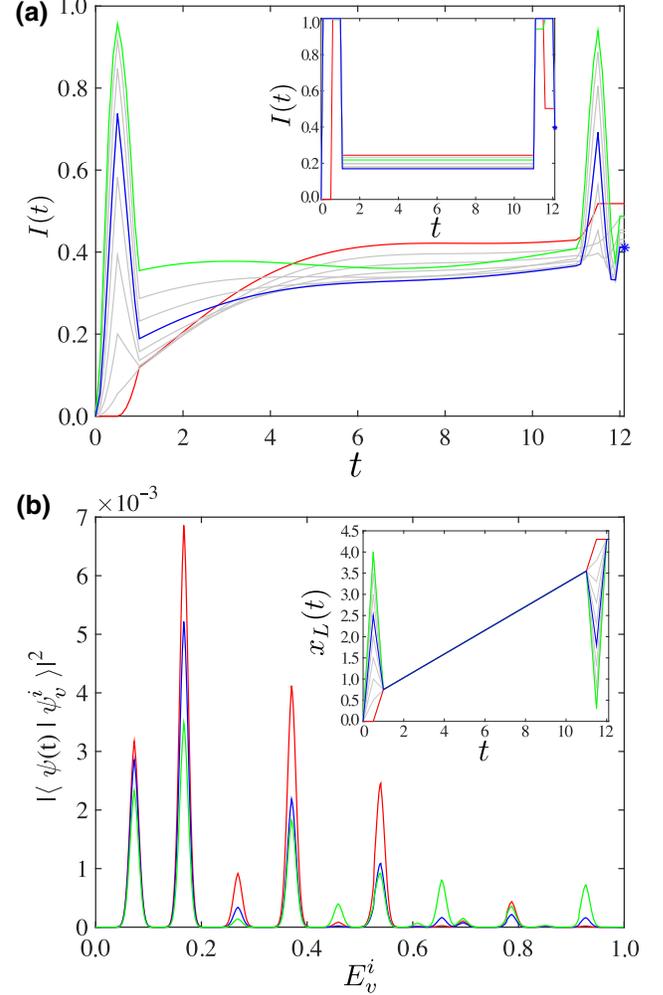


FIG. 6. (a) Instantaneous infidelity $I(t)$ with respect to the static (lab) frame ground state (main panel) and moving-frame ground state (inset) as a function of the protocol time t for various dressed jumps in regime (I), which are depicted by the different colors (red, blue, green, gray) and shown in the inset of (b). The optimal dressed jump (blue) has the minimum infidelity with the moving-frame ground state and becomes the strategy with the minimum static frame infidelity after about $t \approx 4$. (b) The time-evolved state $|\psi(t)\rangle$ expanded in the single and two particle excitations $|\psi_v^i\rangle$ of the moving-frame basis with energies E_v^i at $t = T/2$ for a forward jump (red), an optimal dressed jump (blue), and nonoptimal dressed jump (green) for regime (I) as given in the inset. The occupation probabilities (δ functions) are convoluted for visualization purposes. The forward jump strategy (red) occupies the excitations close to the bottom of the band in the moving frame ($E_v^i \approx 0.08$) more compared to the dressed jumps (blue and green). In this figure all times are in units of $1/\omega$.

time t) of the static frame Hamiltonian $\mathcal{H}(t)$ [Eq. (1)] or the moving-frame Hamiltonian $H_v(t)$ [Eq. (A2)]. The static frame perspective shows that the optimal jump finds a state that slowly becomes the lowest infidelity state after some time. However, a more revealing picture emerges if we examine the same set of protocols in a moving frame (inset). Here we see that the objective of the initial jumps is to maximize the overlap with the ground state of the moving frame.

Using the same methodology we can also reveal what moving-frame bulk excitations $|\psi_v^i\rangle$ with many-body energy E_v^i (restricting here to single- and two-particle excitations only) the dressed protocol occupies during the evolution, see Fig. 6(b). For protocols close to the bare JMJ strategy (shown in red) one finds that the dominant excitations have energies close to the moving-frame bulk gap ($E_v^i \approx 0.08$). The dressed jumps (in blue and green), however, are able to lower the amplitudes of these excitations, but at the cost of also exciting some higher-energy quasi-particles. For dressed jumps that are too large (in green) these higher-energy excitations eventually dominate and one gets an increase in the ground-state infidelity.

Before ending this section, we now briefly summarize why similar JMJ strategies work in regimes (I)–(III). The first advantage of the JMJ strategy is that, with abrupt initial and final costs, it allows for an effective reduction in the distance l . This is especially obvious in the case of regime (I), where, by definition, simple linear motion implies a need to move above the critical velocity.

The other crucial aspect of the dressed protocols is that they can better target the ground state of the moving system through the combination of jumping and moving. This allows the system to be rapidly accelerated to a near constant velocity frame where there is no additional infidelity cost. For a fixed distance l , the JMJ strategy therefore generally allows for a reduced time by cutting out the slow acceleration parts of regime (IV), allowing regime (III), regime (II) and, to a lesser extent, regime (I) exploit the property of superadiabaticity.

Of course it is always possible to define scenarios deep in regime (I) where the distance and time constraints would lead to large infidelity even for the dressed JMJ protocols. However, they do allow, under the same spatial and temporal constraints, for Majoranas to be moved more efficiently than that of linear motion or slow ramp-up ramp-down protocols.

VI. ROBUSTNESS OF THE JMJ PROTOCOL WITH RESPECT TO INTERACTIONS AND DISORDER

In the space between our simplified theoretical setups [the Kitaev chain in Eq. (1) and proximity-coupled semiconductor in Appendix F] and actual topological quantum devices there are many additional layers of complexity,

and one might wonder, which aspects of the JMJ protocol remain robust. In this section, we assess the robustness of our optimal JMJ protocols obtained in the clean noninteracting system when applied in a system with interactions (see, e.g., [135–141]), and separately applied in a system with disorder, see [76,142,143]. We find that the fidelity loss due to the jumps of the JMJ protocol remains robust and that the overall protocol still outperforms the naive linear benchmark protocol significantly. Finally, we comment on the extension of our ML methods to include these more complex scenarios.

A. The effect of interactions

Interactions are modeled by adding the nearest-neighbor density-density interaction term $H_{\text{int}} = \sum_{x=1}^{N-1} u_x c_x^\dagger c_x c_{x+1}^\dagger c_{x+1}$ with interaction strength u_x to the Kitaev-chain Hamiltonian in Eq. (1), $H(t) = \mathcal{H}(t) + H_{\text{int}}$. Since H_{int} is quartic in the creation and annihilation operators, the Bogolyubov-de-Gennes (BdG) time-evolution algorithm used before (see Appendix B) cannot be applied directly and we consider a time-dependent variational principle with matrix product states (MPS TDVP) [144,145] approach to simulate the dynamics.

In Fig. 7 we show how the optimal protocol obtained in the noninteracting system in regime (I) performs for various finite interaction strengths (see Appendix H for the other regimes). From the instantaneous infidelity plot in

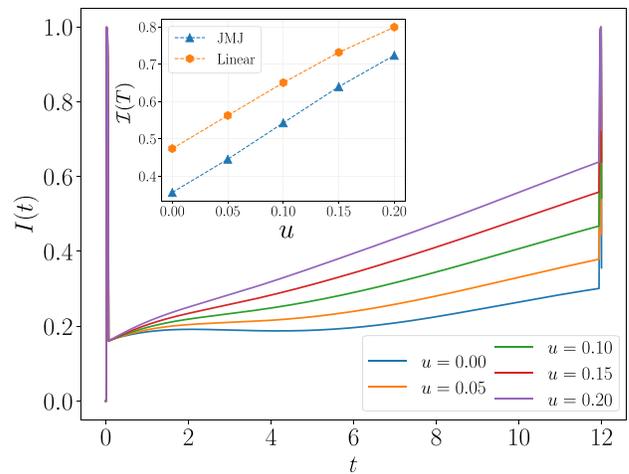


FIG. 7. Main panel: instantaneous infidelity $I(t)$ of the optimal JMJ protocol from the noninteracting system in regime (I) run in a system with several different finite interaction strengths u . The jumps are robust against interactions whereas infidelity increases during the move part of the protocol. Inset: the final infidelity values $\mathcal{I}(T)$ increase approximately linearly with interaction strength for both a linear as well as the best JMJ protocol. For all interaction strengths the JMJ protocol outperforms the linear protocol. The other parameters for these simulations are the same as in Fig. 2 and the parameters of the JMJ protocol can be found in Appendix G.

the main panel we can see that the fidelity loss due to the jumps remains robust with increasing interaction strength whereas in the bulk (move part) the loss increases. As a consequence the final infidelity $\mathcal{I}(T)$ increases with interaction strength (see the inset in Fig. 7), but the performance remains superior to a naive linear protocol at all interaction strengths. The increase of infidelity with interaction can be argued for on the basis of mean-field theory, where one expects repulsive interactions to lower the effective gap (leading to a lower critical velocity) and the appearance of nonuniformity in the effective coupling parameters near the system boundary.

Since we did not perform the optimization directly in the presence of interactions, we cannot argue that the optimal JMJ protocol from the noninteracting system is also the optimal protocol in the interacting system. We note, however, that our ML optimization methods can be directly extended for this task. Firstly, the NES algorithm is a black-box optimization method and can be combined with any method that generates the dynamics such as the MPS-TDVP algorithm. Secondly, to apply DP one needs to make the code for the MPS-TDVP algorithm differentiable. In this regard, we note that the DMRG algorithm to find ground states in interacting systems has been recently made differentiable [116].

B. The effect of disorder

Disorder in the wires can be modeled by adding a random Gaussian noise term $\tilde{\mu}(x)$ with mean 0 and standard deviation λ to the chemical potential $\mu_{\text{dis}}(x) = \mu(x) + \tilde{\mu}(x)$ [76,142,143]. It has been shown that this form of disorder leads to an increase in the effective coherence length ξ_{eff} of the Majorana, which in the continuum limit leads to $1/\xi_{\text{eff}} = 1/\xi - 1/2l_{\text{dis}}$. Here $l_{\text{dis}} = v_F^2/\lambda^2$ is the characteristic disorder length scale. When the disorder strength is such that $2l_{\text{dis}} = \xi$, the effective Majorana coherence length ξ_{eff} diverges and the topological phase is destroyed. Another consequence of disorder is that the critical velocity is significantly reduced [92] and the ground-state fidelity is expected to rapidly decrease upon smoothly moving the Majoranas over a disordered medium.

In Fig. 8 we show the performance of the JMJ protocol in regime (I) for various different disorder strengths compared to a linear benchmark protocol. Regimes (II), (III), and (IV) are detailed in Appendix H. From the averaged instantaneous infidelity $\langle \mathcal{I}(t) \rangle$ in Fig. 8 it can be seen that the infidelity associated with the jumps is robust to the presence of disorder but there is a relatively strong fidelity loss associated with the linear motion of the move part. This is in line with the expectations given the lower critical velocity and longer coherence length induced by disorder. While the final infidelity value $\mathcal{I}(T)$ of the JMJ protocol remains reasonably high in regime (I), for longer time and distance regimes [e.g., deep in regimes (II) and (IV)]

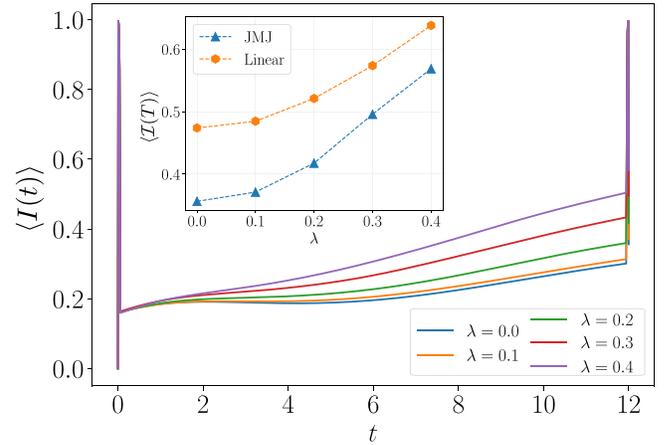


FIG. 8. Main panel: disorder averaged instantaneous infidelity $\langle \mathcal{I}(t) \rangle$ of the optimal JMJ protocol from the clean wire in regime (I) for different disorder strengths λ . The jumps are robust against weak disorder whereas the infidelity increases during the move part of the protocol. The disorder averaging is done over 500 realizations. Inset: the final infidelity values $\langle \mathcal{I}(T) \rangle$ increase approximately quadratically with disorder strength for both a linear as well as the best JMJ protocol. For all disorder strengths the JMJ protocol outperforms the linear protocol. The other parameters for these simulations are the same as in Fig. 2 and the parameters of the JMJ protocol can be found in Appendix G.

$\mathcal{I}(T)$ worsens significantly, which suggests that additional protocol optimization in the presence of disorder may be required to achieve comparable infidelities.

The protocol optimization in the presence of disorder can be achieved through the NES method since the computational complexity grows linearly with the number of disorder realizations, all of which can be simulated in parallel. The DP method can also be applied because the disorder averaging is a differentiable operation, i.e., DP can access the disorder-averaged infidelity $\nabla_{\theta} \langle \mathcal{I}(T) \rangle$ with respect to the control parameters. For example, Ref. [146, 147] considers DP for stochastic optimization tasks.

VII. CONCLUSION AND FURTHER WORK

We apply two state-of-the-art machine-learning techniques, namely differentiable programming and natural evolution strategies, to the problem of manipulating Majorana bound states in a topological superconductor. For DP we show that the entire dynamical evolution of any free-fermion system can be functionally differentiated, allowing for efficient optimization protocols. This in turn provides an ability to tackle computationally harder problems and allows the dynamical optimization to be integrated seamlessly with both direct and neural-network parameterizations of quantum control protocols.

In addition to this, we show how the Majorana control problem itself can be naturally formulated as a game. This allows the application of reinforcement-learning

approaches and NES to zero-mode manipulation. The key advantage here is both speed and flexibility compared to, e.g., the standard MC methods such as simulated annealing. Beyond the advantages observed in our numerical experiments there are a number of additional conceptual and practical ingredients, which can be taken advantage of in the reinforcement-learning setup, including a flexible exploration step, delayed rewards, partial observations, and the ease of accounting for the stochastic nature of the reward function, all of which may be particularly relevant to experimental setups.

Note that these machine-learning approaches are not specific to Majorana bound-state control, and could also be applied to different dynamical many-body problems with radically different motivations and cost functions (e.g., log likelihood [148], Fisher information [149], or the trace distance [150]).

On the theoretical level we introduce a framework for dividing up the Majorana control problem into four distinct transport regimes. This has deepened the understanding of the problem and helped to establish a connection with optimal control theory more broadly. Remarkably, our numerical simulations displays an innate awareness of this rich landscape and uncovered important hidden aspects of the underlying physics. In particular, they have identified a new class of protocol that radically outperforms other known strategies (adiabatic and bang-bang) when there are both spatial and temporal constraints. We show that this family of protocols, and the associated classification of Majorana motion, also holds for the richer proximity-coupled semiconductor nanowire model, which is directly relevant for experiments. In principle, since the defining criteria are related to topological gap, similar protocols should also apply to related models such as the Su-Schrieffer-Heeger chain [151] or multichannel topological wires variants [152].

The core of this alternative strategy is a sudden jump followed by a period of constant velocity motion and then another jump. A theoretical analysis of this dynamics shows that the strategy simultaneously exploits the models stability at constant velocities below a certain threshold, and the property that small instantaneous changes in the domain-wall position do not dramatically reduce the ground-state fidelity. Further analysis of dressed JMJ protocols reveals that more complex initial jump sequences can, in addition to allowing a subcritical velocity, better target the ground state of the moving frame.

The protocol resembles the parallel recent development of the Bang-Anneal-Bang protocols [96] for ground-state preparation in noisy intermediate-scale quantum devices. Our theoretical analysis illuminates why this type of strategy works so successfully and we are hopeful that this approach will motivate more general theoretical analysis of related protocols. In this context it would also be interesting to see the same types of protocols also emerge in

more realistic materials. Our results for the semiconductor model suggest that similar protocols are indeed relevant. However, in that example, it is important to note that boundary motion does not strongly couple between bands with different spin orientations. In real materials one might expect that sudden boundary change would couple the ground state to bands at larger energies and this may restrict the range of motions that one can achieve.

Another natural direction to explore is the application of our methods to the problem of Majorana braiding, which is the ultimate goal behind the study and control of zero modes in the context of topological quantum computing. It remains an interesting open question whether aspects of our protocols remain robust if one models specific devices more closely. We emphasize, however, that the strength of the RL setup is that one can easily encode other experimental restrictions on the control by including them in the reward or action spaces [153]. As such the application of these methods can be readily extended to other more complicated models, such as models including the effect of interactions [135], different simulation techniques such as the time-dependent density-matrix renormalization group [154], and to measurement-based Majorana devices [106].

In a broader sense we believe our work shows that machine learning for quantum control can be done at a scale, speed, and precision that is relevant to modern experimental devices. As such we believe that it will be used to overcome obstacles to realizing large-scale quantum computation, quantum-communication networks, quantum thermal machines, analog quantum simulations, and control of other quantum many-body dynamical systems more broadly.

ACKNOWLEDGMENTS

We thank Andreas Buchleitner, Steve Campbell, Aaron Conlon, Shane Dooley, Ian Jubb, Kevin Kavanagh, Jinguo Liu, Xiuzhe Luo, Frank Schäfer, Jiayu Shen, and Lei Wang for helpful discussions regarding this project. L.C. and G.K. acknowledge Science Foundation Ireland for financial support through Career Development Award 15/CDA/3240. G.K. is also supported by a Schrödinger Fellowship. J.C. acknowledges support from Natural Sciences and Engineering Research Council of Canada (NSERC), the Shared Hierarchical Academic Research Computing Network (SHARCNET), Compute Canada, Google Quantum Research Award, and the Canadian Institute for Advanced Research (CIFAR) AI chair program. D.L and BKC acknowledge support from the Department of Energy Grant DOE desc0020165.

APPENDIX A: DERIVATION OF THE CRITICAL VELOCITY AND MOVING-FRAME PICTURE

In this appendix we first introduce some background theory related to the critical velocity and the

Majorana-mode wave function in the moving frame. Afterwards we derive the resonance time scale T_{res} for Majorana shuttling from the Fermi golden rule and finally we show the emergence of this resonance time scale in numerical simulations for forward oscillating motion.

1. Dispersion and Majorana mode in the moving frame

To find the critical velocity we start from the periodic ($N+1=1$) Kitaev-chain Hamiltonian [Eq. (1) with $V(x, t) = 0$] in the continuum, which reads after a Fourier transformation

$$H_k = \vec{d}_k \cdot \vec{\sigma} = \begin{pmatrix} \frac{k^2}{2m} - \mu & i\Delta k \\ -i\Delta k & -\frac{k^2}{2m} + \mu \end{pmatrix}, \quad (\text{A1})$$

with $\vec{d}_k = [0, -\Delta k, (k^2/2m) - \mu]$, momentum k , and $\vec{\sigma}$ the Pauli spin matrices. Solving for the energy eigenvalues gives mode dispersion $\epsilon_k = \pm|\vec{d}_k| = \pm\sqrt{(k^2/2m - \mu)^2 + \Delta^2 k^2}$. The moving frame can now be obtained by a Galilean transformation in terms of the unitary rotation $U(t) = e^{-ik \int_0^t v(t') dt'}$, which is the time-dependent translation operator that rotates to a frame with domain-wall velocity $v(t)$. Applying this to Eq. (A1) gives the moving-frame Hamiltonian

$$H_v(t) = U^\dagger(t) H_k U(t) + i \frac{dU^\dagger}{dt} U(t) = \begin{pmatrix} \frac{k^2}{2m} - \mu + v(t)k & i\Delta k \\ -i\Delta k & -\frac{k^2}{2m} + \mu + v(t)k \end{pmatrix}. \quad (\text{A2})$$

The moving frame has an extra term $v(t)k$ on the diagonal, which means the moving-frame energy dispersion is $\epsilon_k = \pm|\vec{d}_k| = \pm\sqrt{(k^2/2m - \mu)^2 + \Delta^2 k^2} + vk$ [73]. This results in a tilted dispersion as shown in Fig. 9, when $v = \Delta$ the tilt causes the gap to close and we have a topological phase transition. This velocity defines the critical velocity $v_{\text{crit}} = \Delta$.

Besides the energy dispersion, we can also look at the wave function of the Majorana zero modes in the moving frame. To find this wave function we solve for the zero-energy solutions of the moving-frame Hamiltonian (A2), which results in $\psi = [\phi, -\phi]^T$ with

$$\phi(x) \propto e^{-x/(\gamma\xi)} \sin[\sqrt{k_F^2 + 1/(\gamma\xi)^2} x]. \quad (\text{A3})$$

The difference with the Majoranas in the lab frame is that the localization length $\xi = 1/\Delta m$ is dilated by a factor $\gamma = 1/\sqrt{1 - v^2/\Delta^2}$ causing the Majorana modes to become spatially more extended (delocalized) in the moving frame. When $v = v_{\text{crit}} = \Delta$ the localization length becomes infinite indicating that the local character is lost and hence a topological phase transition occurs as shown in Fig. 9.

2. Derivation and numerical analysis of the resonant time scale

To derive the time scale for resonant Majorana motion we consider a scenario in which we are shuttling the left Majorana back and forth by using

$$v_L(t) = v_{\text{max}} \sin \omega t \quad (\text{A4})$$

for the velocity of the left domain wall in Eq. (2). When v_{max} is not too large, i.e., the amplitude of the left domain-wall position can be considered small compared to the

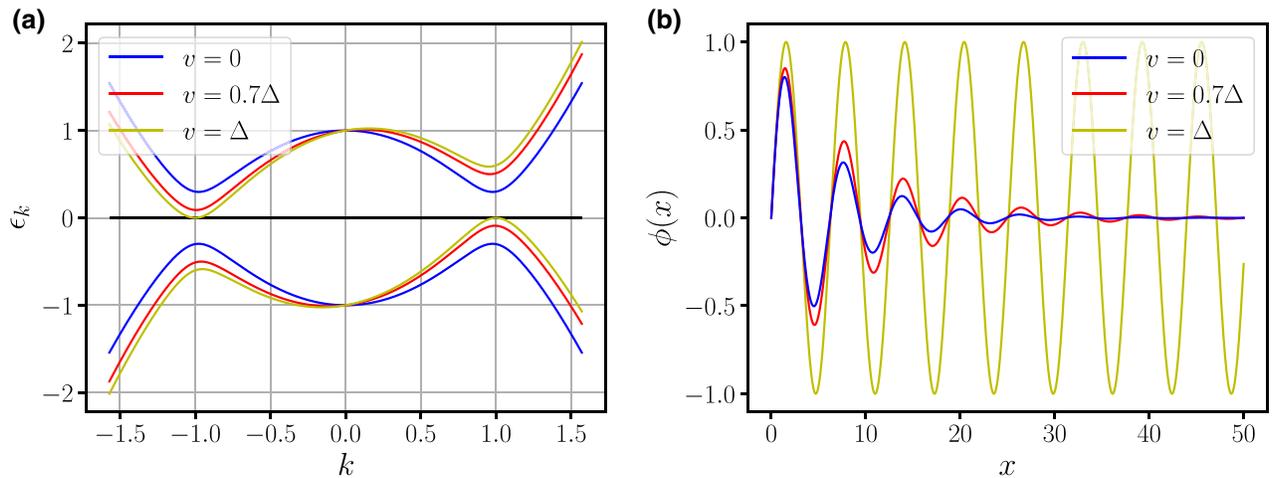


FIG. 9. (a) Mode dispersion in the moving frame for several different velocities. (b) Localization of the Majorana zero modes in the moving frame for several different velocities. It can be seen that the gap closes for $v = v_{\text{crit}} = \Delta$ and simultaneously the Majoranas delocalize in the moving frame.

total length of the wire N , we can treat this motion as a time-dependent perturbation and apply perturbation theory. In this case we write the external potential as $V = V_0 + \delta V \sin \omega t$ in which V_0 is the static domain-wall profile in Eq. (2) and $\delta V \sin \omega t$ a time-dependent fluctuation on top of it. From Fermi's golden rule for a harmonic perturbation [155] we can find the infidelity rate to be

$$\begin{aligned} \lim_{T \rightarrow \infty} \frac{\mathcal{I}(T)}{T} &= \lim_{T \rightarrow \infty} \frac{1}{T} [1 - |\langle \psi_0 | \exp -i \int_0^T \mathcal{H}(t) dt | \psi_0 \rangle|^2] \\ &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{i \neq 0} |\langle \psi_i | \exp -i \int_0^T \mathcal{H}(t) dt | \psi_0 \rangle|^2 \end{aligned} \quad (\text{A5})$$

$$\begin{aligned} &= 2\pi \sum_{i \neq 0} [\delta(E_i - E_0 + \omega) + \delta(E_i - E_0 - \omega)] \\ &\quad \times |\langle \psi_i | \delta V | \psi_0 \rangle|^2 \end{aligned} \quad (\text{A6})$$

in which $|\psi_i\rangle$ are the eigenstates of the Hamiltonian with the static domain-wall profile V_0 , which have energy E_i .

From this equation we can read off that there are resonances in the infidelity whenever $\omega_{\text{res}} = \pm(E_i - E_0)$, for the first excited state $\omega_{\text{res}} = E_1 - E_0 = \Delta k_F$ is equal to the gap in the system and we arrive at the resonance time $T_{\text{res}} = 2\pi/\omega_{\text{res}} = 2\pi/\Delta k_F$. The resonances for the higher energy bands are suppressed with the transition amplitude $|\langle \psi_i | \delta V | \psi_0 \rangle|^2$. When $\omega \mapsto \infty$ there are no resonances and the infidelity becomes zero; the intuition for this is that the static system H_0 does not notice the harmonic perturbation because it oscillates at a much higher frequency. Similarly, when $\omega < \omega_{\text{res}}$ there are also no resonances

because the frequencies are smaller than the gap which corresponds to adiabatic Majorana motion.

In the Majorana motion problem in the main text we are, however, looking at moving the left domain wall forward from x_A to x_B . As shown in Fig. 10 we see that the same resonance time scale emerges for forward motion with

$$v_L(t) = v_{\text{max}} \frac{1 - \cos \omega t}{2}. \quad (\text{A7})$$

For $\omega \mapsto \infty$ we get the same rate as constant motion with $\frac{v_{\text{max}}}{2}$, which can be explained for low velocities from adiabaticity.

APPENDIX B: ALGORITHM AND ANALYSIS OF THE DERIVATIVE OF THE INFIDELITY

In this appendix we first outline the algorithm we use to compute the many-body overlap, Eq. (3), for the Majorana wire setup. Afterwards we argue that all operations in this algorithm are differentiable and we can obtain the derivative of the fidelity with respect to control parameters using automatic differentiation. In the end we show the numerical obtained derivatives and check them with finite-difference methods.

To find the fidelity we first write the Hamiltonian, Eq. (1), in Bogolyubov-de-Gennes (single-particle) form [79,156]

$$\mathcal{H}(t) = \frac{1}{2} C^\dagger H_{\text{BdG}}(t) C, \quad (\text{B1})$$

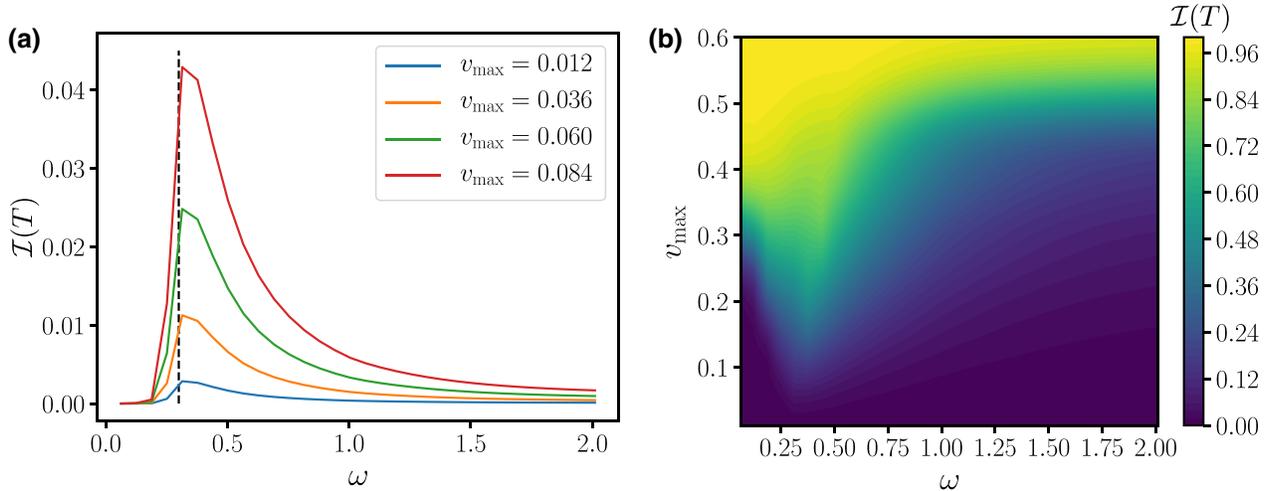


FIG. 10. (a) Infidelity as a function of frequency ω , the black dashed line indicates the resonance frequency $\omega_{\text{res}} = \Delta k_F = 0.3$. At $\omega \mapsto \infty$ the infidelity is the same as for constant linear motion with $x_L(t) = (v_{\text{max}}/2)t$. (b) Infidelity as a function of frequency and v_{max} . For velocities $v_{\text{max}} > v_{\text{crit}} = 0.3$ there are no low adiabatic frequencies anymore. The parameters used in these simulations are $N = 140$, $\mu = 1$, $w = 1$, $\Delta = 0.3$, and $\sigma = 1$.

with $C^\dagger \equiv [c_1^\dagger \cdots c_i^\dagger \cdots c_N^\dagger \quad c_1 \cdots c_i \cdots c_N]$ and diagonalize it

$$H_{\text{BdG}}(t)W(t) = E(t)W(t) \quad (\text{B2})$$

in which the columns of $W(t) = \begin{bmatrix} u(t) & v(t)^* \\ v(t) & u(t)^* \end{bmatrix}$ are the eigenmodes (quasiparticle modes) of the BdG Hamiltonian and $E(t)$ is the diagonal matrix of the mode energies $[E]_{ii}(t) = \epsilon_i(t)$. To compute the final many-body overlap at time $t = T$ we then use the Onishi formula

$$\mathcal{F}(T) = |\langle \psi_B | \psi(T) \rangle|^2 = \det[u_B^* u(T) + v_B^* v(T)], \quad (\text{B3})$$

which relates the BdG quasiparticle picture to the many-body picture. In this $u(T)$ and $v(T)$ can be obtained from $W(T) = \mathcal{T} e^{-i \int_0^T H_{\text{BdG}}(t') dt'} W(0)$ and u_B and v_B from the diagonalization of the instantaneous BdG Hamiltonian with the domain wall at position x_B . $\mathcal{F}(T)$ can now be used to compute the infidelity $\mathcal{I}(T)$ in Eq. (3) in the main text.

From a programming perspective, a code to compute the infidelity from Eqs. (B3) and (3) consists of a series of elementary operations (primitives) f_i that map the input control $x_L(t)$ with $t \in [0, T]$ to the output $\mathcal{I}(T) = f_n \circ f_{n-1} \circ \cdots \circ f_1[x_L(t)]$. The required primitives for this algorithm are matrix multiplications, diagonalization of a hermitian matrix Eq. B2 and taking the determinant Eq. (B3). Each of these operations has a known forward and reverse mode derivative [157] that can be recursively assembled in the chain rule

$$\frac{d\mathcal{I}}{dx_L} = \frac{df_n}{df_{n-1}} \frac{df_{n-1}}{df_{n-2}} \cdots \frac{df_1}{dx_L} \quad (\text{B4})$$

to find the derivative of the infidelity with respect to the control.

An automatic differentiation package, like the JAX library [124], computes all the derivatives of the primitives in a code automatically and assembles them in the chain rule to evaluate the total gradient of a computer program. We apply this method to the code to compute the infidelity and use the gradient for the optimizations described in the main text. In Fig. 11(a) we show an example for the gradient obtained in this way for a linear motion protocol in regime (I), which we check with finite difference to make sure it is working correctly. We observe that for this protocol the gradient is the biggest at the boundary, which might explain the jumping behavior at the beginning and end of the protocols we found with the ML optimizations (Fig. 2).

To compute the matrix exponential required for the computation of $W(T) = \mathcal{T} e^{-i \int_0^T H_{\text{BdG}}(t') dt'} W(0)$ we discretize the continuous time variable t in individual time steps of size dt and apply the Trotter-Suzuki expansion to write $\mathcal{T} e^{-i \int_0^T H_{\text{BdG}}(t') dt'} \approx \prod_{j=1}^{T/dt} e^{-i H_{\text{BdG}}(j dt) dt}$. This means that while the time complexity of the back-propagation algorithm is similar to the forward evaluation, the memory complexity grows linearly with the number of discrete time steps T/dt . A method to circumvent this growth in memory complexity is the use of adjoint sensitivity methods [158], which avoid back propagating through the time-evolution differential equation (Schrödinger equation) by solving a second differential equation backwards in time. In practice, the memory complexity did not turn out to be a bottleneck in our simulations, however, and it is sufficient to use the backpropagation algorithm as described above. For extending our methods to include disorder and interactions or for simulating very long times $T \sim \mathcal{O}(10^3)$ the adjoint method might be a good option to explore and we refer the interested reader to Ref. [159] for comparative studies between the adjoint and backpropagation algorithms.

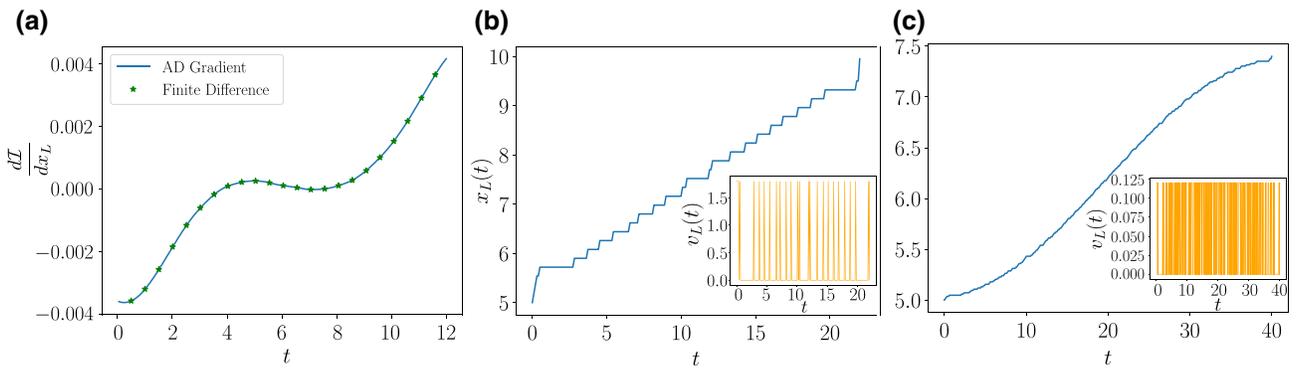


FIG. 11. (a) Comparison between the gradient of the infidelity \mathcal{I} with respect to the control x_L obtained with automatic differentiation (AD) and finite difference. The gradients $d\mathcal{I}/dx_L$ are evaluated for a linear protocol $x_L(t) = v_{\text{avg}} t$ as a function of time t in regime (I). The AD gradient matches up with the finite-difference gradient and we note that at the beginning and end of the protocol the magnitude of the gradient is the largest. (b),(c) Optimized parameterized bang-bang protocols in regimes (II) (b) and (IV) (c) with the simulated annealing method. The position profiles $x_L(t)$ are plotted in the main panels and the velocity profiles $v_L(t)$ in the insets. The bang-bang protocols seem to approximate the JMJ protocol in regime (II) and superadiabatic protocol in regime (IV). The parameters for these simulations are the same as in Fig. 2 in the main text.

APPENDIX C: GRADIENT FORMULA FOR NATURAL EVOLUTION STRATEGIES

Denote $p_\theta(\phi)$ for the Gaussian distribution $\phi \sim \mathcal{N}(\theta, \sigma^2 I)$ with σ fixed. It follows that [127]

$$\begin{aligned} \nabla_\theta \mathbb{E}_{\phi \sim \mathcal{N}(\theta, \sigma^2 I)} [\mathcal{I}(T, \phi)] \\ = \nabla_\theta \int p_\theta(\phi) \mathcal{I}(T, \phi) d\phi \end{aligned} \quad (C1)$$

$$= \int p_\theta(\phi) \nabla_\theta \log p_\theta(\phi) \mathcal{I}(T, \phi) d\phi \quad (C2)$$

$$= \int p_\theta(\phi) \frac{(\phi - \theta)}{\sigma^2} \mathcal{I}(T, \phi) d\phi. \quad (C3)$$

With change of variable $\phi = \theta + \sigma \epsilon$, we have

$$\begin{aligned} \nabla_\theta \mathbb{E}_{\phi \sim \mathcal{N}(\theta, \sigma^2 I)} [\mathcal{I}(T, \phi)] &= \int p_\theta(\epsilon) \mathcal{I}(T, \theta + \sigma \epsilon) \epsilon / \sigma d\epsilon \\ &= \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} [\mathcal{I}(T, \theta + \sigma \epsilon) \epsilon / \sigma]. \end{aligned} \quad (C4)$$

$$(C5)$$

APPENDIX D: COMPARING STANDARD BENCHMARK PROTOCOLS AND DIFFERENT PARAMETERIZATIONS OF THE CONTROL

In addition to the SA benchmark results in the main paper, in this appendix we compare the ML results to optimized parameterized standard bang-bang [74,95] and superadiabatic ramp-up-down protocols [75]. We also show results for ML optimizations of different parameterizations of the control $x_L(t)$.

For a standard bang-bang protocol the velocity of the domain wall v_L is at all times only allowed to take either the value $v_L = v_{\min}$ or the value $v_L = v_{\max}$, a switch between these two discrete velocities is called a bang. To search for the optimal bang-bang protocols we fix the minimum time between two consecutive bangs to be $\Delta t = 0.1$ and perform a simulated annealing search similar to the one described in the main text. In this SA search we impose an additional constraint (compared to the free search in the main text) in which each update steps consist of changing v_{\max} to v_{\min} and vice versa to retain the bang-bang character of the protocols. We note that a version of this

method is used in Ref. [132] to search for optimal bang-bang protocols in a spin system. We use $v_{\min} = 0$ similar to the bang-bang protocols in Ref. [74] and we scan different values of the maximum velocity $0 < v_{\max} < 8v_{\text{avg}}$. This means that for some of the bang-bang protocols in regimes (I) and (II) $v_{\max} \geq v_{\text{crit}}$, which is not the case for the protocols in Ref. [74], however, we observe that increasing v_{\max} leads to lower infidelities even when making v_{\max} bigger than the critical velocity.

The ramp-up-down protocols are a family of superadiabatic protocols in which the velocity is slowly built up from zero to some maximum velocity v_{\max} and then ramped down again to zero as given by

$$v_L(t) = \begin{cases} v_{\max} \frac{1 - \cos \omega t}{2}, & 0 \leq t \leq \frac{\pi}{\omega} \\ v_{\max}, & \frac{\pi}{\omega} \leq t \leq \frac{\pi}{\omega} + T \\ v_{\max} \frac{1 - \cos(\omega t - \omega T)}{2}, & \frac{\pi}{\omega} + T \leq t \leq \frac{2\pi}{\omega} + T \\ 0, & \text{otherwise} \end{cases} \quad (D1)$$

in which ω is the parameter that determines how quickly the velocity is accelerated to v_{\max} . To find the best ramp-up-down protocol in each regime we scan over a large range of values of the free parameter ω (v_{\max} is fixed by the average velocity constraint).

In Table I we summarize the best results (lowest infidelity) of these protocols. We compare these results to protocols obtained with the ML methods as described in the main text and see that in all regimes the ML protocols outperform the standard benchmark protocols significantly. The optimal bang-bang protocols obtained with the SA method in regimes (II) and (IV) are plotted in Fig. 11. It can be seen that the protocols start to approximate the optimal JMJ and superadiabatic protocols obtained with the other methods. To be able to fully approximate the optimal JMJ protocols with a bang-bang protocol one needs to test even higher maximum velocities than we did here and also consider a negative minimum velocity $v_{\min} < 0$.

We also test the NES and DP optimization algorithms for different parameterizations of the control of the domain wall on a smaller system size $N = 50$ before fixing the specific parameterizations as used and described in the main

TABLE I. Results for the infidelity $\mathcal{I}(T)$ of different benchmark protocols (first three columns) in regimes (I)–(IV) compared to results obtained with SA, AD, and NES. For these simulations (l, T) are $\{(4.32, 12), (4.95, 22), (0.48, 8), (2.4, 40)\}$, $N = 110$ and the other parameters of the Majorana wire setup are the same as for Fig. 2 in the main text.

Regime	Linear	Bang-bang	Ramp-up-down	SA	DP	NES
(I)	0.4738	0.3923	0.4817	0.3801	0.3546	0.3431
(II)	0.2236	0.1713	0.2350	0.1637	0.1549	0.1514
(III)	0.0120	0.0067	0.0125	0.0062	0.0056	0.0071
(IV)	0.0077	0.0004	0.0045	0.0004	0.0005	0.0009

TABLE II. Comparison of the lowest infidelity values between optimizations of AD and NES with respect to different parameterizations of the control [position $x_L(t)$, velocity $v_L(t)$ and neural net $x_L(t) = \text{NN}_\theta(t)$]. Based on these testing values we determine which parameterizations to use for the optimizations in the main text. These simulations are run for a system size of $N = 50$ and all the other parameters are the same as the parameters in Fig. 2 in the main text.

Regime	DP neural net	DP position	NES position	NES velocity	NES neural net
(I)	0.3910	0.4159	0.3945	0.4008	0.4586
(II)	0.1541	0.1531	0.1453	0.1509	0.1842
(III)	0.0064	0.0070	0.0120	0.0069	0.0084
(IV)	0.0001	0.0024	0.0148	0.0004	0.0007

text for the big system $N = 110$. We compare parameterizing the control by the position $x_L(t)$, the velocity $v_L(t)$ or a neural network $x_L(t) = \text{NN}_\theta(t)$ and show the results in Table II. The optimization with NES gives the lowest infidelity when we parameterize by position in regimes (I) and (II) and by velocity in regimes (III) and (IV). For the DP optimization the best infidelity optimizes a neural network in regimes (I), (III), and (IV) and optimizes over the position in regime (II).

APPENDIX E: DISCUSSION ON OTHER METHODS FOR QUANTUM CONTROL

We now briefly discuss the NES and DP methods in relation to traditional quantum-control algorithms and RL methods. First we note that DP provides a general framework to automatically evaluate gradients of computer programs, which can be integrated into traditional quantum-control techniques such as the Krotov method [160], Gradient-based pulse engineering [161], and other gradient-based control techniques. In this context, a key advantage of DP is its ability to ease the implementation of gradient-based techniques, where, e.g., the incorporation of physical constraints, such as the restrictions in the amplitude of the controls, requires only adding the constraints to the main objective function. This avoids complex and potentially error-prone analytical calculations. In addition, the implementation of control protocols to widely different physical scenarios (e.g., free fermion or boson Hamiltonians or alternative simulation strategies) typically requires only the coding of the forward propagation of the simulation with minimal changes to the optimization sub-routines. In addition, quantum optimal control algorithms based on DP can easily harness the acceleration afforded by graphics processing units (GPUs), which may speedup quantum-control calculations by more than an order of magnitude [120]. This is most easily achieved using freely available open-source packages such as JAX [124].

Compared to RL and other gradient-free methods, the DP method used here remains a powerful method as long as the physical model is differentiable, as is the case in our work since the BdG algorithm (Appendix B) used to compute the many-body infidelity is fully differentiable.

DP and all gradient-based methods provide a strong optimization signal and typically converge faster than RL techniques, which often require high sample complexity for good convergence. In addition, the use of gradients avoids the reward function design problem in RL [162,163]. This allows us to efficiently extend DP to the many-particle Hamiltonian required for the Majorana transport.

On the other hand, NES exhibits multiple advantageous features. First, as a gradient-free black-box optimization method, NES is the most flexible among the methods implemented in our work since it can be applied to controls that are discrete and continuous. Second, it can optimize nondifferentiable control problems, as well as problems where the gradient estimation is numerically unstable in a way that makes the application of DP challenging. NES is also easily parallelizable across multiple processors for a fast collection of sufficient samples for optimization. Lastly, compared to RL, we mention that NES can be directly applied to any desired objective function. This avoids the problem of the design of reward function commonly encountered in RL algorithms, which is necessary to avoid sparse learning signals that may hamper the efficiency of RL methods in quantum control.

We also briefly mention that RL techniques such as the Watkins Q -learning algorithm [70] and policy gradient, tabular Q learning, and deep Q learning have recently been shown to be capable of achieving high performance in other quantum-control setups [66]. Reference [66] provides a detailed comparison between tabular Q learning, deep Q learning, and policy gradient [70], and gradient descent and Krotov algorithms applied to the problem of preparing a target quantum state. Gleaned from their numerical experiments, the authors in Ref. [66] conclude that the deep Q learning and policy gradient algorithms tend to outperform other RL and traditional control algorithms when the problem allows for discrete values of the control parameters, and when the problem is scaled up to large system sizes.

Finally, we comment on well-established adiabatic shortcut methods [12], in particular, counter diabatic driving (CDD) and Lewis-Riesenfeld invariants (LRI), in the context of the Majorana game problem. We note that an attempt to implement these in our setup encounters an array of problems. In CDD for example, the control

protocols required a level of Hamiltonian control far beyond what would be experimentally feasible, where a full time-dependent and spatial control of a problem Hamiltonian is required. Lastly, LRI requires finding a suitable invariant, which to the best of our knowledge, is not known for the Majorana wire.

APPENDIX F: SEMICONDUCTING NANOWIRE SPIN-ORBIT COUPLED TO *S*-WAVE SUPERCONDUCTOR

The Kitaev-chain model for topological *p*-wave superconductivity [25] used for the simulations in the main text can be realized effectively in a semiconducting nanowire proximity coupled to a conventional *s*-wave superconductor in a Zeeman field [15–17]. In this appendix, we first introduce the proximity-coupled model and discuss in which limits the Kitaev-chain toy model can be obtained. Then we show how the critical velocity v_{crit} and resonance frequency can be obtained such that we can define the same Majorana motion regimes as in the main text. Finally we show numerical optimization results obtained with AD and NES in each of those regimes for two different sets of parameters of the proximity-coupled model (with one set outside the Kitaev-chain limit).

1. Hamiltonian and energy dispersion

The Hamiltonian for the proximity-coupled semiconducting nanowire (NW) for periodic boundary conditions, a homogeneous potential $V(x) = V$ and chemical potential $\mu(x) = \mu$ in the continuum momentum representation is

given by

$$\mathcal{H}_{\text{NW}} = \frac{1}{2} \int dk \Psi_k^\dagger H^{\text{BdG}}(k) \Psi_k, \quad (\text{F1})$$

with $\Psi_k^\dagger = [c_{k\uparrow}^\dagger c_{k\downarrow}^\dagger c_{-k\uparrow} c_{-k\downarrow}]$ and

$$H^{\text{BdG}}(k) = \left(\frac{k^2}{2m} + V - \mu + B\sigma_x + \alpha k\sigma_y \right) \tau_z + \Delta\sigma_y \tau_x \quad (\text{F2})$$

in which σ_i acts on the spin degrees of freedom and τ_i on the particle-hole space. In here we have an external magnetic field B in the x direction, spin-orbit coupling with strength α in the y direction and a standard *s*-wave superconducting term with gap Δ . Note the main difference with the Kitaev chain in Eq. (A1) is the spin degrees of freedom $s = \{\uparrow\downarrow\}$, which result in Zeeman splitting, spin-orbit interaction and *s*-wave superconductivity instead of spinless *p*-wave superconductivity.

The energy dispersion can be found by solving the characteristic equation for H^{BdG} Eq. (F2), which leads to

$$\begin{aligned} \epsilon_{\pm}^2(k) &= \xi^2 + B^2 + \alpha^2 k^2 + |\Delta|^2 \\ &\pm 2\sqrt{B^2|\Delta|^2 + \xi^2 B^2 + \alpha^2 k^2 \xi^2} \end{aligned} \quad (\text{F3})$$

in which we define $\xi(k) = k^2/2m + V - \mu$. In this four-band model various parameter choices lead to a gapped dispersion with in some phases Majorana zero modes, see for a complete discussion, e.g., Ref. [17]. Moreover, in two specific limits (large Zeeman field $B \gg \alpha, \Delta$ or for large spin-orbit coupling $\alpha \gg B, \Delta$) this proximity-coupled NW

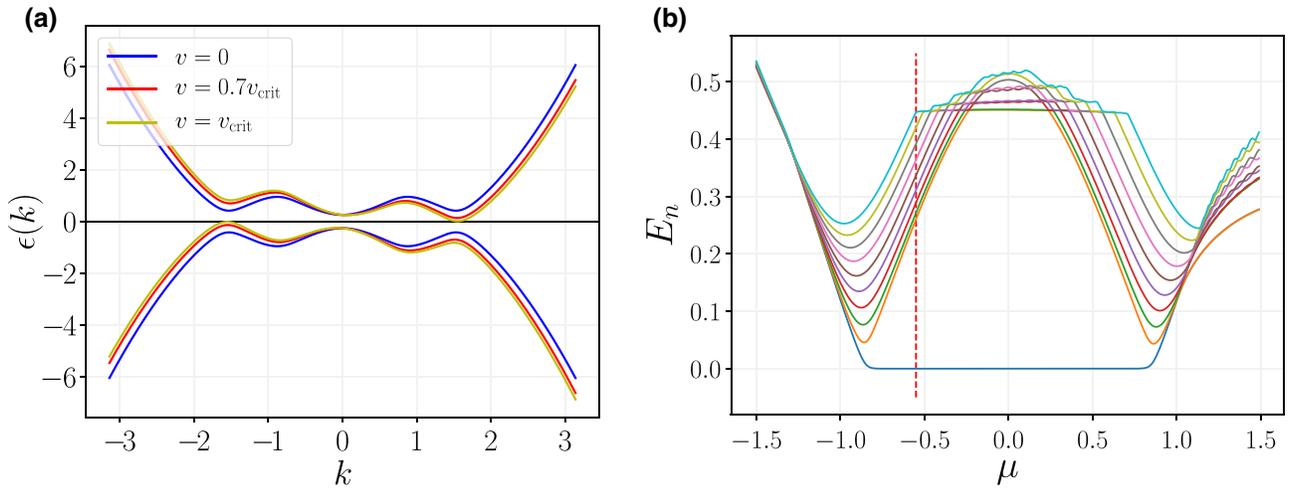


FIG. 12. (a) Energy dispersion of the Hamiltonian, Eq. (F4), in the moving frame for three different velocities. The velocity tilts the dispersion and the gap closes when $v = v_{\text{crit}}$. The model parameters used for this dispersion are $B = \alpha = 1.0$, $\Delta = 0.5$ and $\mu = -0.55$, $V = 0.0$. (b) Excitation energy eigenvalues E_n of the NW model on a lattice for open-boundary conditions as a function of the chemical potential μ . The red dashed line indicates $\mu = -0.55$. The other parameters are the same as for (a). We get a zero mode (blue line) but also some Andreev bound states (e.g., orange line) that come down when we bring μ closer to the bottom of the band.

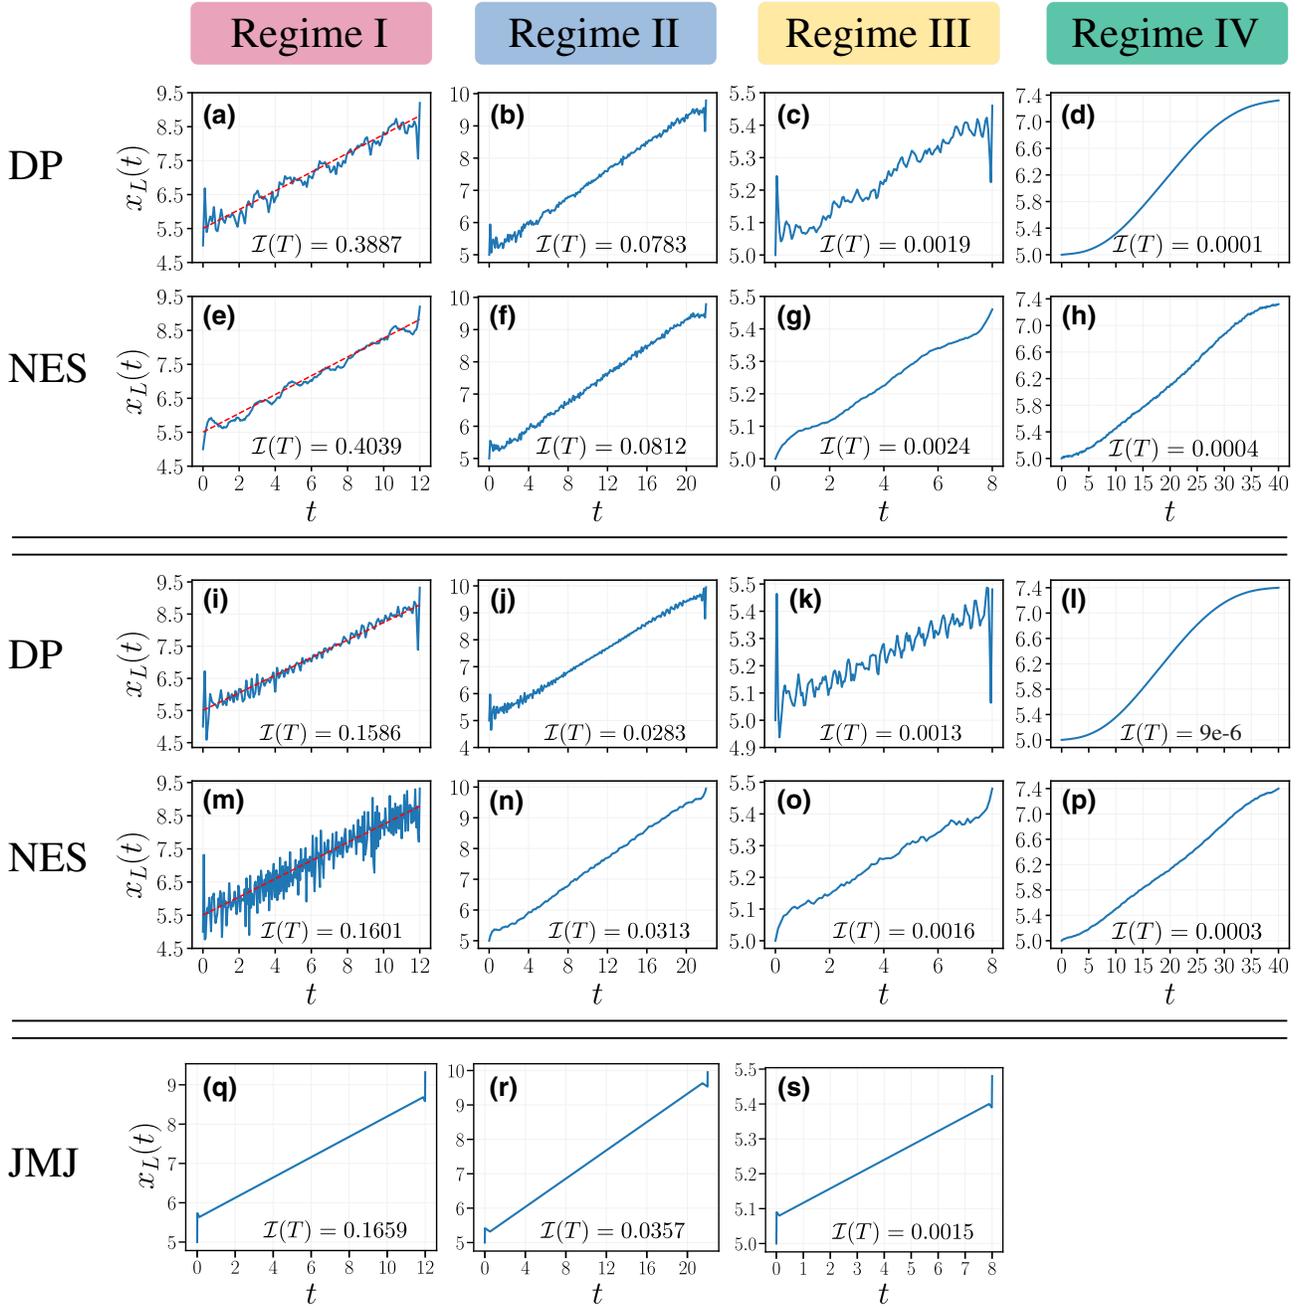


FIG. 13. Overview of the optimal controls in the four different Majorana motion regimes (I–IV) obtained after optimization with DP, NES, and the parameterized JMJ protocols for the proximity-coupled NW model as given by Eq. (F2). In all panels we plot the domain-wall position $x_L(t)$ as a function of time starting from $x_L(0) = x_A = 5.0$. The red dotted line in regime (I) has a slope equal to the critical velocity $v_{\text{crit}} \approx 0.27$ for (a) and (e) and $v_{\text{crit}} \approx 0.28$ for (i) and (m). The obtained strategies are similar to the strategies obtained for the Kitaev chain in the main text with pulselike motion at the beginning and end in regimes (I)–(II)–(III) being smoother motion in the middle and in regime (IV) a smooth adiabatic protocol. In the simulations for (a)–(h) we set $N = 100$, $\mu = 0.0$, $w = 1$, $\Delta = 0.8$, $B = 2$, $\alpha = 0.8$, $V_{\text{height}} = 30.1$, and $\sigma = 1$ and for (i)–(s) we set $N = 100$, $\mu = -0.55$, $w = 1$, $\Delta = 0.5$, $B = 1$, $\alpha = 1.0$, $V_{\text{height}} = 30.1$, and $\sigma = 1$. The constraint parameters (l, T) are given by $\{(4.32, 12), (4.95, 22), (0.48, 8), (2.4, 40)\}$ for regimes (I), (II), (III), and (IV), respectively.

model can be projected onto the lower bands (reduces to two similar blocks) resulting in the Kitaev-chain model with p -wave superconducting gap $\Delta' = \alpha\Delta/B$ (large B) or $\Delta' = \Delta/m\alpha$ (large α), respectively [73,152].

2. Critical velocity and resonance frequency

To bring the NW in Eq. (F2) into the moving frame we apply the same techniques as for the Kitaev chain discussed earlier in Appendix A, i.e., we apply the

time-dependent translation operator $U(t) = e^{-ik \int_0^t v(t') dt'}$ to rotate the Hamiltonian Eq. (F2), which results in the moving-frame Hamiltonian

$$H_v^{\text{BdG}}(k) = \left(\frac{k^2}{2m} + V - \mu + B\sigma_x + \alpha k\sigma_y \right) \tau_z + \Delta\sigma_y\tau_x + vk\mathbb{1}_4. \quad (\text{F4})$$

Like for the Kitaev chain this transformation gives an additional contribution vk to the energy dispersion that tilts it as shown in Fig. 12. In the regimes when the system is originally gapped the gap closes due to this tilt when $v \equiv v_{\text{crit}} = E_{\text{gap}}/k_F$. For most parameters finding the general expression for the energy gap and critical velocity is too complicated and we need to obtain it numerically. However, things simplify in the Kitaev-chain limit for $B \gg \alpha, \Delta$, which gives $v_{\text{crit}} \approx \alpha\Delta/B$ and when $\alpha \gg B, \Delta$ we obtain $v_{\text{crit}} \approx \Delta/m\alpha$ [73].

The resonance frequency in this model is like for the Kitaev chain defined as the being equal to the gap $\omega_{\text{res}} = E_{\text{gap}}$. With these associations for the critical velocity and resonance frequency the definition of the Majorana motion regimes originally in terms of parameters of the Kitaev chain can be extended to include the proximity-coupled model as well; regime (I) is defined to be for $v_{\text{avg}} > v_{\text{crit}}$, regime (II) close but below the critical velocity, regime (III) low velocity and short time scales with respect to the energy gap and finally the (super)adiabatic regime (IV) for long times and low velocities with respect to the gap.

3. Optimization results

By putting the NW model, Eq. (F2), on a lattice, opening up the boundary condition and imposing the smooth potential profile in Eq. (2) we acquire the same control of the position of the Majorana zero mode as for the Kitaev chain in the main text. In Fig. 12(b) we show an example of the excitation energies for a set of parameters outside the Kitaev-chain limit in which we have a Majorana zero mode and also Andreev bound states.

We then optimize the transport of the Majoranas in this NW model with the same AD and NES techniques as used before. We show the results for these simulations for the same parameters as in Fig. 12 from differentiable programming and natural evolutionary strategies in Figs. 13(i)–13(p) and from JMJ protocols in Figs. 13(q)–13(s). For completeness in Figs. 13(a)–13(h) we include results obtained by choosing parameters that approximate the Kitaev-chain limit (high B field).

It can be seen that for these two sets of parameters we obtain qualitatively similar strategies for Majorana transport in the proximity-coupled NW as for the simpler Kitaev-chain model as shown in Fig. 2 in the main text. That is, in regions (I)–(II)–(III) we get the JMJ strategy with a pulselike motion at the beginning and end of the

protocol and an on average constant motion with a velocity v_{crit} in the middle of the protocol. For the parameters outside the Kitaev limit, the initial and final dressed jumps are comprised of three parts: a forward jump, a backward jump, and another forward jump. For these parameters we observe a much smoother protocol in regime (II) compared to the Kitaev chain. In regime (IV) for both sets of parameters we again recover the expected smooth adiabatic protocols.

To show that the JMJ strategy (together with its analysis discussed in the main text) is also a robust strategy for the proximity-coupled NW model we scan again over the different parameters $\Delta x_{\text{forward}}$ and Δx_{back} outside the Kitaev limit. The optimal JMJ protocols obtained in this way together with the infidelity values are shown in Figs. 13(q)–13(s). It can be seen that the infidelity values for the dressed JMJ protocol are competitive to the protocols obtained with DP and NES. Compared to the Kitaev-chain results [(d) in Fig. 1] it can be seen that for this proximity-coupled NW model the dressing of the jumps is slightly less pronounced.

APPENDIX G: PARAMETERS AND INTERPOLATION OF THE OPTIMIZED JMJ MODEL STRATEGY

In this appendix we give the parameters of the optimized JMJ strategies in regimes (I), (II), and (III) as shown in Fig. 1(d) in the main text and verify how close these strategies are to the protocols obtained with the ML methods. In the main text we observe that dressed JMJ protocols with the lowest infidelities have a linear relation between $\Delta x_{\text{forward}}$ and $\Delta x_{\text{backward}}$ as Fig. 3 shows. To attain the best JMJ protocols in regimes (I)–(III), we scan through $\Delta x_{\text{forward}}$ and Δx_{back} with respect to the linear relation in each regime and pick the one that gives the overall lowest infidelity. The parameters of the resulting JMJ strategies are provided in Table III and the strategies are shown in Fig. 1.

We linearly interpolate these best JMJ model strategies to the machine-learning protocols in the main paper (Fig. 2) to see whether there are additional protocols that are better in terms of infidelity. The results for this in regime (III) with the DP strategy are shown in Fig. 14 from which we can see that the DP strategy is slightly better than

TABLE III. Best dressed JMJ model strategy parameters and their corresponding infidelities obtained from scans over $\Delta x_{\text{forward}}$, Δx_{back} , and Δt_{back} . The corresponding profiles are shown in Fig. 1 in the main paper.

Regime	$\Delta x_{\text{forward}}$	Δx_{back}	Δt_{back}	Infidelity
(I)	7.992	7.5060	0.05	0.3575
(II)	1.9093	1.5496	0.31	0.1589
(III)	0.5482	0.4571	0.33	0.0057

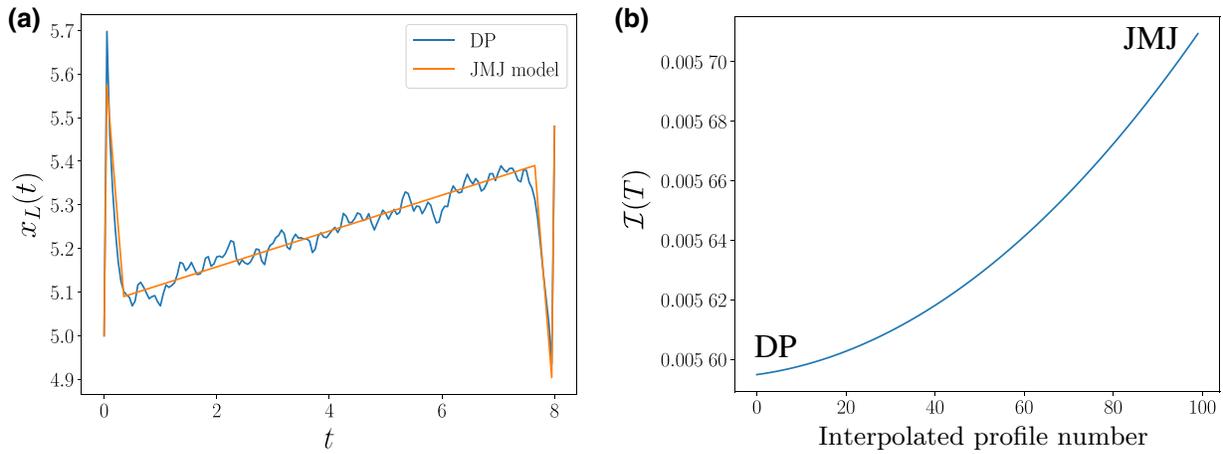


FIG. 14. (a) The DP optimized majorana motion strategy in regime (III) ($T = 8, l = 0.48$) on top of the optimized (scans over the free parameters) JMJ model strategy. The optimal JMJ model strategy captures the main features (jumps) of the DP strategy well. (b) Infidelity $\mathcal{I}(T)$ obtained for 100 linearly interpolated profiles between the DP strategy and JMJ model strategy as shown in (a). We observe a convex decreasing function towards the more flexible DP result.

the JMJ (due to more free parameters) and that there are no other local minima in between them. For the other regimes and interpolations from the NES strategy we obtain similar results.

APPENDIX H: EXTRA RESULTS ON ROBUSTNESS OF THE JMJ STRATEGY WITH RESPECT TO INTERACTIONS AND DISORDER

In Fig. 15 we show simulation results when running the optimal protocols obtained in the clean system in a

system with disorder and separately interactions. In summary we see that in regimes (I), (II), and (III) the ML protocols in the clean noninteracting system outperform a naive linear benchmark protocol. On the other hand, in regime (IV) after large enough interaction strengths or disorder strengths the linear protocols outperform the superadiabatic protocols found before in the literature. Moreover, the infidelity values get worse by more than an order of magnitude in this regime, which emphasizes the importance of looking for alternative strategies

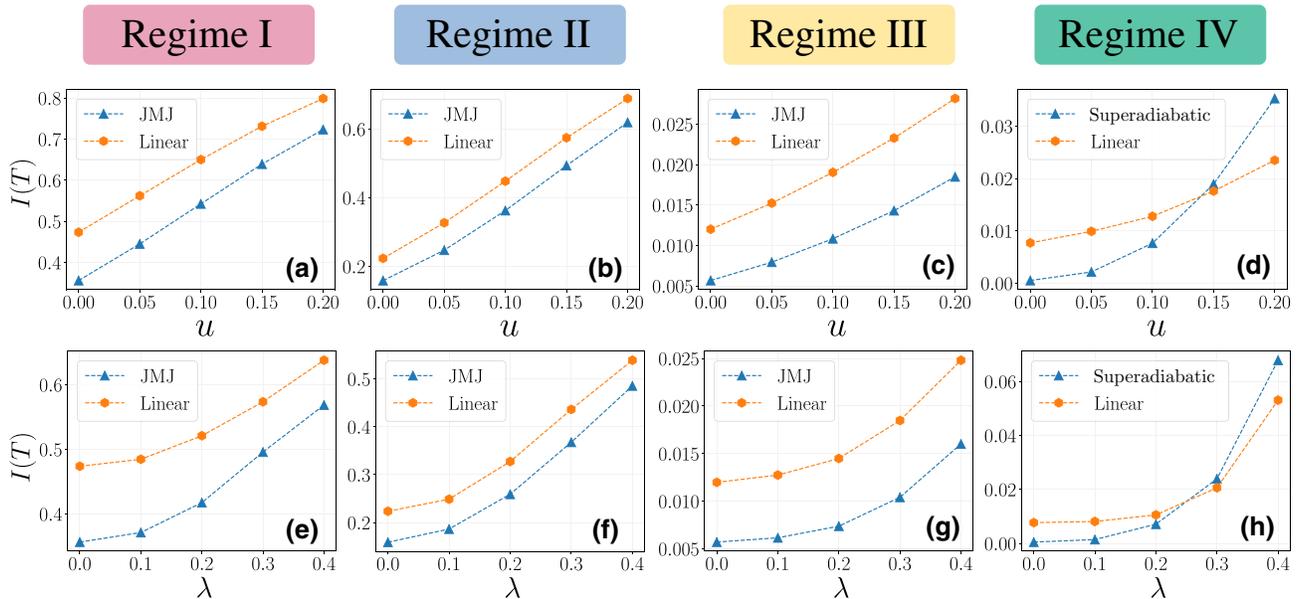


FIG. 15. Overview of simulation results of running the optimal protocols obtained in the clean noninteracting system in a system with interactions (top row) or disorder (bottom row) in all four regimes. In all regimes the infidelity increases with disorder strength λ and also with interaction strength u . The parameters for these simulations are the same as in Fig. 2 in the main text and the disorder averaging is done over 500 disorder realizations.

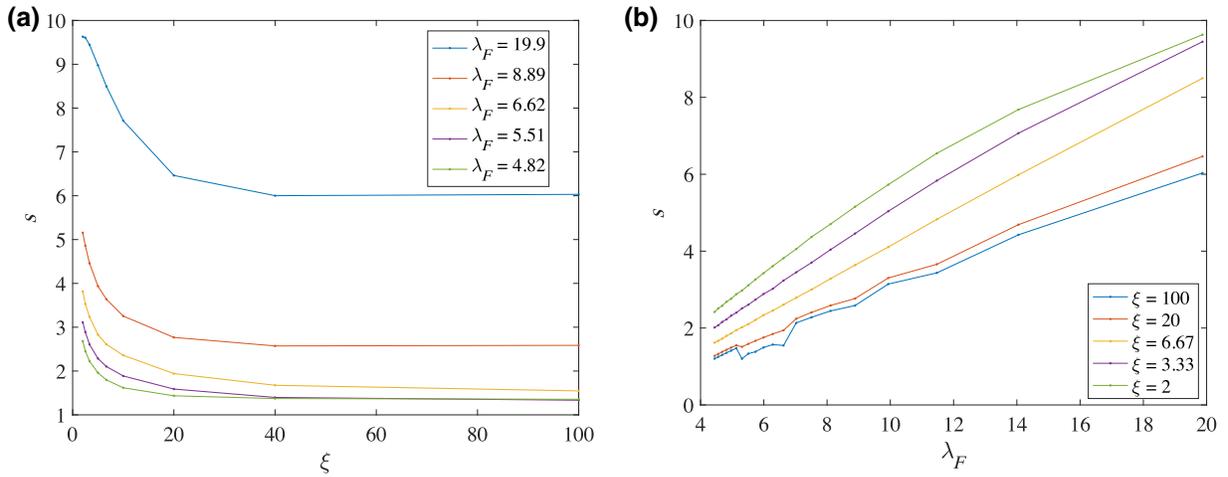


FIG. 16. In the p -wave model the $O_\delta \sim \exp(-\delta^2/s^2)$ have an interesting relationship to the two key length scales. The fit parameter s is reduced by a factor of about 1/2 when the coherence length $\xi = 1/(m\Delta)$ is made very large. On the other hand, its value tends to increase linearly in proportion to $\lambda_F = 2\pi/k_F$. The topological gap $E_{\text{gap}} \sim \Delta k_F = 2\pi/(m\xi\lambda_F)$ is, however, inversely proportional to both length scales.

that can be more efficient than the smooth superadiabatic protocols.

The observation that the infidelity values generally get worse with increasing interaction strength can be explained, at least partially, due to a lowering of effective critical velocity. This comes about in two main ways. Firstly on a mean-field level we know that including repulsive on-site interactions results in a lower topological energy gap E_{gap} (see, e.g., [135]). This would result in, as shown in Ref. [164], the resonance frequency ω_{res} becoming lower, which means that the resonance time scale, and consequently regime (IV), get pushed to longer time scales T .

Another feature that may be at play here is the fact that interactions, on a mean-field level, typically induce nonuniformity in the effective couplings (found, for example, by running some suitable Hartree-Fock-Bogoliubov minimization [156]). The result in this case would be similar to the scenario encountered in the disordered situations, where moving the wall over changing landscape amounts to a time-dependent perturbation (see, e.g., Ref. [92]). For future studies it would be interesting to investigate if one can use NES or DP to improve these values further by optimizing in the presence of disorder or interactions.

APPENDIX I: JUMP INFIDELITY COSTS WITH SYSTEM PARAMETERS

In this appendix we examine the infidelity cost of a single jump as a function of the important physical length scales $\lambda_F = 2\pi/k_F$ and $\xi = 1/(m\Delta)$. Our numerical test suggests a near Gaussian drop off in the infidelity as a function of jump distance: $O_\delta \sim \exp(-\delta^2/s^2)$. Generally speaking, the larger the value of s , the larger one can jump.

In Fig. 16 we plot the value of s that we obtain from a numerical fitting the O_δ drop off. We see that the value of the superconducting parameter Δ (and hence ξ) do play a role: the value of s tends to go down, but not dramatically so, as Δ gets smaller (and ξ gets larger). Perhaps more importantly we see that the s parameter is linearly related to λ_F . This means the jump size can, in principle, get bigger if one tunes the chemical potential towards the bottom of the single particle valence band.

- [1] J. M. Gambetta, J. M. Chow, and M. Steffen, Building logical qubits in a superconducting quantum computing system, *npj Quantum Inf.* **3**, 2 (2017).
- [2] J. I. Cirac and P. Zoller, Quantum Computations with Cold Trapped Ions, *Phys. Rev. Lett.* **74**, 4091 (1995).
- [3] J. G. Bohnet, B. C. Sawyer, J. W. Britton, M. L. Wall, A. M. Rey, M. Foss-Feig, and J. J. Bollinger, Quantum spin dynamics and entanglement generation with hundreds of trapped ions, *Science* **352**, 1297 (2016).
- [4] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, Trapped-ion quantum computing: Progress and challenges, *Appl. Phys. Rev.* **6**, 021314 (2019).
- [5] C. Gross and I. Bloch, Quantum simulations with ultracold atoms in optical lattices, *Science* **357**, 995 (2017).
- [6] 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 (2017).
- [7] F. Casola, T. van der Sar, and A. Yacoby, Probing condensed matter physics with magnetometry based on nitrogen-vacancy centres in diamond, *Nat. Rev. Mater.* **3**, 17088 (2018).
- [8] D. P. DiVincenzo, The physical implementation of quantum computation, *Fortschritte der Physik* **48**, 771 (2000).

- [9] A. Baksic, H. Ribeiro, and A. A. Clerk, Speeding up Adiabatic Quantum State Transfer by Using Dressed States, *Phys. Rev. Lett.* **116**, 230503 (2016).
- [10] H. Rabitz, R. de Vivie-Riedle, M. Motzkus, and K. Kompa, Whither the future of controlling quantum phenomena?, *Science* **288**, 824 (2000).
- [11] X. Wang, M. Allegra, K. Jacobs, S. Lloyd, C. Lupo, and M. Mohseni, Quantum Brachistochrone Curves as Geodesics: Obtaining Accurate Minimum-Time Protocols for the Control of Quantum Systems, *Phys. Rev. Lett.* **114**, 170501 (2015).
- [12] 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).
- [13] A. del Campo and K. Sengupta, Controlling quantum critical dynamics of isolated systems, *Eur. Phys. J. Spec. Top.* **224**, 189 (2015).
- [14] A. G. J. MacFarlane, J. P. Dowling, and G. J. Milburn, Quantum technology: The second quantum revolution, *Philos. Trans. R. Soc. London. Ser. A: Math., Phys. Eng. Sci.* **361**, 1655 (2003).
- [15] L. Fu and C. L. Kane, Superconducting Proximity Effect and Majorana Fermions at the Surface of a Topological Insulator, *Phys. Rev. Lett.* **100**, 096407 (2008).
- [16] R. M. Lutchyn, J. D. Sau, and S. Das Sarma, Majorana Fermions and a Topological Phase Transition in Semiconductor-Superconductor Heterostructures, *Phys. Rev. Lett.* **105**, 077001 (2010).
- [17] Y. Oreg, G. Refael, and F. von Oppen, Helical Liquids and Majorana Bound States in Quantum Wires, *Phys. Rev. Lett.* **105**, 177002 (2010).
- [18] J. Alicea, Y. Oreg, G. Refael, F. von Oppen, and M. P. A. Fisher, Non-abelian statistics and topological quantum information processing in 1d wire networks, *Nat. Phys.* **7**, 412 (2011).
- [19] M. Kjaergaard, K. Wölms, and K. Flensberg, Majorana fermions in superconducting nanowires without spin-orbit coupling, *Phys. Rev. B* **85**, 020503 (2012).
- [20] M. Leijnse and K. Flensberg, Introduction to topological superconductivity and majorana fermions, *Semicond. Sci. Technol.* **27**, 124003 (2012).
- [21] V. Mourik, K. Zuo, S. M. Frolov, S. R. Plissard, E. P. A. M. Bakkers, and L. P. Kouwenhoven, Signatures of majorana fermions in hybrid superconductor-semiconductor nanowire devices, *Science* **336**, 1003 (2012).
- [22] M. T. Deng, C. L. Yu, G. Y. Huang, M. Larsson, P. Caroff, and H. Q. Xu, Anomalous zero-bias conductance peak in a Nb-InSb nanowire-Nb hybrid device, *Nano Lett.* **12**, 6414 (2012).
- [23] A. Das, Y. Ronen, Y. Most, Y. Oreg, M. Heiblum, and H. Shtrikman, Zero-bias peaks and splitting in an Al-InAs nanowire topological superconductor as a signature of majorana fermions, *Nat. Phys.* **8**, 887 (2012).
- [24] R. M. Lutchyn, E. P. A. M. Bakkers, L. P. Kouwenhoven, P. Krogstrup, C. M. Marcus, and Y. Oreg, Majorana zero modes in superconductor-semiconductor heterostructures, *Nat. Rev. Mater.* **3**, 52 (2018).
- [25] A. Y. Kitaev, Unpaired majorana fermions in quantum wires, *Phys.-Usp.* **44**, 131 (2001).
- [26] A. Kitaev, Fault-tolerant quantum computation by anyons, *Ann. Phys. (N. Y.)* **303**, 2 (2003).
- [27] S. D. Sarma, M. Freedman, and C. Nayak, Majorana zero modes and topological quantum computation, *npj Quantum Inf.* **1**, 1 (2015).
- [28] C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma, Non-abelian anyons and topological quantum computation, *Rev. Mod. Phys.* **80**, 1083 (2008).
- [29] T. D. Stanescu and S. Tewari, Majorana fermions in semiconductor nanowires: Fundamentals, modeling, and experiment, *J. Phys.: Condens. Matter* **25**, 233201 (2013).
- [30] D. Aasen, M. Hell, R. V. Mishmash, A. Higginbotham, J. Danon, M. Leijnse, T. S. Jespersen, J. A. Folk, C. M. Marcus, K. Flensberg, and J. Alicea, Milestones toward Majorana-Based Quantum Computing, *Phys. Rev. X* **6**, 031016 (2016).
- [31] D. Rainis and D. Loss, Majorana qubit decoherence by quasiparticle poisoning, *Phys. Rev. B* **85**, 174533 (2012).
- [32] A. Voulodimos, N. Doulamis, A. Doulamis, and E. Protopapadakis, Deep learning for computer vision: A brief review, *Comput. Intell. Neurosci.* **2018**, 7068349 (2018).
- [33] T. Young, D. Hazarika, S. Poria, and E. Cambria, Recent Trends in Deep Learning Based Natural Language Processing, [arXiv:1708.02709](https://arxiv.org/abs/1708.02709) [cs] (2018).
- [34] G. Eraslan, Ž. Avsec, J. Gagneur, and F. J. Theis, Deep learning: New computational modelling techniques for genomics, *Nat. Rev. Genet.* **20**, 389 (2019).
- [35] 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).
- [36] J. Carrasquilla, Machine learning for quantum matter, *Adv. Phys.: X* **5**, 1797528 (2020).
- [37] I. E. Lagaris, A. Likas, and D. I. Fotiadis, Artificial neural network methods in quantum mechanics, *Comput. Phys. Commun.* **104**, 1 (1997).
- [38] J. Carrasquilla and R. G. Melko, Machine learning phases of matter, *Nat. Phys.* **13**, 431 (2017).
- [39] G. Carleo and M. Troyer, Solving the quantum many-body problem with artificial neural networks, *Science* **355**, 602 (2017).
- [40] D. Kochkov and B. K. Clark, Variational optimization in the AI era: Computational graph states and supervised wave-function optimization, [arXiv:1811.12423](https://arxiv.org/abs/1811.12423) (2018).
- [41] D. Luo and B. K. Clark, Backflow Transformations via Neural Networks for Quantum Many-Body Wave Functions, *Phys. Rev. Lett.* **122**, 226401 (2019).
- [42] D. Pfau, J. S. Spencer, A. G. de G. Matthews, and W. M. C. Foulkes, Ab-initio solution of the many-electron Schrödinger equation with deep neural networks, [arXiv:1909.02487](https://arxiv.org/abs/1909.02487) [physics.chem-ph] (2019).
- [43] J. Hermann, Z. Shtzle, and F. No, Deep neural network solution of the electronic Schrödinger equation, [arXiv:1909.08423](https://arxiv.org/abs/1909.08423) [physics.comp-ph] (2019).
- [44] O. Sharir, Y. Levine, N. Wies, G. Carleo, and A. Shashua, Deep Autoregressive Models for the Efficient Variational Simulation of Many-Body Quantum Systems, *Phys. Rev. Lett.* **124**, 020503 (2020).
- [45] M. Hibat-Allah, M. Ganahl, L. E. Hayward, R. G. Melko, and J. Carrasquilla, Recurrent neural network wave functions, *Phys. Rev. Res.* **2**, 023358 (2020).

- [46] S. Lu, X. Gao, and L.-M. Duan, Efficient representation of topologically ordered states with restricted Boltzmann machines, *Phys. Rev. B* **99**, 155136 (2019).
- [47] X. Gao and L.-M. Duan, Efficient representation of quantum many-body states with deep neural networks, *Nat. Commun.* **8**, 662 (2017).
- [48] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, Neural-Network Quantum States, String-Bond States, and Chiral Topological States, *Phys. Rev. X* **8**, 011006 (2018).
- [49] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Neural-network quantum state tomography, *Nat. Phys.* **14**, 447 (2018).
- [50] G. Torlai and R. G. Melko, Latent Space Purification via Neural Density Operators, *Phys. Rev. Lett.* **120**, 240503 (2018).
- [51] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Reconstructing quantum states with generative models, *Nat. Machine Intelligence* **1**, 155 (2019).
- [52] M. Rupp, A. Tkatchenko, K.-R. Müller, and O. A. von Lilienfeld, Fast and Accurate Modeling of Molecular Atomization Energies with Machine Learning, *Phys. Rev. Lett.* **108**, 058301 (2012).
- [53] M. Rupp, Machine learning for quantum mechanics in a nutshell, *Int. J. Quantum. Chem.* **115**, 1058 (2015).
- [54] K. T. Schütt, M. Gastegger, A. Tkatchenko, K.-R. Müller, and R. J. Maurer, Unifying machine learning and quantum chemistry with a deep neural network for molecular wavefunctions, *Nat. Commun.* **10**, 5024 (2019).
- [55] J. C. Snyder, M. Rupp, K. Hansen, K.-R. Müller, and K. Burke, Finding Density Functionals with Machine Learning, *Phys. Rev. Lett.* **108**, 253002 (2012).
- [56] L. Li, J. C. Snyder, I. M. Pelaschier, J. Huang, U.-N. Niranjan, P. Duncan, M. Rupp, K.-R. Müller, and K. Burke, Understanding machine-learned density functionals, *Int. J. Quantum. Chem.* **116**, 819 (2016).
- [57] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, Self-learning monte carlo method, *Phys. Rev. B* **95**, 041101 (2017).
- [58] L. Huang and L. Wang, Accelerated Monte Carlo simulations with restricted Boltzmann machines, *Phys. Rev. B* **95**, 035105 (2017).
- [59] S. Pilati, E. M. Inack, and P. Pieri, Self-learning projective quantum Monte Carlo simulations guided by restricted Boltzmann machines, *Phys. Rev. E* **100**, 043301 (2019).
- [60] A. A. Melnikov, H. P. Nautrup, M. Krenn, V. Dunjko, M. Tiersch, A. Zeilinger, and H. J. Briegel, Active learning machine learns to create new quantum experiments, *Proc. Natl. Acad. Sci.* **115**, 1221 (2018).
- [61] J. Yao, L. Lin, and M. Bukov, Reinforcement Learning for Many-Body Ground State Preparation based on Counter-Diabatic Driving, [arXiv:2010.03655](https://arxiv.org/abs/2010.03655) [cond-mat, physics:physics, physics:quant-ph] (2020).
- [62] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, Reinforcement Learning with Neural Networks for Quantum Feedback, *Phys. Rev. X* **8**, 031084 (2018).
- [63] V. Dunjko and H. J. Briegel, Machine learning & artificial intelligence in the quantum domain: A review of recent progress, *Rep. Prog. Phys.* **81**, 074001 (2018).
- [64] 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).
- [65] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven, Universal quantum control through deep reinforcement learning, *npj Quantum Inf.* **5**, 1 (2019).
- [66] X.-M. Zhang, Z. Wei, R. Asad, X.-C. Yang, and X. Wang, When does reinforcement learning stand out in quantum control? A comparative study on state preparation, *npj Quantum Inf.* **5**, 1 (2019).
- [67] F. Schäfer, M. Kloc, C. Bruder, and N. Lörch, A differentiable programming method for quantum control, *Machine Learning: Sci. Technol.* **1**, 035009 (2020).
- [68] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, S. Petersen, C. Beattie, A. Sadik, I. Antonoglou, H. King, D. Kumaran, D. Wierstra, S. Legg, and D. Hassabis, Human-level control through deep reinforcement learning, *Nature* **518**, 529 (2015).
- [69] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, Mastering the game of go with deep neural networks and tree search, *Nature* **529**, 484 (2016).
- [70] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (A Bradford Book, Cambridge, MA, USA, 2018).
- [71] A. G. Baydin, B. A. Pearlmutter, A. A. Radul, and J. M. Siskind, Automatic differentiation in machine learning: A survey, *J. Mach. Learn. Res.* **18**, 5595 (2017).
- [72] I. Rechenberg, *Evolutionsstrategie: Optimierung Technischer Systeme Nach Prinzipien der Biologischen Evolution* (Frommann-Holzboog, Stuttgart, 1973).
- [73] M. S. Scheurer and A. Shnirman, Nonadiabatic processes in majorana qubit systems, *Phys. Rev. B* **88**, 064515 (2013).
- [74] T. Karzig, A. Rahmani, F. von Oppen, and G. Refael, Optimal control of majorana zero modes, *Phys. Rev. B* **91**, 201404 (2015).
- [75] A. Conlon, D. Pellegrino, J. K. Slingerland, S. Dooley, and G. Kells, Error generation and propagation in majorana-based topological qubits, *Phys. Rev. B* **100**, 134307 (2019).
- [76] P. W. Brouwer, M. Duckheim, A. Romito, and F. von Oppen, Probability Distribution of Majorana End-State Energies in Disordered Wires, *Phys. Rev. Lett.* **107**, 196804 (2011).
- [77] S. B. Chung, H.-J. Zhang, X.-L. Qi, and S.-C. Zhang, Topological superconducting phase and majorana fermions in half-metal/superconductor heterostructures, *Phys. Rev. B* **84**, 060510 (2011).
- [78] S. Nadj-Perge, I. K. Drozdov, B. A. Bernevig, and A. Yazdani, Proposal for realizing majorana fermions in chains of magnetic atoms on a superconductor, *Phys. Rev. B* **88**, 020407 (2013).
- [79] B. Bernevig and T. Hughes, *Topological Insulators and Topological Superconductors* (Princeton University Press, Princeton, 2013).

- [80] N. Read and D. Green, Paired states of fermions in two dimensions with breaking of parity and time-reversal symmetries and the fractional quantum Hall effect, *Phys. Rev. B* **61**, 10267 (2000).
- [81] D. A. Ivanov, Non-Abelian Statistics of Half-Quantum Vortices in p -Wave Superconductors, *Phys. Rev. Lett.* **86**, 268 (2001).
- [82] A. Stern, F. von Oppen, and E. Mariani, Geometric phases and quantum entanglement as building blocks for non-abelian quasiparticle statistics, *Phys. Rev. B* **70**, 205338 (2004).
- [83] M. Stone and S.-B. Chung, Fusion rules and vortices in $p_x + ip_y$ superconductors, *Phys. Rev. B* **73**, 014505 (2006).
- [84] J. C. Budich, S. Walter, and B. Trauzettel, Failure of protection of majorana based qubits against decoherence, *Phys. Rev. B* **85**, 121405 (2012).
- [85] F. Konschelle and F. Hassler, Effects of nonequilibrium noise on a quantum memory encoded in majorana zero modes, *Phys. Rev. B* **88**, 075431 (2013).
- [86] H. T. Ng, Decoherence of interacting majorana modes, *Sci. Rep.* **5**, 12530 (2015).
- [87] F. L. Pedrocchi and D. P. DiVincenzo, Majorana Braiding with Thermal Noise, *Phys. Rev. Lett.* **115**, 120402 (2015).
- [88] Y. Hu, Z. Cai, M. A. Baranov, and P. Zoller, Majorana fermions in noisy kitaev wires, *Phys. Rev. B* **92**, 165118 (2015).
- [89] C. Knapp, M. Zaletel, D. E. Liu, M. Cheng, P. Bonderson, and C. Nayak, The Nature and Correction of Diabatic Errors in Anyon Braiding, *Phys. Rev. X* **6**, 041003 (2016).
- [90] C. Knapp, T. Karzig, R. M. Lutchyn, and C. Nayak, Dephasing of majorana-based qubits, *Phys. Rev. B* **97**, 125404 (2018).
- [91] Y. Huang, A. M. Lobos, and Z. Cai, Dissipative majorana quantum wires, *iScience* **21**, 241 (2019).
- [92] T. Karzig, G. Refael, and F. von Oppen, Boosting Majorana Zero Modes, *Phys. Rev. X* **3**, 041017 (2013).
- [93] B. Bauer, T. Karzig, R. V. Mishmash, A. E. Antipov, and J. Alicea, Dynamics of majorana-based qubits operated with an array of tunable gates, *SciPost Phys.* **5**, 4 (2018).
- [94] For a more detailed discussion of these effects in more complicated geometries see Refs. [165,166]. For a discussion on alternative braiding schemes see Sec. II A.
- [95] Z.-C. Yang, A. Rahmani, A. Shabani, H. Neven, and C. Chamon, Optimizing Variational Quantum Algorithms Using Pontryagin's Minimum Principle, *Phys. Rev. X* **7**, 021027 (2017).
- [96] L. T. Brady, C. L. Baldwin, A. Bapat, Y. Kharkov, and A. V. Gorshkov, Optimal protocols in quantum annealing and QAOA problems, [arXiv:2003.08952](https://arxiv.org/abs/2003.08952) [quant-ph] (2020).
- [97] M. V. Berry, Quantum phase corrections from adiabatic iteration, *Proc. R. Soc. London A: Math. Phys. Sci.* **414**, 31 (1987).
- [98] M. Deschamps, G. Kervern, D. Massiot, G. Pintacuda, L. Emsley, and P. J. Grandinetti, Superadiabaticity in magnetic resonance, *J. Chem. Phys.* **129**, 204110 (2008).
- [99] J. D. Sau, D. J. Clarke, and S. Tewari, Controlling non-abelian statistics of majorana fermions in semiconductor nanowires, *Phys. Rev. B* **84**, 094505 (2011).
- [100] B. I. Halperin, Y. Oreg, A. Stern, G. Refael, J. Alicea, and F. von Oppen, Adiabatic manipulations of majorana fermions in a three-dimensional network of quantum wires, *Phys. Rev. B* **85**, 144501 (2012).
- [101] B. van Heck, A. R. Akhmerov, F. Hassler, M. Burrello, and C. W. J. Beenakker, Coulomb-assisted braiding of majorana fermions in a josephson junction array, *New J. Phys.* **14**, 035019 (2012).
- [102] M. Burrello, B. van Heck, and A. R. Akhmerov, Braiding of non-abelian anyons using pairwise interactions, *Phys. Rev. A* **87**, 022343 (2013).
- [103] P. Bonderson, M. Freedman, and C. Nayak, Measurement-Only Topological Quantum Computation, *Phys. Rev. Lett.* **101**, 010501 (2008).
- [104] P. Bonderson, Measurement-only topological quantum computation via tunable interactions, *Phys. Rev. B* **87**, 035113 (2013).
- [105] S. Vijay and L. Fu, Teleportation-based quantum information processing with majorana zero modes, *Phys. Rev. B* **94**, 235446 (2016).
- [106] T. Karzig, C. Knapp, R. M. Lutchyn, P. Bonderson, M. B. Hastings, C. Nayak, J. Alicea, K. Flensberg, S. Plugge, Y. Oreg, C. M. Marcus, and M. H. Freedman, Scalable designs for quasiparticle-poisoning-protected topological quantum computation with majorana zero modes, *Phys. Rev. B* **95**, 235305 (2017).
- [107] S. Plugge, A. Rasmussen, R. Egger, and K. Flensberg, Majorana box qubits, *New J. Phys.* **19**, 012001 (2017).
- [108] C. Knapp, J. I. Väyrynen, and R. M. Lutchyn, Number-conserving analysis of measurement-based braiding with majorana zero modes, *Phys. Rev. B* **101**, 125108 (2020).
- [109] C. Zeng, G. Sharma, T. D. Stanescu, and S. Tewari, Feasibility of measurement-based braiding in the quasi-majorana regime of semiconductor-superconductor heterostructures, *Phys. Rev. B* **102**, 205101 (2020).
- [110] E. Torrontegui, S. Ibáñez, S. Martínez-Garaot, M. Modugno, A. del Campo, D. Guéry-Odelin, A. Ruschhaupt, X. Chen, and J. G. Muga, in *Advances In Atomic, Molecular, and Optical Physics*, Advances in Atomic, Molecular, and Optical Physics Vol. 62, edited by E. Arimondo, P. R. Berman, and C. C. Lin (Academic Press, Cambridge Massachusetts, 2013), Chap. 2, p. 117.
- [111] T. Karzig, F. Pientka, G. Refael, and F. von Oppen, Shortcuts to non-abelian braiding, *Phys. Rev. B* **91**, 201102 (2015).
- [112] G. Kells, V. Lahtinen, and J. Vala, Kitaev spin models from topological nanowire networks, *Phys. Rev. B* **89**, 075122 (2014).
- [113] We note that the control can also be represented by the velocity or a neural-network parameterization of the domain wall from which the position $x_L(t)$ can be extracted.
- [114] L. Coopmans, D. Luo, G. Kells, B. K. Clark, and J. Carrasquilla, Protocol discovery for the quantum control of majoranas by differential programming and natural evolution strategies, [github repository](https://github.com/loopmans) (2020).
- [115] R. E. Wengert, A simple automatic derivative evaluation program, *Commun. ACM* **7**, 463 (1964).
- [116] H.-J. Liao, J.-G. Liu, L. Wang, and T. Xiang, Differentiable Programming Tensor Networks, *Phys. Rev. X* **9**, 031041 (2019).

- [117] H. Xie, J.-G. Liu, and L. Wang, Automatic differentiation of dominant eigensolver and its applications in quantum physics, *Phys. Rev. B* **101**, 245139 (2020).
- [118] S. Sorella and L. Capriotti, Algorithmic differentiation and the calculation of forces by quantum Monte Carlo, *J. Chem. Phys.* **133**, 234111 (2010).
- [119] S.-X. Zhang, Z.-Q. Wan, and H. Yao, Automatic differentiable Monte Carlo: Theory and application, [arXiv:1911.09117](https://arxiv.org/abs/1911.09117) [physics.comp-ph] (2019).
- [120] N. Leung, M. Abdelhafez, J. Koch, and D. Schuster, Speedup for quantum optimal control from automatic differentiation based on graphics processing units, *Phys. Rev. A* **95**, 042318 (2017).
- [121] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, Learning representations by back-propagating errors, *Nature* **323**, 533 (1986).
- [122] W. Baur and V. Strassen, The complexity of partial derivatives, *Theor. Comput. Sci.* **22**, 317 (1983).
- [123] J. Morgenstern, How to compute fast a function and all its derivatives: A variation on the theorem of Baur-Strassen, *SIGACT News* **16**, 60 (1985).
- [124] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, and S. Wanderman-Milne, JAX: Composable transformations of Python+NumPy programs (2018).
- [125] D. P. Kingma and J. Ba, in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7–9, 2015, Conference Track Proceedings*, edited by Y. Bengio and Y. LeCun (2015).
- [126] H.-P. Schwefel, *Numerische Optimierung von Computer-Modellen mittels der Evolutionsstrategie*, ISR Vol. 26 (Birkhaeuser, Basel/Stuttgart, 1977), p. 390.
- [127] D. Wierstra, T. Schaul, T. Glasmachers, Y. Sun, J. Peters, and J. Schmidhuber, Natural evolution strategies, *J. Mach. Learn. Res.* **15**, 949 (2014).
- [128] T. Salimans, J. Ho, X. Chen, S. Sidor, and I. Sutskever, Evolution Strategies as a Scalable Alternative to Reinforcement Learning, [arXiv:1703.03864](https://arxiv.org/abs/1703.03864) [cs, stat] (2017).
- [129] R. S. Sutton and A. G. Barto: *Reinforcement Learning: An Introduction* (The MIT Press, 2018), 2nd ed.
- [130] Here we discuss the results obtained in the Kitaev chain while in Appendix F we show the results obtained for the more general proximity-coupled semiconducting nanowire.
- [131] R. Chakrabarti and H. Rabitz, Quantum control landscapes, *Int. Rev. Phys. Chem.* **26**, 671 (2007).
- [132] A. G. R. Day, M. Bukov, P. Weinberg, P. Mehta, and D. Sels, Glassy Phase of Optimal Quantum Control, *Phys. Rev. Lett.* **122**, 020601 (2019).
- [133] D. Ghosh, A. Gupta, A. Reddy, J. Fu, C. Devin, B. Eysenbach, and S. Levine, Learning to reach goals via iterated supervised learning, [arXiv:1912.06088](https://arxiv.org/abs/1912.06088) [cs.LG] (2020).
- [134] H.-L. Huang, M. Narożniak, F. Liang, Y. Zhao, A. D. Castellano, M. Gong, Y. Wu, S. Wang, J. Lin, Y. Xu, H. Deng, H. Rong, J. P. Dowling, C.-Z. Peng, T. Byrnes, X. Zhu, and J.-W. Pan, Emulating Quantum Teleportation of a Majorana Zero Mode Qubit, *Phys. Rev. Lett.* **126**, 090502 (2021).
- [135] E. M. Stoudenmire, J. Alicea, O. A. Starykh, and M. P. Fisher, Interaction effects in topological superconducting wires supporting majorana fermions, *Phys. Rev. B* **84**, 014503 (2011).
- [136] E. Sela, A. Altland, and A. Rosch, Majorana fermions in strongly interacting helical liquids, *Phys. Rev. B* **84**, 085114 (2011).
- [137] R. M. Lutchyn and M. P. A. Fisher, Interacting topological phases in multiband nanowires, *Phys. Rev. B* **84**, 214528 (2011).
- [138] F. Crépin, G. Zaránd, and P. Simon, Nonperturbative phase diagram of interacting disordered majorana nanowires, *Phys. Rev. B* **90**, 121407 (2014).
- [139] N. M. Gergs, L. Fritz, and D. Schuricht, Topological order in the Kitaev/majorana chain in the presence of disorder and interactions, *Phys. Rev. B* **93**, 075129 (2016).
- [140] G. Kells, N. Moran, and D. Meidan, Localization enhanced and degraded topological order in interacting p -wave wires, *Phys. Rev. B* **97**, 085425 (2018).
- [141] A. Moon and B. Nachtergaele, Stability of gapped ground state phases of spins and fermions in one dimension, *J. Math. Phys.* **59**, 091415 (2018).
- [142] P. W. Brouwer, M. Duckheim, A. Romito, and F. von Oppen, Topological superconducting phases in disordered quantum wires with strong spin-orbit coupling, *Phys. Rev. B* **84**, 144526 (2011).
- [143] W. DeGottardi, D. Sen, and S. Vishveshwara, Majorana Fermions in Superconducting 1d Systems Having Periodic, Quasiperiodic, and Disordered Potentials, *Phys. Rev. Lett.* **110**, 146404 (2013).
- [144] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pižorn, H. Verschelde, and F. Verstraete, Time-Dependent Variational Principle for Quantum Lattices, *Phys. Rev. Lett.* **107**, 070601 (2011).
- [145] S. Paeckel, T. Köhler, A. Swoboda, S. Manmana, U. Schollwöck, and C. Hubig, Time-evolution methods for matrix-product states, *Ann. Phys. (N. Y.)* **411**, 167998 (2019).
- [146] F. Schäfer, P. Sekatski, M. Koppenhöfer, C. Bruder, and M. Kloc, Control of stochastic quantum dynamics by differentiable programming, *Mach. Learn.: Sci. Technol.* **2**, 035004 (2021).
- [147] M. Abdelhafez, D. I. Schuster, and J. Koch, Gradient-based optimal control of open quantum systems using quantum trajectories and automatic differentiation, *Phys. Rev. A* **99**, 052327 (2019).
- [148] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)* (Springer-Verlag, Berlin, Heidelberg, 2006).
- [149] D. Petz and C. Ghinea, in *Quantum Probability and Related Topics* (World Scientific, Singapore, 2011), Vol. 27, p. 261.
- [150] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, New York, NY, USA, 2011), 10th ed.
- [151] W. P. Su, J. R. Schrieffer, and A. J. Heeger, Solitons in Polyacetylene, *Phys. Rev. Lett.* **42**, 1698 (1979).
- [152] M.-T. Rieder, G. Kells, M. Duckheim, D. Meidan, and P. W. Brouwer, Endstates in multichannel spinless p -wave superconducting wires, *Phys. Rev. B* **86**, 125423 (2012).

- [153] C. Liu, X. Xu, and D. Hu, Multiobjective reinforcement learning: A comprehensive overview, *IEEE Trans. Syst., Man, and Cybernetics: Syst.* **45**, 385 (2015).
- [154] S. R. White and A. E. Feiguin, Real-Time Evolution Using the Density Matrix Renormalization Group, *Phys. Rev. Lett.* **93**, 076401 (2004).
- [155] J. J. Sakurai, *Modern Quantum Mechanics* (Addison-Wesley, Reading, MA, 1994), Rev. ed.
- [156] P. Ring and P. Schuck, *The Nuclear Many-Body Problem*, Physics and Astronomy Online Library (Springer, Berlin Heidelberg, 2004).
- [157] M. B. Giles, in *Advances in Automatic Differentiation*, edited by C. H. Bischof, H. M. Bücker, P. Hovland, U. Naumann, and J. Utke (Springer, Berlin, Heidelberg, 2008), p. 35.
- [158] R. T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. Duvenaud, Neural ordinary differential equations, [arXiv:1806.07366](https://arxiv.org/abs/1806.07366) [cs.LG] (2019).
- [159] C. Rackauckas, Y. Ma, V. Dixit, X. Guo, M. Innes, J. Revels, J. Nyberg, and V. Ivaturi, A Comparison of Automatic Differentiation and Continuous Sensitivity Analysis for Derivatives of Differential Equation Solutions, [arXiv:1812.01892](https://arxiv.org/abs/1812.01892) (2018).
- [160] V. F. Krotov, in *Advances in Nonlinear Dynamics and Control: A Report from Russia*, Progress in Systems and Control Theory, edited by A. B. Kurzhanski (Birkhäuser, Boston, MA, 1993), p. 74.
- [161] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, Optimal control of coupled spin dynamics: Design of NMR pulse sequences by gradient ascent algorithms, *J. Magn. Reson.* **172**, 296 (2005).
- [162] G. Matheron, N. Perrin, and O. Sigaud, The problem with ddpG: Understanding failures in deterministic environments with sparse rewards, [arXiv:1911.11679](https://arxiv.org/abs/1911.11679) [cs.LG] (2019).
- [163] G. Dulac-Arnold, D. Mankowitz, and T. Hester, Challenges of real-world reinforcement learning, [arXiv:1904.12901](https://arxiv.org/abs/1904.12901) [cs.LG] (2019).
- [164] L. Coopmans, K. Kavanagh, I. Jubb, S. Dooley, and G. Kells, On dynamical phase error in interacting topological quantum memories, [arXiv:2101.00022](https://arxiv.org/abs/2101.00022) [cond-mat.str-el] (2020).
- [165] A. E. Antipov, A. Bargerbos, G. W. Winkler, B. Bauer, E. Rossi, and R. M. Lutchyn, Effects of Gate-Induced Electric Fields on Semiconductor Majorana Nanowires, *Phys. Rev. X* **8**, 031041 (2018).
- [166] B. D. Woods, T. D. Stanescu, and S. Das Sarma, Effective theory approach to the Schrödinger-Poisson problem in semiconductor majorana devices, *Phys. Rev. B* **98**, 035428 (2018).