

Controlling nonergodicity in quantum many-body systems by reinforcement learning

Li-Li Ye¹ and Ying-Cheng Lai^{1,2,*}

¹*School of Electrical, Computer and Energy Engineering, Arizona State University, Tempe, Arizona 85287, USA*

²*Department of Physics, Arizona State University, Tempe, Arizona 85287, USA*



(Received 30 August 2024; accepted 17 February 2025; published 10 March 2025)

Finding optimal control strategies to suppress quantum thermalization for arbitrarily initial states, the so-called quantum nonergodicity control, is important for quantum information science and technologies. Previous control methods relied largely on theoretical model of the target quantum system, but invertible model approximations and inaccuracies can lead to control failures. We develop a model-free and deep reinforcement-learning (DRL) framework for quantum nonergodicity control. It is a machine-learning method with the unique focus on balancing exploration and exploitation strategies to maximize the cumulative rewards so as to preserve the initial memory in the time-dependent nonergodic metrics over a long stretch of time. We use the paradigmatic one-dimensional tilted Fermi-Hubbard system to demonstrate that the DRL agent can efficiently learn the quantum many-body system solely through the interactions with the environment. The optimal policy obtained by the DRL provides broader control scenarios for managing nonergodicity in the phase diagram as compared to, e.g., the specific protocol for Wannier-Stark localization. The continuous control protocols and observations are experimentally feasible. The model-free nature of DRL and its versatile search space for control functions render promising nonergodicity control in more complex quantum many-body systems.

DOI: [10.1103/PhysRevResearch.7.013256](https://doi.org/10.1103/PhysRevResearch.7.013256)

I. INTRODUCTION

Quantum nonergodicity has been recognized as a central concept in out-of-equilibrium quantum dynamical systems [1–5]. Relevant physical phenomena include spatial localization such as Anderson [1] and Wannier-Stark localization [2,6], and Hilbert space localization such as quantum many-body scars (QMBS) [7] and many-body localization (MBL) [3]. The unique attribute of quantum nonergodicity in suppressing thermalization has implications to fields ranging from statistical mechanics [8] to quantum information science and technologies [9]. Nonergodicity in quantum many-body systems can be generated by a number of physical mechanisms, each leading to rich and complex quantum phases. Because of the potential of broad applications, controlled generation of quantum nonergodicity has attracted a great deal of recent attention [10–26].

Quantum nonergodicity control aims to find optimal control protocols to suppress quantum ergodicity for diverse initial states. For example, in Floquet engineering, controlled realization of QMBS and MBL through periodic driving is experimentally feasible, but the space of control functions is often limited due to its periodic nature, such as sinusoidal driving [15–18], periodic pulse control [27,28], square

wave [29] and binary driving control of two distinct noncommuting Hamiltonians applied in sequence [19–26], as well as Floquet automata circuit control [30]. In Floquet engineering, the optimization of the periodic control has generally been based on theoretical models [31], where control is optimized by the prior knowledge about the physics of the target system, such as the intrinsic dynamics of the quantum many-body system including subharmonic response and discrete time-crystalline [17,27] as well as conservation and symmetry properties [21,32,33] of such systems. Alternatively, adiabatic approaches and their extensions such as counterdiabatic driving [34–36] and quantum leakage minimization [37,38] combine analytical and numerical methods to find the optimal protocol along a trajectory in the parameter space. These methods provide physical insights and interpretation but inevitably suffer from model inaccuracies and approximations. It is worth noting that traditional optimal control methods such as gradient-free optimal control (e.g., chopped random basis [39–41]) and gradient-based optimal control (e.g., gradient ascent pulse engineering [42–45]) were also model based. While there were works on controlling quantum systems to a specific target state [46,47] or in entanglement engineering [48], to our knowledge, model-free approaches have not been investigated for nonergodicity control in quantum many-body systems.

In this paper, we develop a deep reinforcement learning (DRL) based framework for quantum nonergodicity control. In machine-learning based control, reinforcement learning (RL) has emerged as an effective model-free approach with the capability of finding the optimal strategy in a vast and versatile search space of control functions and the ability to adaptively discover control strategies that the traditional methods

*Contact author: Ying-Cheng.Lai@asu.edu

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.

tend to overlook [49,50]. RL employs a trial-and-error learning process to maximize the cumulative rewards through exploration and exploitation in search for a globally optimal policy. DRL further enhances these capabilities by using deep neural networks to optimize the RL agent. In quantum systems, model-free RL control is capable of generating policies or value functions based solely on the interactions with the quantum environment, without any prior knowledge of the model of that environment [51], in contrast to model-based DRL methods [52] that employ a prebuilt model of the environment to guide policy decisions. One issue is choosing an algorithm that is particularly appropriate for quantum nonergodicity control. We choose the proximal policy optimization (PPO) algorithm [53,54], which is justified, as follows.

In modern machine learning, a number of DRL algorithms have been developed, including those incorporating the trust region proximal optimization (TRPO) algorithm [55], deep-Q network (DQN) [56], and deep deterministic policy gradient (DDPG) [57]. Specifically, TRPO provides a common solution to the local minima challenge in optimal policy search. By confining policy updates within a trust region, TRPO ensures that policies do not deviate too far from previous policies. This mechanism not only enhances stability during the training of the RL agent but also ensures a monotonic improvement in policy search. Among the DQN and DDPG algorithms, PPO stands out as a state-of-the-art algorithm [53,54]. As a hybrid actor-critic approach, PPO adeptly navigates the delicate balance between reducing the variance of policy gradients and diminishing bias linked to the value functions. Operating within the framework of deep neural networks, PPO effectively addresses the curse of dimensionality, thereby enhancing its applicability and scalability in complex environments. PPO achieves data efficiency and reliable performance of TRPO but with a first-order optimization procedure, facilitating implementation with reduced computational complexity.

To demonstrate model-free DRL to achieve nonergodicity control in quantum many-body systems in a concrete setting, we employ the one-dimensional (1D) tilted Fermi-Hubbard model, a paradigm in the study of quantum many-body systems capable of generating a spectrum of weak ergodicity-breaking phenomena [11,12,14]. Our PPO agent relies exclusively on the observables and rewards it receives at each step to make decisions, without requiring an explicit physical model for policy decisions. There are two possible physical observables. The first is spin-resolved imbalance [14], the normalized differences in the occupation numbers of spin-up and spin-down particles between odd and even lattice sites. The second is fidelity that measures the square of the norm of the overlap between the time-evolved quantum state and the initial quantum state based on the full chain, partial chain, or even just a single site [58–60]. A common attribute of quantum nonergodic quantities is their retention of the initial memory over the long time, so the reward function at each time step is directly linked to the observation and can be maximized when the time-evolved observation is consistent with its initial value. The policy is optimized by maximizing the accumulated reward. From this perspective, the DRL agent focuses solely on minimizing the discrepancy between the accumulated fidelity or imbalance

and its initial value. As a result, the agent discovers the optimal policy through direct engagement with the quantum environment, a hallmark of model-free DRL, making it possible for the control algorithm to be implemented in experiments. Another distinct feature of our work, which facilitates experimental implementation, is the use of partial observations or even just a single site observation to control nonergodicity in many-body quantum systems, in contrast to previous works [46,47] in this field that required the complete observations of the current quantum state.

In Sec. II A, we introduce the 1D tilted Fermi-Hubbard model and outline the experimental setting for the observation and control space. In Sec. II B, we clarify the basic concept of PPO agent for quantum nonergodicity control. In Sec. III, we demonstrate the performance of DRL with partial observations. Section IV provides a physical interpretation and understanding of the optimal policy derived from DRL. Conclusions and a discussion about the limitations and potential future research are offered in Sec. V.

II. MODEL AND CONTROL METHOD

A. One-dimensional tilted Fermi-Hubbard model

The Fermi-Hubbard or Bose-Hubbard model with a tilted potential has garnered a great deal of recent interest due to the emerging new physics that may arise commonly in many other quantum many-body systems [10–14,58,61–69]. For example, the tilted potential breaks the translational invariance and integrability, and so can induce subdiffusive transport due to its coupling to mass transport in mass-imbalanced 1D or 2D tilted Fermi-Hubbard models [61–65]. The subdiffusive property is related to the nature of nonergodic dynamics [61]. Moreover, the tilted potential can lead to phenomena such as Hilbert space fragmentation [10,11], QMBS [12,13], and deconfinement dynamics of fractons [66] and non-Fermi liquids [67]. The tilted potential model finds applications across diverse systems such as ultracold fermions in tilted optical lattices [11,14], trapped ions [68], and superconducting qubits [58]. Motivated by experimental breakthroughs and the nonergodicity induced by a titled potential, we use the 1D tilted Fermi-Hubbard model to investigate model-free nonergodicity control.

The Hamiltonian of the 1D tilted Fermi-Hubbard chain model is [12]

$$\hat{H} = \sum_{j,\sigma=\uparrow(\downarrow)} (-J\hat{c}_{j+1,\sigma}^\dagger \hat{c}_{j,\sigma} + \text{h.c.} + \Delta j \hat{n}_{j,\sigma}) + U \sum_j \hat{n}_{j,\uparrow} \hat{n}_{j,\downarrow}, \quad (1)$$

where J denotes the nearest-neighbor coupling, Δ is a uniform tilted potential distributed in position space, and U is the on-site Coulomb interaction. The Hamiltonian includes the fermionic creation ($\hat{c}_{j,\sigma}^\dagger$) and annihilation ($\hat{c}_{j,\sigma}$) operators, as well as the number operator $\hat{n}_{j,\sigma} = \hat{c}_{j,\sigma}^\dagger \hat{c}_{j,\sigma}$. For simplicity, we consider a lattice with \mathcal{N} sites, where the spin-up and spin-down fermions are equally distributed, denoted by

$$\mathcal{N}_\uparrow = \mathcal{N}_\downarrow = \mathcal{N}/2.$$

This setup implies an electron filling factor of $\nu = 1$, as in a previous work in the Fock basis [12]. We assume periodic

boundary conditions, where the spin direction is maintained when a particle hops across the boundary. At any given lattice site, the occupation by the spin-up or spin-down electrons is represented as $|\uparrow\rangle$ or $|\downarrow\rangle$, respectively, while an empty site is indicated as $|0\rangle$. A site simultaneously occupied by both spin up and down electrons, known as a doublon [12], is denoted by $|\uparrow\downarrow\rangle$.

Experimentally, a number of platforms are available for controlling the system described by the 1D tilted Fermi-Hubbard chain model. For example, in an optical lattice, a tilted potential can be modulated by a magnetic field gradient [11,14] and the on-site Coulomb interaction U is tunable via a magnetic Feshbach resonance [14,70]. Specifically, the tilted Fermi-Hubbard model in an optical lattice can be characterized by the imbalance [14]:

$$\mathcal{I}^{\uparrow(\downarrow)} \equiv \frac{\mathcal{N}_o^{\uparrow(\downarrow)} - \mathcal{N}_e^{\uparrow(\downarrow)}}{\mathcal{N}_o^{\uparrow(\downarrow)} + \mathcal{N}_e^{\uparrow(\downarrow)}}, \quad (2)$$

where $\mathcal{N}_o^{\uparrow(\downarrow)}$ and $\mathcal{N}_e^{\uparrow(\downarrow)}$ are the occupation numbers of the spin-up and spin-down electrons at the odd and even lattice sites, respectively. There are also experimental techniques [18,70–73] that allow for an independent manipulation of $\Delta(t)$ and $U(t)$ over time, offering precise control over the system's dynamics. In addition, superconducting qubits quantum simulators [58–60] with an integrated and programmable large-scale platform offer flexibility in the control protocols, where quantum tomography measurements offer direct experimental access to the components of the reduced density matrix [58–60]. This capability enables precise measurements of the fidelity [58,60] for both the subchain and the full chain, denoted as \mathcal{F}_{sub} and $\mathcal{F}_{\text{full}}$, respectively. In fact, the time evolution of Von Neumann entanglement entropy for the subchain has been observed [60]. The computational approaches for evaluating the nonergodic metrics are described in Appendix A.

For convenience and clarity, in the following, we present our results using the time unit $\tau \equiv \hbar/J$ with the reduced Planck constant \hbar and the nearest-neighbor coupling strength J . The potential terms Δ and U are expressed in units of J , as outlined in Appendix A. The quantum dynamics of the 1D tilted Fermi-Hubbard chain are governed by the Schrödinger equation. We employ Trotter decomposition [74,75] with the discrete time step of size $dt = 0.005\tau$. A detailed accuracy analysis can be found in Appendix A.

B. Method of DRL based quantum nonergodicity control

Quantum nonergodicity describes out-of-equilibrium phenomenon that arise when a quantum system resists thermalization or equilibration even after long time evolution. It implies that time-evolved quantum states retain the memory of their initial conditions, remaining closely aligned with them over prolonged duration. Training a DRL agent to realize quantum nonergodicity control thus entails maintaining the time-evolved quantum states as closely as possible to their initial, pure, and unentangled states throughout the evolution process. The scenario of DRL training is illustrated in Fig. 1.

The DRL training contains the following components.

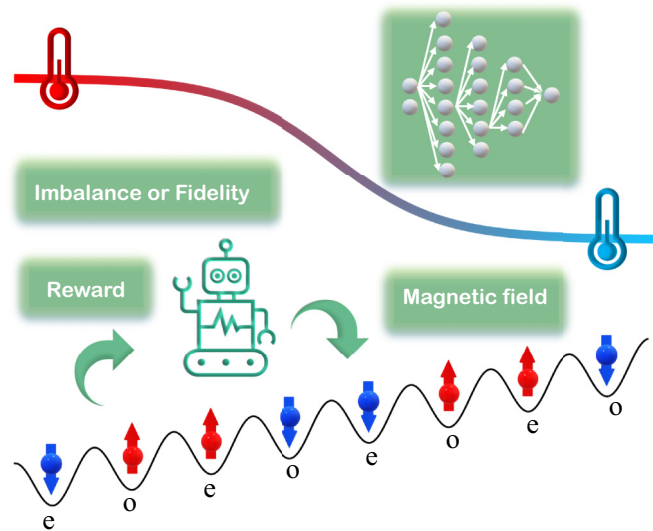


FIG. 1. Scenario of quantum nonergodicity control by model-free DRL. The PPO algorithm coupled with neural networks is used for control. The task involves training a randomly initialized agent to discover the optimal control protocol for steering quantum nonergodicity in the 1D tilted Fermi-Hubbard model. This system is initialized in a specific quantum state within the Fock space. The observation space is one of the following: (1) imbalance metrics $[\mathcal{I}^{\uparrow}, \mathcal{I}^{\downarrow}]$ representing the normalized differences in the occupation numbers of spin-up and spin-down particles between the odd and even lattice sites, (2) partial fidelity \mathcal{F}_{sub} that provides a partial view of the lattice chain dynamics, and (3) full fidelity $\mathcal{F}_{\text{full}}$ that offers a complete observation of the entire lattice chain. The agent receives the corresponding rewards, with the maximum value specifically designed to incentivize the maintenance of the initial state throughout the time evolution. Following policy updates aimed at maximizing the cumulative reward over time, the agent intervenes in the quantum many-body system by applying a designated tilted potential and an on-site Hubbard potential, which can be experimentally implemented via a physical control field, such as a magnetic field. As the accumulated reward converges, the degree of quantum ergodicity gradually diminishes.

Initialization. The environment for the DRL PPO agent to learn is the quantum many-body system: the 1D tilted Fermi-Hubbard chain formulated in Eq. (1). The quantum state is initialized in the Fock space. For example, for a lattice of size $N = 8$, two initial states are

$$\begin{aligned} | - + - + \rangle &= |\downarrow\uparrow\uparrow\downarrow\downarrow\uparrow\uparrow\downarrow\rangle, \\ | + - + - \rangle &= |\uparrow\downarrow\downarrow\uparrow\uparrow\downarrow\downarrow\rangle, \end{aligned}$$

where $|-\rangle \equiv |\downarrow\rangle$ and $|+\rangle \equiv |\uparrow\rangle$. For conciseness, we denote the two kinds of initial states as $| - + \rangle$ and $| + - \rangle$, respectively. In a lattice system with open boundaries, under the approximation $\Delta \approx U \gg J$ these states are in fact QMBS states in the corresponding effective Hamiltonian model [12]. Other permutations of spin configurations in the Fock space tested in our work include $|\uparrow\downarrow\downarrow\uparrow\uparrow\downarrow\downarrow\rangle$, $|\uparrow\downarrow\uparrow\uparrow\downarrow\downarrow\rangle$, $|\downarrow\downarrow\uparrow\uparrow\downarrow\downarrow\rangle$, $|\uparrow\downarrow\uparrow\downarrow\uparrow\uparrow\downarrow\rangle$, and $|\uparrow\uparrow\downarrow\downarrow\uparrow\uparrow\downarrow\rangle$.

Observation Space. Three physical quantities are used for the DRL agent to observe the environment: (1) the imbalance

vector $[\mathcal{I}^\uparrow(t), \mathcal{I}^\downarrow(t)]$ with $\mathcal{I}^{\uparrow(\downarrow)}(t) \in [-1, 1]$, (2) partial-chain fidelity $\mathcal{F}_{\text{sub}}(t) \in [0, 1]$, and (3) full-chain fidelity $\mathcal{F}_{\text{full}}(t) \in [0, 1]$. It is worth noting that, in experiments, detecting imbalance or partial fidelity, or even simply observing a single lattice site, can be more efficient than observing the full-chain fidelity. However, having only partial information about the quantum system poses challenges for optimizing control protocols.

Action Space. The configuration of the observation and action spaces determines computational and control complexity. We adopt continuous observation and action spaces for optimal policy search. Specifically, the action space comprises the tilted potential $\Delta(t)$ and the on-site Hubbard potential $U(t)$, each ranging from -10J to 10J . The limited range of these global control fields is chosen to avoid trivial solutions, such as ideal Wannier-Stark localization ($\Delta \gg J$ and $U = 0$), which would lead to nonergodic behavior but are restricted by experimental feasibility. Despite the discretization of time evolution, the value at each time step remains continuous within the range $[-10, 10]\text{J}$ and can be parameterized by a deep neural network. This defines a continuous action space of infinite dimension.

Reward Design. The design of reward functions is tailored to the quantity of observation. For the imbalance vector, the reward function takes the form:

$$R(t) = -|\mathcal{I}^\downarrow(t) - \mathcal{I}^\downarrow(0)| - |\mathcal{I}^\uparrow(t) - \mathcal{I}^\uparrow(0)|. \quad (3)$$

Alternatively, if the observation space involves the sub- or full-chain fidelity, the reward function becomes

$$R(t) = -|\sqrt{\mathcal{F}(t)} - 1|. \quad (4)$$

Under this setup, the agent incurs a negative penalty for deviations from the initial states, encouraging it to maintain proximity to the initial configuration.

Training. The training of the DRL agent relies on a delicate balance of exploration and exploitation strategies, which is crucial for learning the optimal policy to maintain quantum nonergodicity. In particular, achieving this balance is essential for uncovering the effective strategies for sustaining quantum nonergodicity. The DRL agent must explore diverse actions to comprehend their impact on the quantum system, while also exploiting established strategies to maximize the reward. This iterative process requires that the agent interact with the quantum system, observe the resulting states, and improve its policy based on the received rewards, as shown in Fig. 1. As outlined in Appendix B, both the actor and critic utilize the independent neural networks with the identical size. For various tasks, we adopt two alternative neural-network configurations with three hidden layers in a multilayer perceptron: **Config A** - NN = [256, 128, 64] with a learning rate of 10^{-4} ; **Config B** - a smaller neural network NN = [128, 64, 32] with the learning rate 0.5×10^{-3} . The PPO agent is implemented using the Reinforcement Learning Toolbox in MATLAB.

III. DRL BASED QUANTUM NONERGODICITY CONTROL: RESULTS

A. Illustration of control performance

The distinct feature of quantum nonergodicity lies in its capacity to preserve the memory of the initial unentangled states. In the ideal Wannier-Stark localization scenario [2,6] with an infinitely strong tilted potential ($\Delta \gg J$) and zero on-site Coulomb interaction ($U = 0$), the nonergodic property can be represented by time-evolved quantities, including the full-chain fidelity $\mathcal{F}(t) = 1$, the half-chain entropy $\mathcal{S}(t) = 0$, and the imbalance $\mathcal{I}(t) = \mathcal{I}(0)$ and $\mathcal{I}^{\uparrow(\downarrow)}(t) = \mathcal{I}^{\uparrow(\downarrow)}(0)$ for arbitrarily long time. (Numerical verification for the near-ideal Wannier-Stark localization is described in Appendix A 6 for $\Delta = 100J$ and $U = 0J$.) As a result, the average quantity over each episode with the maximum time horizon \mathcal{T} should satisfy

$$\langle \mathcal{F} \rangle_{\mathcal{T}} = 1, \quad \langle \mathcal{S} \rangle_{\mathcal{T}} = 0, \quad \langle \mathcal{I} \rangle_{\mathcal{T}} = \mathcal{I}(0), \quad \text{and} \\ \langle \mathcal{I}^{\uparrow(\downarrow)} \rangle_{\mathcal{T}} = \mathcal{I}^{\uparrow(\downarrow)}(0).$$

These physical quantities serve as nonergodic metrics, delineating the deviation from the truth. Evaluating the performance of the trained DRL agent relies on its capability to maintain quantum nonergodicity over an extended time horizon, which involves measuring deviations in the nonergodic metrics.

In principle, the full-chain fidelity encapsulates the full quantum information about the quantum state. However, partial observation can be more efficient and feasible in experimental settings. For instance, the 1D tilted Bose-Hubbard model has been successfully realized in superconducting processors [58]. Moreover, quantum tomography measurements in superconducting qubits allow for direct acquisition of the elements of the reduced density matrix [60], enabling observations such as the half-chain entropy and subchain fidelity. In optical lattices, ultracold fermions can be controlled by a magnetic field to simulate the 1D tilted Fermi-Hubbard model with spin-resolved imbalance [14]. These developments make nonergodic metrics accessible in experiments.

The nonergodic metrics are expected to iteratively approach the nonergodic truths during the training phase and demonstrate the retention of the initial state memory during the testing phase. In the training phase with the time horizon $\mathcal{T} = 5\tau$, the observation results in two distinct tasks: the \mathcal{F}_{sub} task with $\mathcal{N}_{\text{sub}} = 1$ or 4 sites and the $[\mathcal{I}^\uparrow, \mathcal{I}^\downarrow]$ task. Figures 2(a) and 2(b) show, respectively, the convergence of the subchain fidelity $\langle \mathcal{F} \rangle_{\mathcal{T}}$ and the spin-resolved average imbalance $\langle \mathcal{I}^{\uparrow(\downarrow)} \rangle_{\mathcal{T}}$. Figure 2(c) shows that a convergence of the consequent full-chain fidelity $\langle \mathcal{F} \rangle_{\mathcal{T}}$ has been achieved, agreeing with the behaviors in Figs. 2(a) and 2(b). During the testing phase for $0 < t \leq 5\tau$, the full-chain fidelity $\mathcal{F}(t)$ and half-chain entropy $\mathcal{S}(t)$, and the imbalances $\mathcal{I}(t)$ and $\mathcal{I}^{\uparrow(\downarrow)}(t)$ exhibit oscillations about their respective nonergodic truths, as shown in Figs. 2(d), 2(g), and 2(h). The oscillatory behavior is originated from Bloch oscillations [14,76] and the optimal control protocol, as shown in Figs. 2(e) and 2(f). The full-chain fidelity closely approaches the nonergodic truth value during the training, where quantum many-body thermalization

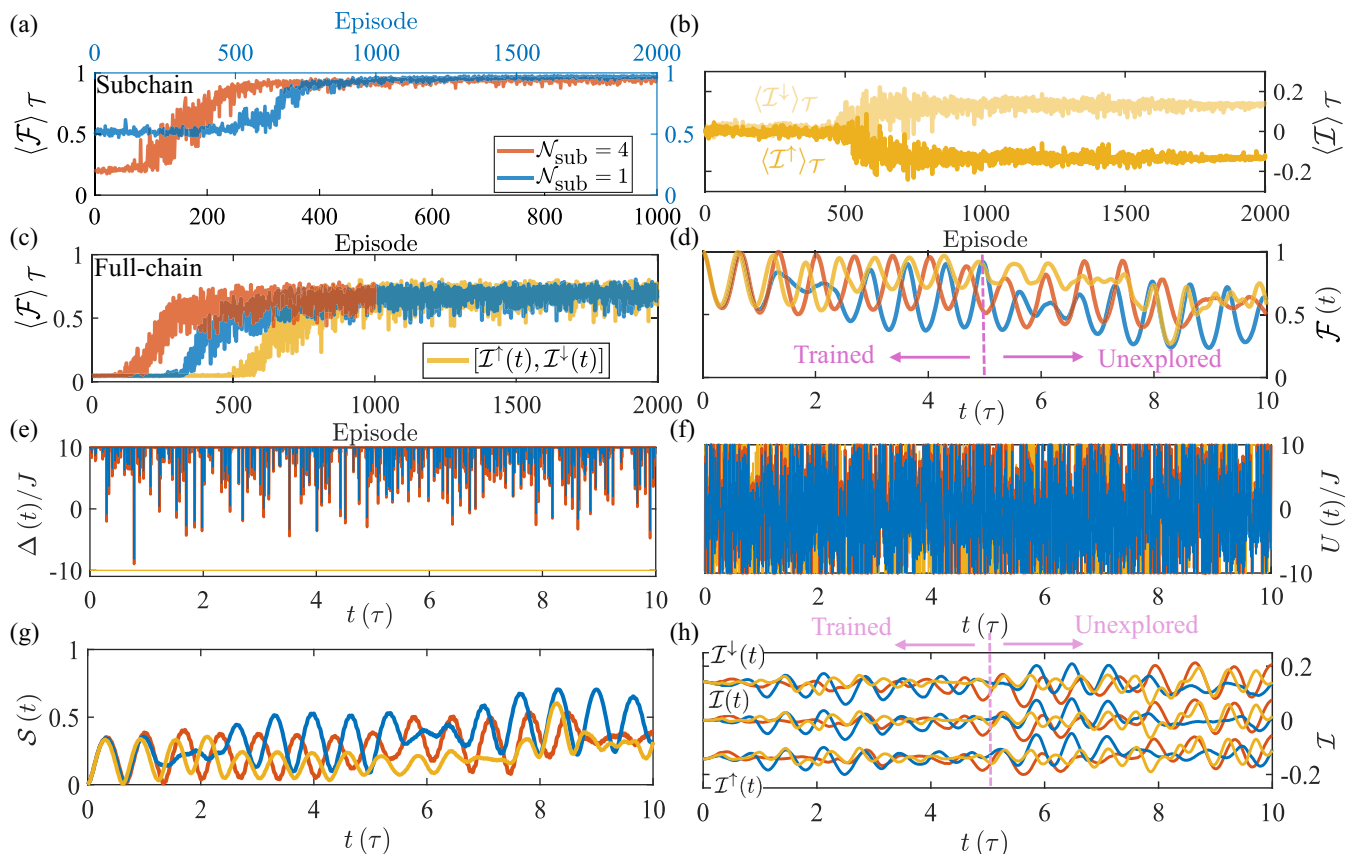


FIG. 2. Evaluation of quantum nonergodicity control by DRL agent. The 1D tilted Fermi-Hubbard chain has the size $\mathcal{N} = 14$. The observations are performed using (1) subchain fidelity \mathcal{F}_{sub} with $\mathcal{N}_{\text{sub}} = 1$ or 4 sites from the left-hand side of the chain, and (2) the imbalance vector $[\mathcal{I}^{\uparrow}, \mathcal{I}^{\downarrow}]$ from the entire chain. These options constitute distinct tasks shown with different colors in (a)–(h), with $\mathcal{N}_{\text{sub}} = 1, 4$ by \mathcal{F}_{sub} , and the $[\mathcal{I}^{\uparrow}, \mathcal{I}^{\downarrow}]$ tasks in blue, orange, and yellow curves, respectively. The initial state for both tasks is identical: $|-+-+--+-\rangle$. (a),(b) Training phase for the \mathcal{F}_{sub} and $[\mathcal{I}^{\uparrow}, \mathcal{I}^{\downarrow}]$ tasks, respectively. The learning curves chart the average fidelity $\langle \mathcal{F} \rangle_{\mathcal{T}}$ and the average imbalance $\langle \mathcal{I} \rangle_{\mathcal{T}}$ over the time horizon $\mathcal{T} = 5\tau$ in episodes, reflecting episodic learning with policy updates and quantum state resets after each episode. (c) Convergence of the resulting $\langle \mathcal{F} \rangle_{\mathcal{T}}$ within the full chain, which agrees with that of the subchain task in (a) and the imbalance task in (b). (d)–(h) Results from the testing phase, where the test time horizon is $\mathcal{T} = 10\tau$, encompassing the unexplored capabilities of the DRL agent. The nonergodic metrics, including $\mathcal{F}(t)$, $S(t)$, and $\mathcal{I}(t)$, demonstrate the success of nonergodicity control through the optimal action flow discovered by the DRL agent, as depicted in (e) and (f). The neural network size for the two tasks with $\mathcal{N}_{\text{sub}} = 1$ and $\mathcal{I}^{\uparrow(\downarrow)}$ is from **Config A**, and that with $\mathcal{N}_{\text{sub}} = 4$ is from **Config B** (specified in Sec. II B).

is greatly suppressed, as can be seen from evolution of the half-chain entropy. The spin-resolved imbalances also oscillate about the nonzero initial value with a small amplitude. The three nonergodic metrics, distinguished by three different colors, exhibit comparable behaviors, indicating similar performance for the \mathcal{F}_{sub} and $[\mathcal{I}^{\uparrow}, \mathcal{I}^{\downarrow}]$ tasks. In the untrained region $5\tau < t \leq 10\tau$, the nonergodic metrics oscillate more wildly with a slightly increased amplitude as compared to the trained time region in Figs. 2(d), 2(g), and 2(h), implying the potential role of time prediction and controllability of DRL in the unexplored region.

B. Scalability and complexity of deep reinforcement learning

In general, RL deals with sequential decision-making problems, so the complexity of the PPO agent algorithm involves not only the number of samples but also the quality and variety of the quantum environment that the agent encounters. For the tilted Fermi-Hubbard chain, the lattice size determines

the dimension of the Hilbert space in which rich quantum states or phases arise. A longer time horizon increases the time complexity for the PPO agent. When applied to a quantum many-body system, a key attribute of the PPO agent algorithm is the scalability of performance with the lattice size, training time complexity, and the sample number. These factors can directly influence the feasibility of the DRL control in larger systems. To study the scalability, we define the number of samples as the sample complexity, in which the time complexity can be directly encoded, and systematically test the performance for various lattice sizes and training time horizons. The results are shown in Fig. 3. Specifically, Figs. 3(a)–3(c) present the results for lattice sizes $\mathcal{N} = [6, 8, 10, 12, 14]$, demonstrating stable performance of DRL even in the unexplored time horizons and suggesting the feasibility of extending DRL nonergodicity control into a larger Hilbert space. For the worst-case sample complexity, i.e., the number of observation points at each time over the whole training process, is about $10^6 - 10^7$ for

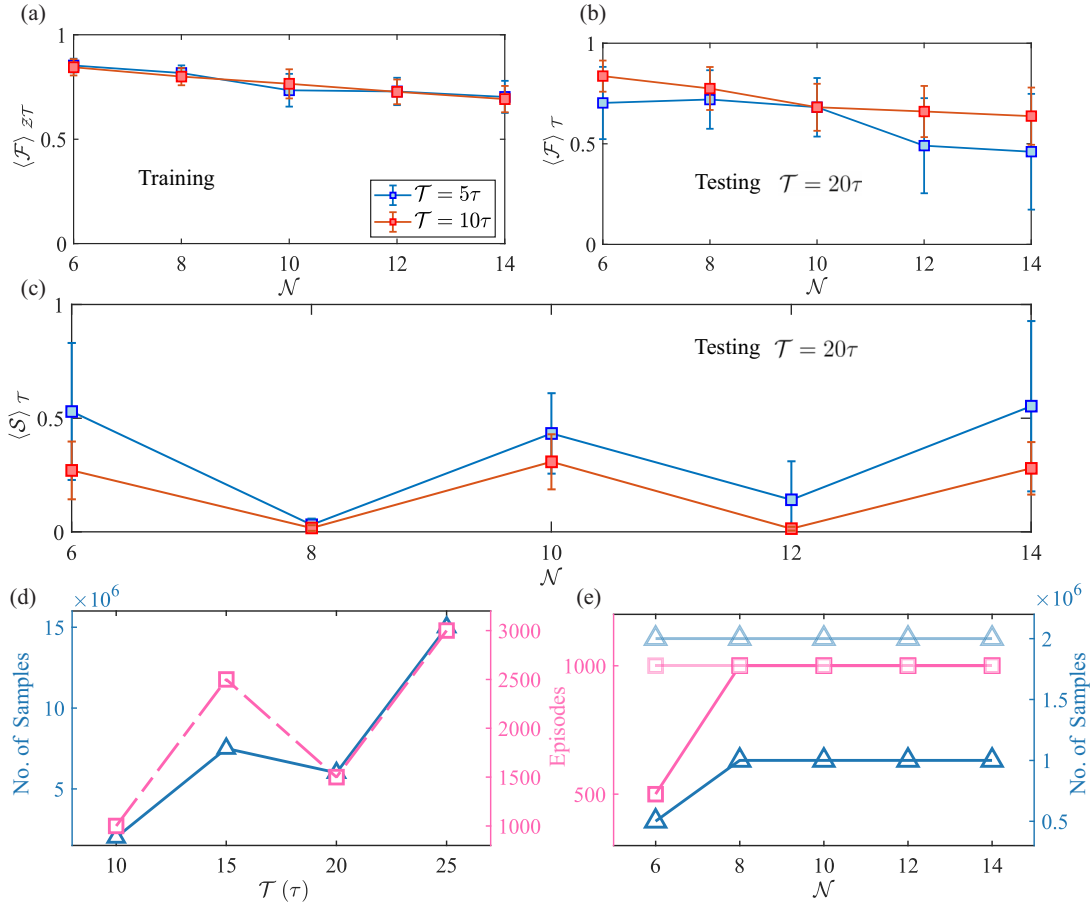


FIG. 3. Scalability and complexity of DRL performance and training. The observation is full-chain fidelity for lattice sizes $\mathcal{N} = [6, 8, 10, 12, 14]$ with the initial state $|-\ +\rangle$. (a) The training outcomes, denoted by $\langle \mathcal{F} \rangle_{\mathcal{Z}\mathcal{T}}$, are assessed through $\mathcal{Z} = 10$ episodes and calculated from the last ten episodes of the training phase. This metric provides a quantitative measure of the performance of the PPO agent in multiple training sessions. Two different training time horizons are also used: $\mathcal{T} = 5\tau$ and 10τ . The average full fidelity $\langle \mathcal{F} \rangle_{\mathcal{Z}\mathcal{T}}$ and the standard deviation characterize the scalability and stability of the learning process. (b), (c) Testing results for the time horizon $\mathcal{T} = 20\tau$ (including the unexplored horizon), in terms of the full-chain fidelity and half-chain entropy. The PPO agent is trained under two different time horizons: $\mathcal{T} = 5\tau$ and 10τ . The error bars denote the standard deviation within a single episode, providing a measure of the variability and reliability of the agent's performance across different training duration. (d) Worst-case sample complexity and the corresponding training episodes across various training time horizons. The worst-case scenarios are represented by the blue curves and corresponding axes, which is from seven distinct initial states that include $|-\ +\ -\ +\rangle$, $|+\ -\ +\ -\rangle$, and five other permutations of spin configurations: $|\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow\rangle$, $|\uparrow\downarrow\uparrow\downarrow\downarrow\uparrow\rangle$, $|\downarrow\uparrow\downarrow\uparrow\downarrow\uparrow\rangle$, $|\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow\rangle$, and $|\uparrow\downarrow\downarrow\uparrow\uparrow\downarrow\rangle$. The training episodes are determined by assessing the qualitative convergence of the full-chain fidelity $\langle \mathcal{F} \rangle_{\mathcal{T}}$ for a chain of size $\mathcal{N} = 8$, as represented by the pink curve and axis. The neural network configuration is **Config A**. (e) Sample complexity and the corresponding training episodes versus lattice chain size for the fixed initial state $|-\ +\rangle$. The full-chain fidelity varies distinctly across two time horizons. For $\mathcal{T} = 5\tau$, the DRL employs the neural network configuration of **Config B**, shown by dark blue and pink curves. For a longer time horizon of $\mathcal{T} = 10\tau$, the training DRL utilizes the **Config A**, as illustrated by light blue and pink curves.

the seven tested initial states in Fig. 3(d). Especially, the sample complexity shows that the performance is largely independent of the lattice size, as a result of the nonergodicity control mechanism, i.e., the tendency to gradually converge to the approximated control in the single-particle picture.

IV. PHYSICAL INTERPRETATION AND ROBUSTNESS OF DRL

To understand the mechanism of quantum nonergodicity control, we recall the phenomenon of Anderson localization [1,77,78] in the single-particle picture. It arises from

independent and identical random chemical potentials assigned to each lattice site described by the tight-binding model. The disorders characterized by the magnitude of the range in the random on-site potential, if sufficiently strong, will disrupt the quantum ergodicity. Single-particle localization can also occur without disorders. For example, substituting the random potential with a uniform electric field can lead to Wannier-Stark localization [2,6]. This effect has been observed under a sufficiently strong tilted potential in a superconducting quantum processor [58].

The picture of single-particle localization provides insights into the phenomenon of MBL [3]: the augmentation of Anderson localization with constant on-site Coulomb interac-

tion [73]. It was also found that quasirandom disorders in the chemical potentials in the Aubry-Andre model [4,5,70] can lead to MBL. Moreover, a strong random Coulomb interaction at each site has been demonstrated to facilitate the onset of MBL [73,79,80]. The spatial distribution of the electric field further facilitates MBL in the presence of many-body interactions. For instance, Stark many-body localization [81] emerges under a nonuniform electric field when a harmonic term is present that breaks the pure linearity of the electric field. In addition, sufficiently strong random fields superimposed on a uniform electric field can trigger MBL [82]. It is worth noting that MBL resides within the realm of strong ergodicity breaking, whereas QMBS corresponds to weak ergodicity breaking, both violating the eigenstate thermalization hypothesis [83–85], where the quantum dynamics of QMBS depend on the initial conditions rooted in the disconnected structure in the Hilbert space.

Deep RL delivers optimal control to induce nonergodicity in a quantum many-body system. Figures 2(e) and 2(f) show the optimally controlled trajectories for $\Delta(t)$ and $U(t)$, respectively. For the control based on $\Delta(t)$, it tends to converge to a constant value: either $\Delta/J = 10$ or -10 , where the sign is due to the different titled directions within the lattice chain. Alternatively, incorporating random perturbations into the constant Δ also represents a potential control protocol by DRL. For optimal protocol based on $U(t)$, it oscillates within the original search range: $-10 \leq U(t)/J \leq 10$. To understand the quantum phases that DRL learns and why it converges to some specific values as exemplified in Figs. 2(e) and 2(f), we simplify the optimal protocol of $\Delta(t)$ as $\Delta(t)/J = \Delta_0/J = 10$, where Δ_0 is a constant, or incorporates random perturbation as $\Delta(t) = \Delta_0 + hw(t)$ with the constant $h/J \in [0, 10]$ and random number $w(t) \in [-1, 0]$ at each time step, ensuring it stays within the range $\Delta(t)/J \in [0, 10]$ as in Fig. 2(e). For the optimal protocol of $U(t)$, we simplify it as $U(t) = Uw'(t)$ with the constant $U/J \in [0, 10]$ and random numbers $w'(t) \in [-1, 1]$ at each time step, limiting $U(t)/J$ between -10 and 10 . For convenience, we use the term “Deep RL-aligned protocol” to denote the simplified control protocols. The quantum phase generated by the DRL-aligned protocol is referred to as the “Deep RL-aligned phase”. Comparing the DRL-aligned and other quantum phases, especially the QMBS with the constant $\Delta \approx U \gg J$, entails testing four pairs of actions: $(\Delta, Uw'(t))$, (Δ, U) , $(\Delta_0 + hw(t), U)$, and $(\Delta_0 + hw(t), Uw'(t))$, as illustrated in Fig. 4.

We use the average fidelity $\langle \mathcal{F} \rangle_{\mathcal{T}}$ for the whole chain to interpret DRL and characterize the quantum phases, which is justified, as follows. In episodic learning, the DRL agent collects a sequence of observations, rewards, and actions within each episode, subsequently updating its policy to maximize the accumulated reward for future training. Observations could consist of fidelity or imbalance, with the reward function directly linked to the observation and aimed at converging to its initial value. As a result, DRL is designed to focus solely on minimizing the discrepancy between the accumulated fidelity or imbalance and its initial value. Plotting the average fidelity over one episode offers a way to understand the physical mechanisms of learning with DRL. While imbalance only provides partial information about the quantum state, fidelity encompasses the entire chain, making it an appropriate

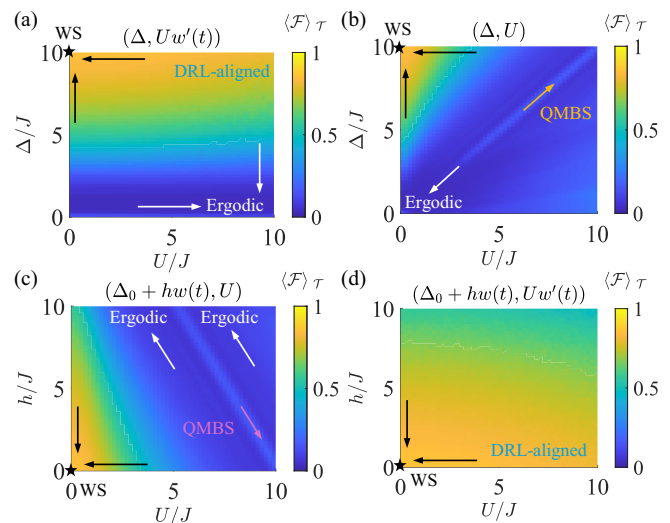


FIG. 4. Physical interpretation and robustness of DRL for quantum nonergodicity control in terms of the phase diagrams of the fidelity. (a)–(d) Control parameter space for the full-chain fidelity $\langle \mathcal{F} \rangle_{\mathcal{T}}$ in different parameter planes: $(\Delta, Uw'(t))$, (Δ, U) , $(\Delta_0 + hw(t), U)$, and $(\Delta_0 + hw(t), Uw'(t))$, where the random numbers $w(t) \in [-1, 0]$ and $w'(t) \in [-1, 1]$ are independent, identical, and uniformly distributed at each time step for a fixed $\Delta_0 = 10J$, and constants Δ , U , and h . The fidelity $\langle \mathcal{F} \rangle_{\mathcal{T}}$ is calculated using $\mathcal{Z} = 100$ disorder averages, with the system initialized in the state $| - + - + \rangle$ in a lattice of size $\mathcal{N} = 8$ over the time horizon $\mathcal{T} = 10\tau$ that matches the testing time horizon of DRL in Fig. 2.

indicator. We note that the average fidelity was used to characterize the quantum phases in a previous work [82].

To understand why the optimal protocol converges to a specific region, as demonstrated in Figs. 2(e) and 2(f), we examine the quantum phase diagram to reveal what the DRL agent has learned for nonergodicity control tailored to the quantum many-body system. The testing time horizon \mathcal{T} dictates the temporal span for constructing the quantum phase diagram of $\langle \mathcal{F} \rangle_{\mathcal{T}}$. Despite the short observation time, the phase diagram highlights the DRL-aligned regime and distinct quantum phases, as illustrated in Fig. 4 for the full lattice model, where QMBS, thermalization or ergodic, and Wannier-Stark phases are displayed. More details are shown by the related time-dependent nonergodic metrics in Fig. 5.

For $\Delta \approx U \gg J$, the first-order Schrieffer-Wolff transformation [86] is applicable, which can be used to derive the effective model of the full lattice system [12]. The QMBS states $| - + \rangle$ and $| + - \rangle$ associated with weak ergodic breaking exhibit switching dynamics within a hypergrid structure characterized by the time-dependent nonergodic metrics, as described in a previous work [12]. The inherent tower structure of the overlap between QMBS states and the eigenstates [12] is indicative of a concentration about some specific energy levels and violation of eigenstate thermalization hypothesis. In our work, we use the full lattice model with periodic boundary conditions and the approximation $\Delta \approx U \approx 10J$ to find QMBS states with behavior similar to that of the corresponding states in the effective model, as shown in Figs. 4(b) and 4(c) and Figs. 5(a)–5(d). The phase

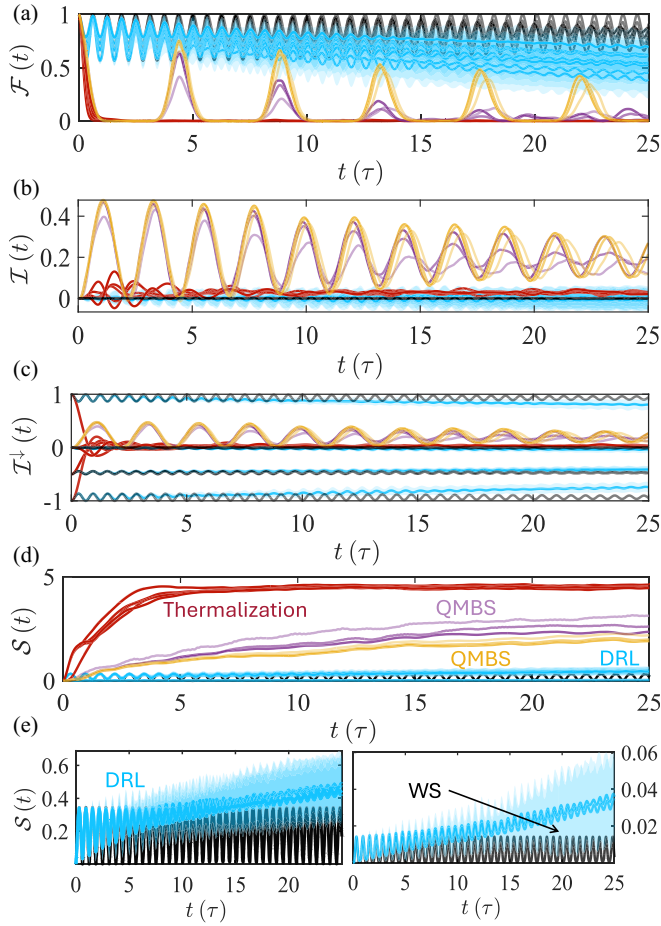


FIG. 5. Comparison of DRL control with other protocols in terms of the quantum phases by using time-dependent nonergodic metrics. These metrics include the (a) full-chain fidelity $\mathcal{F}(t)$, (b) imbalance $\mathcal{I}(t)$, (c) spin-down imbalance $\mathcal{I}^\downarrow(t)$, and (d,e) the half-chain Von Neuman entropy $S(t)$. The analysis is conducted over the time horizon of $\mathcal{T} = 25\tau$ in a lattice of size $\mathcal{N} = 8$. Beginning with the QMBS state $| - + \rangle$ [12], the QMBS phase is illustrated by the yellow curves for $\Delta/J = U/J = 6, 8, 10$ (from lighter to darker shades). A random perturbation yields the QMBS state in purple for $(\Delta_0 + hw(t), U) = (10 + 10w(t), 5.2)J, (10 + 5.8w(t), 7.3)J, (10 + 0.5w(t), 10)J$, displayed in progressively darker curves. Also shown are DRL results and two other quantum phases, the Wannier-Stark and thermalization phases. Seven initial states are used, including the typical QMBS, $| - + - + \rangle, | + - + - \rangle$ and five other permutations of spin configurations: $|\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow\rangle, |\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow\rangle, |\downarrow\uparrow\downarrow\uparrow\downarrow\uparrow\rangle, |\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow\rangle, |\uparrow\uparrow\downarrow\downarrow\uparrow\uparrow\rangle$ (plotted in identical colors). The Wannier-Stark phase is characterized by constants $\Delta/J = 10$ and $U/J = 0$, while thermalization features constants $\Delta/J = 1, U/J = 2$. The three quantum phases serve as benchmarks for evaluating the performance of DRL. The blue curves depict the testing outcomes of DRL, averaged over $\mathcal{Z} = 100$ independent testing episodes. The standard deviations are also shown. Deep RL is trained to learn the quantum system using the full-chain fidelity observable with the training time of $\mathcal{T} = 25\tau$ and neural network size of **Config A** as described in Sec. II B.

diagram of $\langle \mathcal{F} \rangle_{\mathcal{T}}$ in Fig. 4(b) reveals the presence of the QMBS phase along the direction indicated by the orange arrow, with a darker blue area nearby. This finding agrees with an earlier result [12] on the $U - \Delta$ phase diagram featuring the first peak of the imbalance. Figure 4(c) reveals that random perturbations to the tilted potential result in a similar QMBS pattern, suggesting the robustness of the QMBS states. Additional results supporting the robustness are shown by the time-dependent nonergodic metrics in Figs. 5(a)–5(d), where the perturbations lead to a quicker decay of the revival amplitude in the full-chain fidelity and imbalance, though with a slight rise in the entanglement entropy.

Within the time horizon $\mathcal{T} = 10\tau$, there is a tendency towards ergodic behavior in the quantum dynamics, as indicated by the white arrows in the phase diagrams in Figs. 4(a)–4(c). For $\Delta/J = 1$ and $U/J = 2$, the short-time thermalization process is demonstrated in Figs. 5(a)–5(d). There is a swift decline and stabilization in the fidelity and imbalance, accompanied by a rapid convergence of the entanglement entropy. The Wannier-Stark phase can usually be characterized by $\Delta \gg J, U$ but, due to the limited action range, the action pair of $\Delta = 10J$ and $U = 0J$ emerges as the closest approximation to the ideal Wannier-Stark phase, as shown in both Figs. 4 and 5. The regions in the vicinity of the Wannier-Stark point in the phase diagrams exhibit a tendency towards the Wannier-Stark phase, as indicated by the black arrows in Figs. 4(a)–4(d). In terms of the time-dependent nonergodic metrics, the Wannier-Stark phase serves as an ideal benchmark for initial memory retention and nonergodicity control. It maintains the initial values of the fidelity, imbalance, and entropy, in spite of the Bloch oscillations [76] of the period $t_B = 2\pi/\Delta$. Within the search space $\Delta/J \in [-10, 10]$ and $U/J \in [-10, 10]$, the Wannier-Stark phase can simply be regarded as a specific point at $\Delta/J = 10$ and $U/J = 0$, which is sensitive to constant perturbations from the on-site Coulomb interaction U as indicated in Figs. 4(b) and 4(c) but is robust against perturbations $hw(t)$ in the tilted potential Δ_0 to some extent, as shown in Figs. 4(c) and 4(d). The DRL-aligned protocol also reveals the robustness against perturbations for $Uw'(t)$, as shown in Figs. 4(a) and 4(d), highlighting a broad control scheme where DRL converges to maximize the accumulated fidelity.

While the results in Fig. 4 are from the DRL-aligned protocol, the real DRL control flow and the corresponding performance of the nonergodic metrics are demonstrated in Fig. 5. In particular, Fig. 5(e) reveals a consistency between the DRL protocol and Wannier-Stark fidelity in the short term, exhibiting the same period of Bloch oscillations. However, in the long run, a slight deviation in the nonergodic metrics emerges, representing the trade-off between robustness and performance of nonergodicity control in the DRL action protocol.

The control method can be extended to training over longer time horizons, where some unexplored characteristics of the Wannier-Stark phase can be revealed. Over an extended period, the Wannier-Stark phase exhibits both short-term Bloch oscillations, as illustrated in Fig. 5, and envelope oscillations in the time-dependent series, as shown in Fig. 6(a). The period

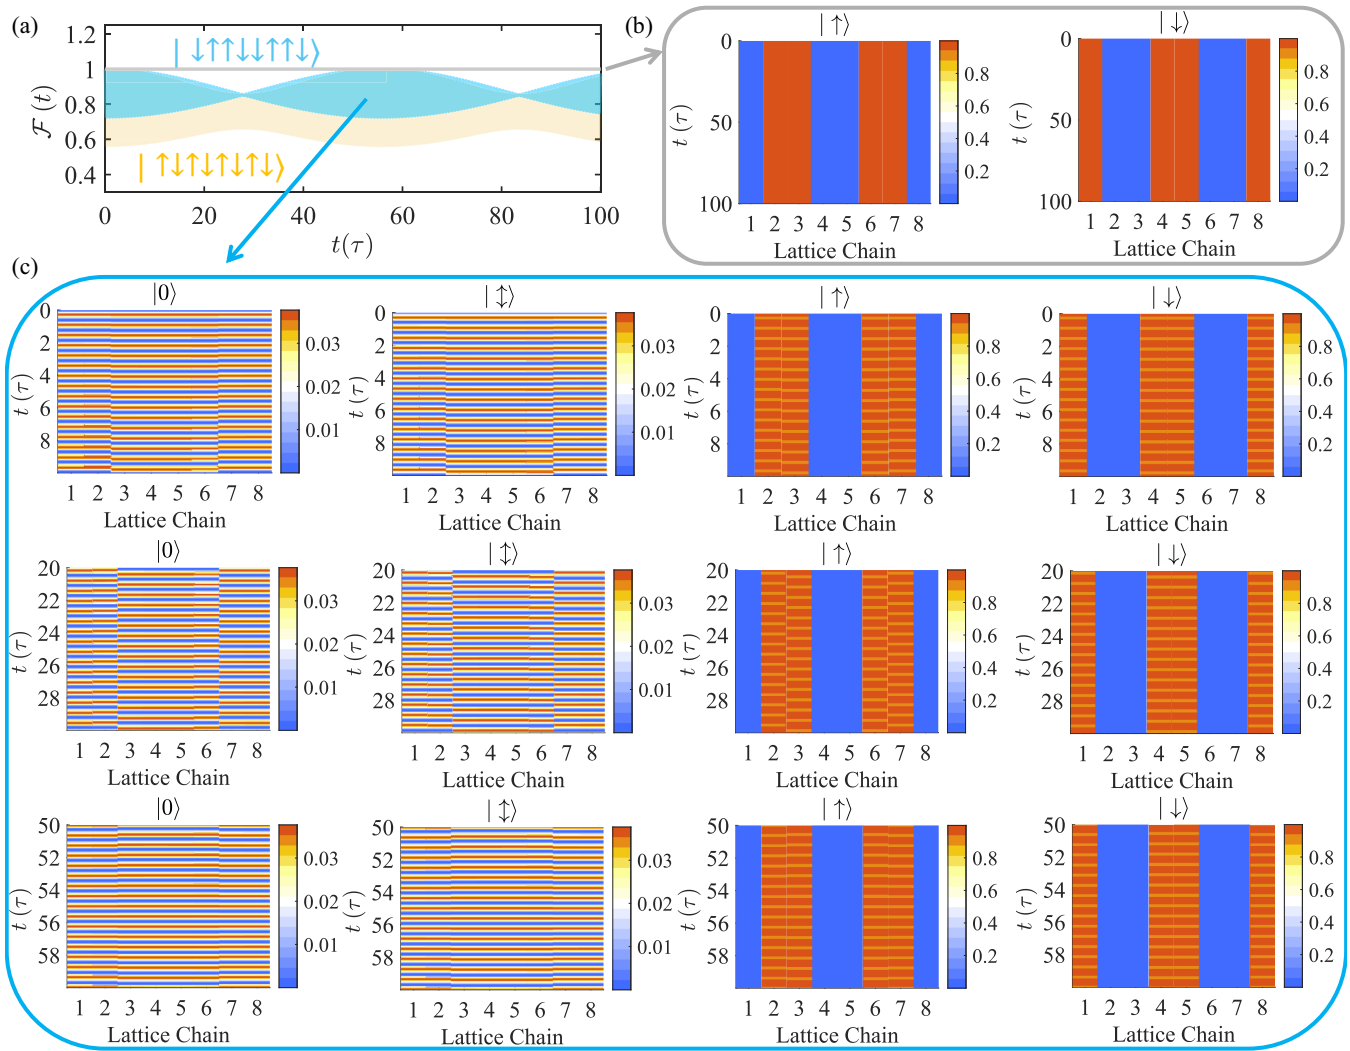


FIG. 6. Interpretation of the envelope oscillation of the Wannier-Stark fidelity in a lattice chain with $\mathcal{N} = 8$ sites. (a) Periodic oscillations in the envelope amplitude of the time-dependent fidelity of the full lattice chain, along with the short-time Bloch oscillation. The grey and blue curves represent simulations starting from the same initial state $| - + - + \rangle$, with different action pairs: $(\Delta, U) = (100, 0)J$ and $(10, 0)J$, respectively. The yellow curves are initialized with another state $|\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow\rangle$, with the action pair $(\Delta, U) = (10, 0)J$. (b) Ideal Wannier-Stark phase visualized by the probability distribution in the 2D parameter plane spanning space and time. (c) Time slice of the space-time panel for $t \in [0, 10]\tau$, $[20, 30]\tau$, $[50, 60]\tau$, which displays the probability distribution in the basis $|0\rangle$, $|\downarrow\rangle$, $|\uparrow\rangle$, and $|\downarrow\rangle$ at each site location.

of the envelope oscillations is positively correlated with the tilted potential, while the period of the Bloch oscillations has a negative correlation with it. The amplitude of the Bloch oscillations and the shape of the envelope oscillations are influenced by the initial state conditions, as illustrated in Figs. 5(e) and 6(a), respectively. The ideal Wannier-Stark phase is shown in Fig. 6(b), which can be used as a benchmark. Figure 6(c) reveals that the spatial oscillations of a quantum state contribute to the envelope oscillation. In general, simulating quantum many-body systems is challenging due to the difficulty of exponential growth in the computational complexity with the system size, but experiments are possible. With the numerical validation of the feasibility of DRL for controlling a small system, the training methodology can in principle be extended to experiments with larger systems.

V. DISCUSSION

In complex quantum many-body systems, thermalization leading to ergodicity is a major source of decoherence. Developing methods to suppress thermalization to achieve nonergodicity is essential for applications, e.g., in quantum information science and technology. AI-based optimal control provides the possibility of controlling complex quantum many-body systems to achieve nonergodicity. Utilizing the 1D tilted Fermi-Hubbard model as a paradigm, we developed a model-free DRL approach to controlling quantum nonergodicity, where the DRL agent interacts with the quantum environment stipulated by the quantum many-body system. During the online training phase, the DRL agent, specifically a PPO agent, collects time series data in real time within one episode, which include observations, rewards, and actions. A

criterion to choose the type of observations is their experimental accessibility in real time. For the 1D tilted Fermi-Hubbard chain, both the full or partial chain fidelity [58,60] and the spin-resolved imbalance [14] satisfy this criterion. The reward values are determined by the observable variable and a function tailored to meeting the nonergodic objective that the metrics preserve the initial memory over a long time. Consequently, the reward function is crafted to maintain the time-dependent nonergodic metric as close to its initial value as possible. The stochastic policy dictates subsequent actions, $\Delta(t)$ and $U(t)$ for the tilted Fermi-Hubbard chain, can physically be implemented through a properly designed magnetic field [11,14,70] within an episode. This policy is a stochastic probability distribution over all possible actions conditioned on the given observable at that time. The policy is updated after each episode to maximize the accumulated reward over one episode. The well-trained PPO agent is saved after a predetermined number of training episodes designed to ensure the complete convergence of the mean reward training curve. In the online testing phase during which the optimal policy determined by training is not updated, the well-trained PPO agent applies the optimal nonergodic control to the same quantum system in real time. The optimal policy of the well-trained PPO agent gives the subsequent actions based on its observations. The control is completely data driven in the sense that, in the whole training and testing process, no prior knowledge about the target quantum many-body system is required: all needed is model-free DRL with experimentally available observables. The quantum phases that the DRL agent has learned can be used to understand the physical mechanisms underlying the optimal control policy.

Two patterns arising from the DRL control are that, initially, the absolute value of the titled potential term tends to reach a constant maximum value, while the on-site Coulomb interaction appears random within the original search range (DRL-aligned policy). We explored all possible combinations of these two patterns and discovered a rich array of quantum phases through various phase diagrams, including ergodic, QMBS, and Wannier-Stark phases. Both the QMBS and Wannier-Stark phases are robust against perturbations to the tilted potential, but the latter is sensitive to the constant on-site Coulomb potential and tends to thermalize. In contrast, the DRL-aligned policy offers a broad control scenario for perturbations to either the tilted term or the on-site Coulomb interaction. By comparing these phases with the actual DRL policy over time series, we observed that the DRL policy closely aligns with the Wannier-Stark phase. A physical analysis indicates that, under the condition of an infinite tilted term, the Wannier-Stark phase approaches the ideal state, which is indicative of single-particle localization. This provides a simplistic protocol for our control task. In general, the DRL protocol offers superior control robustness, with performance comparable to the Wannier-Stark phase under the nonergodicity control objective. Another appealing feature of DRL nonergodicity control is that the observations from even just one site suffice for realizing the control goal, facilitating experimental implementation.

In a recent work [87], the authors employed a model-based machine learning method, specifically the variational entanglement-enhancing field, to optimize the magnetic field

in a quantum spin chain system, thereby enabling persistent ballistic entanglement spreading. By adjusting the local parameters in the Hamiltonian, this approach accelerates the saturation of ergodicity beyond that which would typically occur in a homogeneous, time-independent Hamiltonian. This result is somewhat distinct from that of our work. In particular, while the authors of Ref. [87] sought to expedite ergodicity from an initial product state, the goal of our work is to prevent its onset. In spite of the different control objectives, both their work and ours highlight the potential of using a time-dependent control field to enable a wide range of physical implementations, whether for accelerating or suppressing ergodicity. The use of machine learning to optimize control fields in quantum systems provides a powerful and versatile tool for tackling more complex quantum control tasks.

ACKNOWLEDGMENTS

We thank Prof. L. Ying for comments and suggestions, Dr. J.-L. Wang for a discussion on the entanglement-entropy calculation, and Dr. A. Yang for general discussions on quantum many-body systems. We also thank the group of Prof. M. Aidelburger for providing their published codes. This work was supported by the Air Force Office of Scientific Research under Grant No. FA9550-21-1-0438 and by the Office of Naval Research under Grant No. N00014-24-1-2548.

DATA AVAILABILITY

The data and code for this work are available at GitHub [88]

APPENDIX A: 1D TILTED FERMION-HUBBARD MODEL

1. Fock basis

The behavior of interacting fermions, constrained to move along a 1D lattice and interacting via the on-site Coulomb interaction, is described by the 1D tilted Fermi-Hubbard model in the presence of an external linear potential. Here, fermions are represented as spin-1/2 particles with spin up (down), with the respective numbers \mathcal{N}_\uparrow and \mathcal{N}_\downarrow among the \mathcal{N} lattice sites. Accordingly, the number of bases for the spin up (down), denoted by d_\uparrow (d_\downarrow), is determined as

$$d_{\uparrow(\downarrow)} = \binom{\mathcal{N}}{\mathcal{N}_{\uparrow(\downarrow)}}. \quad (\text{A1})$$

In the Fock space, the entire basis is constructed by combining the basis states for spin up and spin down:

$$\hat{c}_1^\dagger \hat{c}_2^\dagger \dots \hat{c}_{i_{\mathcal{N}_\uparrow}}^\dagger \hat{c}_{j_1}^\dagger \hat{c}_{j_2}^\dagger \dots \hat{c}_{j_{\mathcal{N}_\downarrow}}^\dagger |0\rangle,$$

which corresponds to one-to-one pairs of tuples:

$$((i_1, i_2, \dots, i_{\mathcal{N}_\uparrow}), (j_1, j_2, \dots, j_{\mathcal{N}_\downarrow})) = (\alpha, \beta), \quad (\text{A2})$$

where a symbol, such as $i_{\mathcal{N}_\uparrow}$ or $j_{\mathcal{N}_\downarrow}$, records the occupied site location of the \mathcal{N}_\uparrow -th spin up or \mathcal{N}_\downarrow -th spin down particle at the lattice site.

The total number of possible pairs of tuples is $d_\uparrow \times d_\downarrow$, representing the Hilbert space dimension in the 1D tilted Fermi-Hubbard model with particle-number conservation.

The general quantum state over the 1D spin-lattice can be expanded in the Fock basis as

$$|\psi\rangle = \sum_{\alpha, \beta} |\alpha, \beta\rangle \langle \alpha, \beta | \psi \rangle \equiv \sum_{\alpha, \beta} M_{\alpha\beta}^{(\psi)} |\alpha, \beta\rangle. \quad (\text{A3})$$

Computationally, the quantum state can be denoted by the matrix $M^{(\psi)}$ of dimension $d_{\uparrow} \times d_{\downarrow}$.

The Hamiltonian governing the time evolution of the quantum state is

$$\begin{aligned} \hat{H} = & -J \sum_j (\hat{c}_{j+1, \uparrow}^{\dagger} \hat{c}_{j, \uparrow} + \text{h.c.}) - J \sum_j (\hat{c}_{j+1, \downarrow}^{\dagger} \hat{c}_{j, \downarrow} + \text{h.c.}) \\ & + U \sum_j \hat{n}_{j, \uparrow} \hat{n}_{j, \downarrow} + \Delta \sum_j j \hat{n}_{j, \uparrow} + \Delta \sum_j j \hat{n}_{j, \downarrow}, \end{aligned}$$

where \hat{c}_j and \hat{c}_j^{\dagger} are the fermionic annihilation and creation operators, respectively, and \hat{n}_j is the particle number operator distinguished by spin up and down. The parameters J , U , and Δ denote the nearest-neighbor hopping strength, the on-site Coulomb interaction, and the strength of the tilted potential, respectively. We use periodic boundary conditions so as to maintain the continuity of the state across the boundary. For example, the transition of a spin-up fermion from the last site to the first site is represented as $\hat{c}_{1, \uparrow}^{\dagger} \hat{c}_{L, \uparrow}$. The periodic boundary conditions ensure that the lattice behaves as if it were looped, allowing for seamless transitions of particles across the boundary.

To obtain the time evolution of the Schrödinger system, we refer to previous works [14] and modify the Hamiltonian as

$$\begin{aligned} \hat{H} = & (\hat{H}_{\uparrow}^{\text{hop}})_{d_{\uparrow} \times d_{\uparrow}} \otimes \mathbb{I}_{d_{\downarrow} \times d_{\downarrow}} + \mathbb{I}_{d_{\uparrow} \times d_{\uparrow}} \otimes (\hat{H}_{\downarrow}^{\text{hop}})_{d_{\downarrow} \times d_{\downarrow}} \\ & + (\hat{H}^{\text{diag}})_{(d_{\uparrow} \times d_{\downarrow}) \times (d_{\uparrow} \times d_{\downarrow})}, \end{aligned}$$

where each spin type only hops within its respective subspace \hat{H}^{hop} and the diagonal components of \hat{H}^{diag} are represented as V :

$$\begin{aligned} V_{\alpha, \beta} = & \sum_{k=1}^{N_{\uparrow}} V_{i_k, \uparrow} + \sum_{k=1}^{N_{\downarrow}} V_{j_k, \downarrow} \\ & + U | \{i_1, i_2, \dots, i_{N_{\uparrow}}\} \cap \{j_1, j_2, \dots, j_{N_{\downarrow}}\} |, \end{aligned}$$

where $| \{i_1, i_2, \dots, i_{N_{\uparrow}}\} \cap \{j_1, j_2, \dots, j_{N_{\downarrow}}\} |$ denotes the count of the identical elements between the two sets. Consequently, the Schrödinger equation in the 1D tilted Fermi-Hubbard model is simplified to

$$i\hbar \dot{|\psi\rangle} = (\hat{H}_{\uparrow}^{\text{hop}} \otimes \mathbb{I}_{\downarrow}) |\psi\rangle + (\mathbb{I}_{\uparrow} \otimes \hat{H}_{\downarrow}^{\text{hop}}) |\psi\rangle + \hat{H}^{\text{diag}} |\psi\rangle, \quad (\text{A4})$$

where the dot over the quantum state $|\psi\rangle$ denotes its time derivative, specifically:

$$|\dot{\psi}\rangle \equiv \partial |\psi\rangle / \partial t = \sum_{\alpha, \beta} \dot{M}_{\alpha\beta}^{(\psi)} |\alpha\beta\rangle.$$

Normalizing both sides of Eq. (A4) with the hopping strength J , we obtain a dimensionless equation. The time unit is defined as $\tau \equiv \hbar/J$, and the potential terms Δ and U are expressed in units of J .

The Schrödinger equation can be recast in the matrix form, expanded in the Fock basis $|\alpha\beta\rangle$, through the following trans-

formation:

$$\begin{aligned} (\hat{H}_{\uparrow}^{\text{hop}} \otimes \mathbb{I}_{\downarrow}) |\psi\rangle &= \sum_{\gamma\beta} M_{\gamma\beta}^{(\psi)} \left(\sum_{\alpha} H_{\uparrow, \alpha\gamma}^{\text{hop}} |\alpha\rangle_{\uparrow} \right) \otimes |\beta\rangle_{\downarrow} \\ &= \sum_{\alpha\beta} \left(H_{\uparrow}^{\text{hop}} M^{(\psi)} \right)_{\alpha\beta} |\alpha\beta\rangle, \\ (\mathbb{I}_{\uparrow} \otimes \hat{H}_{\downarrow}^{\text{hop}}) |\psi\rangle &= \sum_{\alpha\gamma} M_{\alpha\gamma}^{(\psi)} |\alpha\rangle_{\uparrow} \otimes \left(\sum_{\beta} H_{\downarrow, \beta\gamma}^{\text{hop}} |\beta\rangle_{\downarrow} \right) \\ &= \sum_{\alpha\beta} \left(M^{(\psi)} H_{\downarrow}^{\text{hop}} \right)_{\alpha\beta} |\alpha\beta\rangle, \\ \hat{H}^{\text{diag}} |\psi\rangle &= \sum_{\alpha, \beta} M_{\alpha\beta}^{(\psi)} \hat{H}^{\text{diag}} |\alpha\beta\rangle \\ &= \sum_{\alpha, \beta} (V \circ M^{(\psi)})_{\alpha\beta} |\alpha\beta\rangle. \end{aligned}$$

The matrix form of the Schrödinger equation is given by

$$i\hbar \dot{M}^{(\psi)} = H_{\uparrow}^{\text{hop}} M^{(\psi)} + M^{(\psi)} H_{\downarrow}^{\text{hop}} + V \circ M^{(\psi)}. \quad (\text{A5})$$

Based on the Lie-Trotter-Suzuki product formula [74,75], we can represent the quantum dynamics through iterative time evolution of matrices:

$$M^{(\psi)}(t + \delta t) \approx e^{-i\delta t V(t)} \circ e^{-i\delta t H_{\uparrow}^{\text{hop}}} M^{(\psi)}(t) e^{-i\delta t H_{\downarrow}^{\text{hop}}}, \quad (\text{A6})$$

where the symbol “ \circ ” denotes the element-wise multiplication and the exponential operator in $e^{-i\delta t V(t)}$ means the element-wise exponentiation.

2. Full-chain observable

To illuminate the time evolution process of quantum many-body systems, certain experimentally measurable physical quantities are utilized. One such quantity is fidelity $\mathcal{F}(t)$, which quantifies the overlap between the initial quantum state and its time-evolved counterpart. Mathematically, fidelity is expressed as

$$\mathcal{F}(t) = | \langle \psi_0 | e^{-i\hat{H}t} | \psi_0 \rangle |^2 = | \sum_{\alpha, \beta} M_{\alpha\beta}^{*(0)} M_{\alpha\beta}^{(t)} |^2.$$

When the quantum state extends over the entire chain, the resulting full-chain fidelity \mathcal{F} captures the complete quantum information embedded in the state.

Another metric is imbalance, which represents the normalized differences in the particle occupation numbers between odd and even lattice sites:

$$\mathcal{I} = \frac{\mathcal{N}_o - \mathcal{N}_e}{\mathcal{N}_o + \mathcal{N}_e}, \quad (\text{A7})$$

where \mathcal{N}_o and \mathcal{N}_e denote the occupation numbers at odd and even lattice locations, respectively. The expected value of the imbalance for a specific quantum state is calculated as

$$\langle \mathcal{I} \rangle = \sum_{\alpha\beta} |M_{\alpha\beta}|^2 \mathcal{I}_{\alpha\beta}. \quad (\text{A8})$$

The spin-resolved version of the imbalance is defined as

$$\mathcal{I}^{\uparrow(\downarrow)} = \frac{\mathcal{N}_o^{\uparrow(\downarrow)} - \mathcal{N}_e^{\uparrow(\downarrow)}}{\mathcal{N}_o^{\uparrow(\downarrow)} + \mathcal{N}_e^{\uparrow(\downarrow)}}, \quad (\text{A9})$$

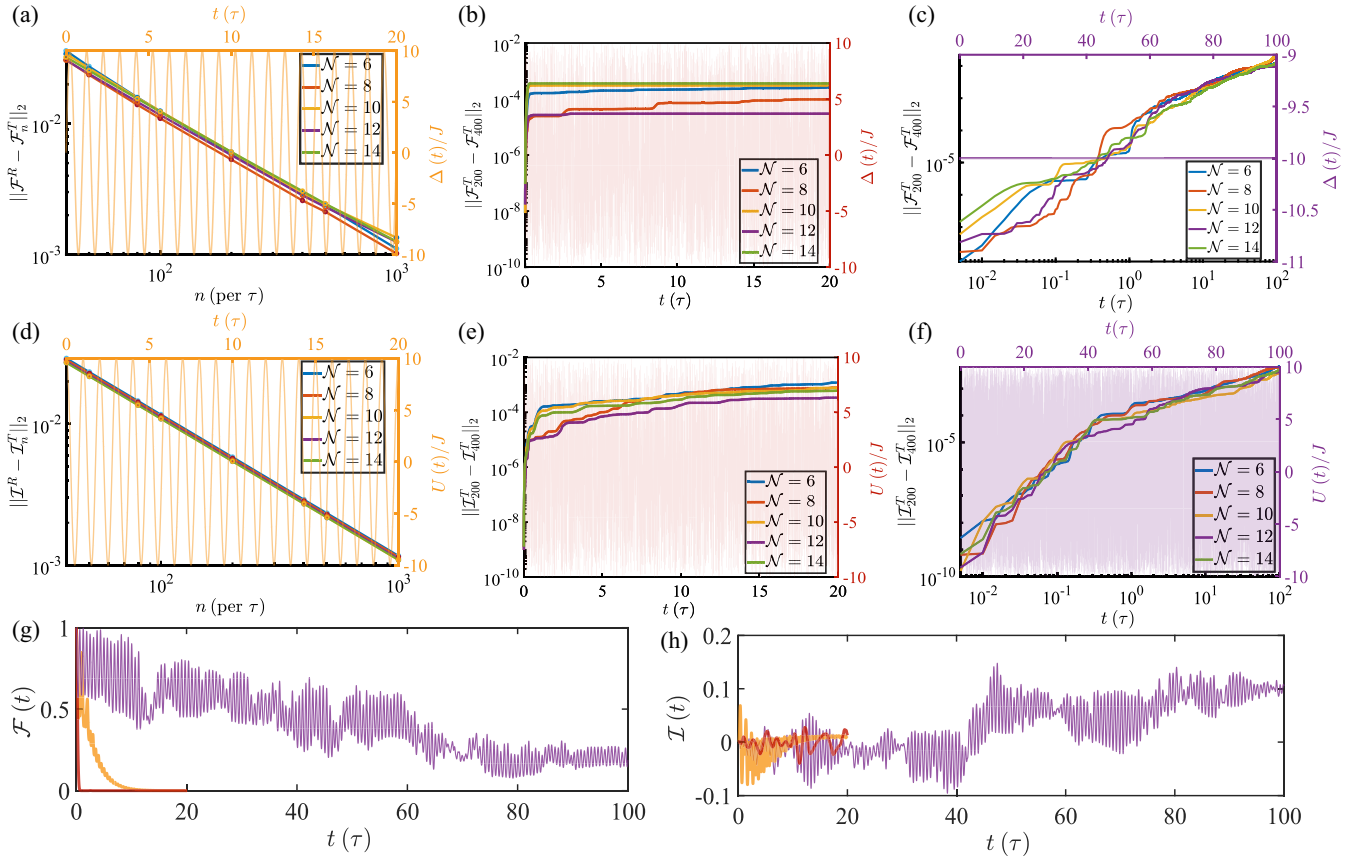


FIG. 7. Numerical error analysis of Trotter decomposition for the full-chain fidelity and imbalance over short- and long-time horizons with various lattice numbers. (a), (d) Benchmark for Trotter decomposition in the short-time regime ($\mathcal{T} = 20\tau$) obtained by the fourth-order Runge-Kutta method with the time step $dt = 10^{-4}\tau$. This setup discretizes time into $n = 10^4$ steps per time unit τ . The L^2 -norm compares the Trotter and Runge-Kutta methods in terms of the fidelity and imbalance. The system dynamics are periodically driven by $\Delta(t)/J = 10 \sin(2\pi t)$ and $U(t)/J = 10 \cos(2\pi t)$, depicted by the orange curves and axes. (b), (e) L^2 -norm of the Trotter decomposition errors between $n = 200$ and $n = 400$ steps per unit τ in the short time regime. The driving signals, $\Delta(t)$ and $U(t)$, are uniformly and randomly distributed within the range $[-10, 10]J$ over time, represented by the red curves and axes. (c), (f) Numerical error by the Trotter method over a longer time scale ($\mathcal{T} = 100\tau$) for constant $\Delta = -10J$ and randomly driven $U(t)$, uniformly distributed within $[-10, 10]J$, shown as the purple curves and axes. (g, h) The time evolution of the full-chain fidelity and imbalance with $n = 200$ discrete steps per time unit τ in a system of size $\mathcal{N} = 14$.

indicating the occupation imbalance for spin-up and spin-down particles.

3. Accuracy of Trotter decomposition

To quantify the numerical errors of simulated observables derived from the Trotter decomposition, we use the L^p -norm [14, 89]:

$$\begin{aligned} |\mathcal{F}_{\text{Runge}} - \mathcal{F}_{\text{Trotter}}^n|_p &= \left(\int_0^{\mathcal{T}} |\mathcal{F}_R(t) - \mathcal{F}_T^n(t)|^p dt \right)^{1/p}, \\ |\mathcal{I}_{\text{Runge}} - \mathcal{I}_{\text{Trotter}}^n|_p &= \left(\int_0^{\mathcal{T}} |\mathcal{I}_R(t) - \mathcal{I}_T^n(t)|^p dt \right)^{1/p}, \end{aligned}$$

with $p = 1, 2, \dots$. For $p = \infty$, the norms are defined as

$$\begin{aligned} |\mathcal{F}_{\text{Runge}} - \mathcal{F}_{\text{Trotter}}^n|_{\infty} &= \max(|\mathcal{F}_R(t) - \mathcal{F}_T^n(t)|), \\ |\mathcal{I}_{\text{Runge}} - \mathcal{I}_{\text{Trotter}}^n|_{\infty} &= \max(|\mathcal{I}_R(t) - \mathcal{I}_T^n(t)|), \end{aligned}$$

where $\mathcal{F}_{\text{Runge}}$ and $\mathcal{I}_{\text{Runge}}$ are obtained by the fourth-order Runge-Kutta method with the time step $dt = 10^{-4}\tau$, equivalent to $n = 10^4$ steps per time unit τ .

4. Quantum dynamics

The Runge-Kutta method, which directly solves the matrix equation (A5), serves as a benchmark for assessing the Trotter method. Using two different discrete steps $n = 200$ and $n = 400$ per time unit τ can yield the numerical error in the Trotter decomposition. We assess the numerical errors across three distinct driven protocols: (1) periodic driving: $\Delta(t)/J = 10 \sin(2\pi t)$ and $U(t)/J = 10 \cos(2\pi t)$, (2) random driving: $\Delta(t)$ and $U(t)$ uniformly distributed in $\in [-10, 10]J$, and (3) special protocol: constant $\Delta = -10J$ with uniformly random driving $U(t)$ in the interval $[-10, 10]J$. For these scenarios, the L^2 -norm metrics of the fidelity and imbalance reveal certain scaling behaviors across various lattice sizes $\mathcal{N} = 6, 8, 10, 12, 14$, as shown in Figs. 7(a)–7(f). These metrics are consistent with respect to both the number of discrete

steps n per time unit τ and the scaled time $t(\tau)$. It can be seen that the protocol combining constant Δ with random $U(t)$ outperforms the other two control protocols in retaining the memory of the quantum state, as illustrated in Fig. 7(g).

5. Subchain observable

To calculate the observables of a quantum state on a subchain, the bases of the Hilbert space need to be reorganized and divided into two subspaces: l and r . Observables in the subspace l are obtained by tracing out the opposing subspace r . The number of lattice sites in the subchain l , denoted as \mathcal{N}_l and counted from the left, defines the scope of this subspace. Accordingly, the number of bases in the left-hand l and right-hand r subchains are $d_l = 4^{\mathcal{N}_l}$ and $d_r = 4^{\mathcal{N}_r}$, respectively, with the relationship $\mathcal{N}_l + \mathcal{N}_r = \mathcal{N}$. Each lattice site hosts one of the four possible states: empty $|0\rangle$, spin up $|\uparrow\rangle$, spin down $|\downarrow\rangle$, and doublon $|\downarrow\rangle$ (spin up and down simultaneously). The total number of bases is given by

$$d_l \times d_r = 4^{\mathcal{N}}, \quad (\text{A10})$$

reflecting the exponential scaling with the number \mathcal{N} of lattice sites. In this framework, the quantum many-body state can be expressed as

$$|\psi\rangle = \sum_{l,r} \psi_{l,r} |l\rangle \otimes |r\rangle, \quad (\text{A11})$$

where l and r are tuples with

$$l \equiv (l_1, l_2, \dots, l_{\mathcal{N}_l}), \\ r \equiv (r_1, r_2, \dots, r_{\mathcal{N}_r}).$$

Here, $l_i, r_j = 0, 1, 2, 3$ corresponds to four possible states for each site. In the quantum many-body system, particle number is conserved, so the size of the Hilbert space for the configuration with $\mathcal{N}/2$ spins up and down is described by Eq. (A1).

Mapping the focused quantum state into this subchain subspace without particle number conservation, the elements in the density matrix are sparsely distributed and structured as follows:

$$\rho = |\psi\rangle\langle\psi| = \sum_{l,l',r,r'} \psi_{lr} \psi_{l'r'}^* |l\rangle\langle l'| \otimes |r\rangle\langle r'|.$$

After tracing out the right-hand subchain subspace r , the reduced density matrix ρ_l is given by

$$\rho_l \equiv \text{Tr}_r \rho = \sum_{r''} \langle r'' | \psi \rangle \langle \psi | r'' \rangle \\ = \sum_{l,l'} \sum_{r,r',r''} \psi_{lr} \psi_{l'r'}^* \delta_{r,r''} \delta_{r',r''} |l\rangle\langle l'| = \psi \psi^\dagger. \quad (\text{A12})$$

Similarly, the reduced density matrix for subspace r , denoted as ρ_r , can be obtained as $\rho_r = (\psi^\dagger \psi)^T$. In Eq. (A11), the matrix ψ is used to represent the pure quantum many-body state, with the abstract notations l and r corresponding to the row and column in state matrices, respectively.

To enable a calculation of the Von Neumann entanglement entropy in a numerically efficient manner, we carry out a singular value decomposition of the quantum many-body state matrix ψ :

$$\psi_{d_l \times d_r} = U \Sigma V^\dagger, \quad (\text{A13})$$

where $\Sigma_{d_l \times d_r}$ is the rectangular diagonal matrix, $U_{d_l \times d_l}$ and $V_{d_r \times d_r}$ are unitary matrices. The nonzero entries along the main diagonal of Σ represent the real singular values of the matrix ψ . Consequently, the reduced density matrices can be written rewritten as

$$\rho_l = U \Sigma \Sigma^\dagger U^\dagger, \\ \rho_r^T = V \Sigma^\dagger \Sigma V^\dagger.$$

To calculate the half-chain entropy, we have $d_l = d_r = d$ and $\Sigma \Sigma^\dagger = \Sigma^\dagger \Sigma = \Sigma^2$. The diagonal elements Σ_i^2 of the

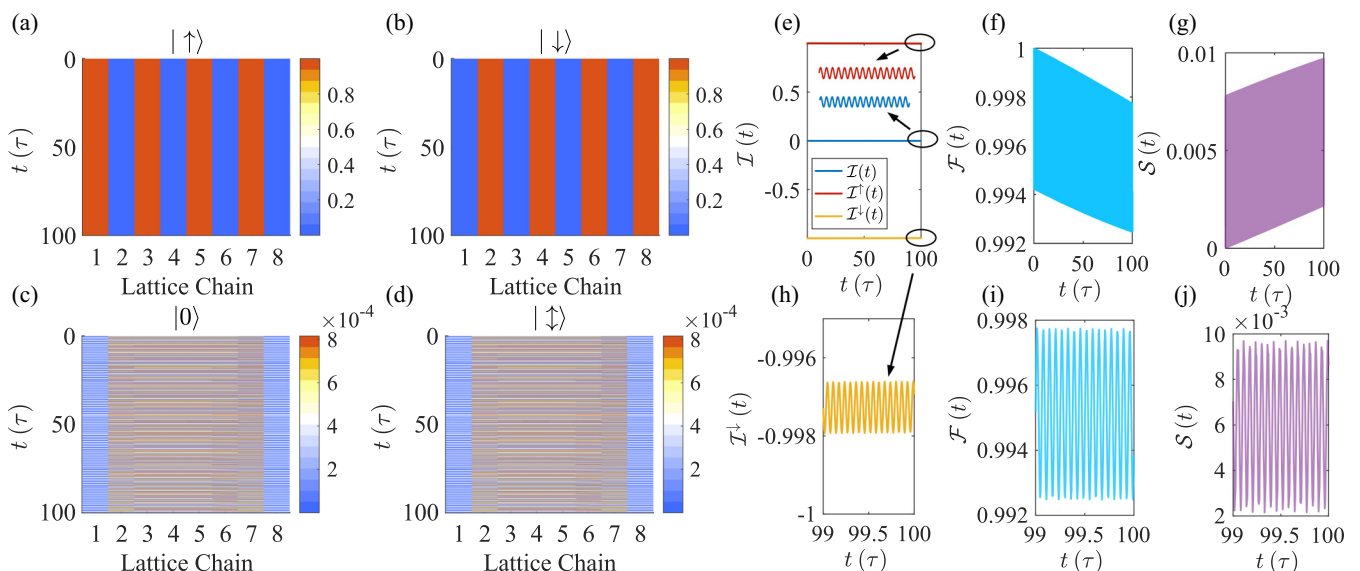


FIG. 8. Near-ideal Wannier-Stark localization. (a)–(d) The probability of observing the quantum many-body state in four distinct states $|\uparrow\rangle$, $|\downarrow\rangle$, $|0\rangle$, and $|\downarrow\rangle$, at each lattice site. (e)–(j) Time evolution of the nonergodic metrics, including imbalance, full-chain fidelity, and half-chain entropy. The initial state is configured as $|\uparrow\downarrow\uparrow\downarrow\uparrow\downarrow\rangle$, with a significantly tilted potential $\Delta = 100J$ and zero on-site Coulomb interaction, $U = 0J$, on an $\mathcal{N} = 8$ lattice.

square matrix Σ^2 correspond exactly to the eigenvalues of the reduced density matrices ρ_l and ρ_r . The Von Neuman entanglement entropy within the half chain can then be determined by

$$S_l = S_r = - \sum_{i=1}^d \Sigma_i^2 \ln \Sigma_i^2.$$

Given that Σ_i^2 can also be interpreted as the square of the singular value Σ_i of the pure quantum many-body state matrix ψ , calculating the entanglement entropy primarily involves determining the singular values of this state matrix, which has the dimensions $d \times d$.

Another key metric for subchains is the fidelity \mathcal{F}_{sub} , defined as the overlap between the reduced density matrices ρ_l at the initial time $\rho_l(0)$ and at a later time $\rho_l(t)$. Mathematically, \mathcal{F}_{sub} is expressed as

$$\mathcal{F}_{\text{sub}}(t) = (\text{Tr}[\sqrt{\sqrt{\rho_l(t)}\rho_l(0)\sqrt{\rho_l(t)}}])^2, \quad (\text{A14})$$

where the square root of ρ_l is computed as

$$\sqrt{\rho_l} = U \Sigma U^\dagger, \quad (\text{A15})$$

where the following has been used:

$$\sqrt{\rho_l}\sqrt{\rho_l} = U \Sigma U^\dagger U \Sigma U^\dagger = U \Sigma^2 U^\dagger = \rho_l.$$

The fidelity metric $\mathcal{F}_{\text{sub}}(t)$ thus quantifies how much the subchain quantum state at time t retains the characteristics of the subchain state at time $t = 0$, according to the evolution of its reduced density matrix.

6. Near-ideal Wannier-Stark localization

Ideal Wannier-Stark localization is typically characterized by the freezing and localized patterns of particles within the spatial lattice space, often due to the influence of a large, tilted electric potential. Figures 8(a)–8(d) illustrate successful particle freezing around the initial state over a long time horizon: $\mathcal{T} = 100\tau$, under a significantly tilted potential $\Delta = 100J$. The subsequent time-evolved series of nonergodic metrics closely approximates the ideal Wannier-Stark case: $\mathcal{I}(t) = \mathcal{I}(0)$, $\mathcal{F}(t) = \mathcal{F}(0) = 1$, and $\mathcal{S}(t) = \mathcal{S}(0) = 0$, in spite of the noticeable decay due to the finite tilt of the potential in numerical computations. Effectively, Fig. 8 displays near-ideal Wannier-Stark localization. Note that Bloch oscillations are still observable in the zoom-in region of Figs. 8(e), 8(h)–8(j).

APPENDIX B: DEEP REINFORCEMENT LEARNING ALGORITHMS

1. PPO pseudoalgorithm

PPO is a standard, policy-gradient based RL method. It works by alternating between gathering data through interactions with the environment and optimizing a surrogate objective function that has been clipped. PPO aims to balance exploration and exploitation by restricting the magnitude of policy updates. This feature contributes to PPO's robustness and efficiency, making it well-suited for a wide range of learning tasks. Here we explain the steps and the mathematics behind of the PPO algorithm.

Pseudoalgorithm: PPO agent

Input: Initial actor $\pi(a|s; \theta)$, critic $V(s; \phi)$, clipping factor ϵ , policy learning rate α_θ , value function learning rate α_ϕ , number of episodes N_{ep} , number of epochs K , and number of minibatches M .
Output: Optimized policy parameters θ .

1: **for** $Episode = 1$ to N_{ep} **do**:

a: Collect trajectory \mathcal{D} with policy π_θ .

b: Compute advantage estimates \hat{A}_t with V_ϕ .

c: Compute return \hat{G}_t .

d: Update policy π_θ by stochastic gradient ascent and value function V_ϕ by stochastic gradient descent:

for $epoch = 1$ to K **do**:

i: Divide \mathcal{D} into M mini-batches.

ii: **for** each $miniBatch$ in \mathcal{D} **do**:

- Compute probability ratio $r_t(\theta)$.

- Compute objective $L^{CLIP}(\theta)$.

- Compute square-error loss $L^{VF}(\phi)$.

- Update $\theta \leftarrow \theta + \alpha_\theta \nabla_\theta L^{CLIP}(\theta)$.

- Update $\phi \leftarrow \phi - \alpha_\phi \nabla_\phi L^{VF}(\phi)$.

2: **return** θ .

Generalized advantage estimation. A trajectory is denoted by \mathcal{D} :

$$\mathcal{D} = (s_0, a_0, R_0, s_1, \dots, s_{\mathcal{T}-1}, a_{\mathcal{T}-1}, R_{\mathcal{T}-1}), \quad (\text{B1})$$

which consists of tuples (state s_t , action a_t , reward R_t). We employ generalized advantage estimation [90], which leverages a value function estimator to calculate the advantage estimates, \hat{A}_t , for each time step within a trajectory. Specifically, the advantage estimate at time t is given by

$$\hat{A}_t = \delta_t + (\gamma\lambda)\delta_{t+1} + \dots + (\gamma\lambda)^{\mathcal{T}-t-1}\delta_{\mathcal{T}-1}, \quad (\text{B2})$$

where the temporal difference error, δ_t , is defined as

$$\delta_t = R_t + \gamma V(s_{t+1}; \phi) - V(s_t; \phi) \quad (\text{B3})$$

and δ_t signifies the immediate advantage of selecting an action under the policy $\pi(a_t|s_t; \theta)$. The stochastic policy $\pi(a_t|s_t; \theta)$ represents the conditional probability distribution over the action space a_t given the state s_t . The value function $V(s_t; \phi)$ is employed to evaluate the quality of state s_t based on the cumulative reward received. The discount factor $\gamma \in (0, 1)$ (typical value $\gamma = 0.997$) and the hyperparameter λ (commonly $\lambda = 0.95$) modulate the weighting of future rewards. In essence, the generalized advantage \hat{A}_t at time t aggregates discounted future advantages up to the terminal stage $\mathcal{T} - 1$, enabling more stable and efficient policy updates.

Return. The return $\hat{G}(\mathcal{D})$ is defined as the cumulative reward collected throughout a trajectory \mathcal{D} , represented by $\hat{G}(\mathcal{D}) = \sum_{t=0}^{\mathcal{T}-1} \gamma^t R_t$ with \mathcal{T} denoting the time horizon. For ease of mathematical treatment, a discounted version is often employed, termed the finite-horizon discounted return:

$$\hat{G}(\mathcal{D}) = \sum_{t=0}^{\mathcal{T}-1} \gamma^t R_t.$$

This formulation acknowledges the contribution of future rewards while assigning them diminishing importance relative to more immediate rewards. The return at each individual time step, \hat{G}_t , is calculated as the sum of the rewards from the

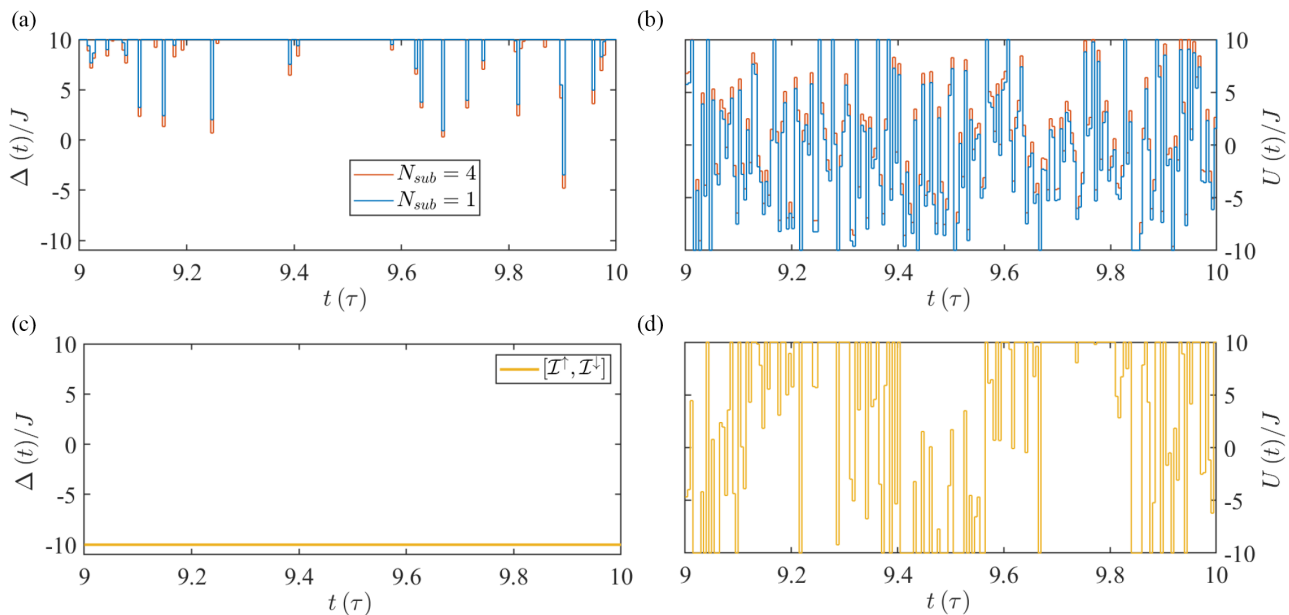


FIG. 9. A magnified view of the optimized nonergodic control flow from 9τ to 10τ shown in Figs. 2(e) and 2(f). (a) and (b) display the optimized control flow when the observation is the fidelity with respect to $N_{\text{sub}} = 4$ and $N_{\text{sub}} = 1$, respectively. (c) and (d) present the optimized control flow when the observation is the imbalance $[I^\uparrow, I^\downarrow]$.

current time step t onwards, discounted by γ to reflect the time value of the rewards:

$$\hat{G}_t = \sum_{k=t}^{\mathcal{T}-1} \gamma^{k-t} R_k.$$

Within the proximal policy optimization framework, this return can be derived from the generalized advantage estimate by

$$\hat{G}_t = \hat{A}_t + V(s_t; \phi), \quad (\text{B4})$$

where \hat{A}_t represents the advantage estimate at time t and $V(s_t; \phi)$ is the value function's estimate of the state's value.

Square-error loss and clipped surrogate objective function. The square-error loss, denoted by $L^{VF}(\phi)$, measures how far the value function's predictions ($\hat{V}(s_t; \phi)$) are from the actual returns (\hat{G}_t) received, which is given by

$$L^{VF}(\phi) = \hat{\mathbb{E}}_t[(\hat{V}(s_t; \phi) - \hat{G}_t)^2], \quad (\text{B5})$$

where $\hat{\mathbb{E}}_t[\]$ is the empirical average over a minibatch of data. The clipped surrogate objective function in PPO can be expressed as

$$L^{CLIP}(\theta) = \hat{\mathbb{E}}_t[\min(r_t(\theta)\hat{A}_t, \text{clip}(r_t(\theta), [1 - \epsilon, 1 + \epsilon])\hat{A}_t)], \quad (\text{B6})$$

which ensure that updates to the policy (how the agent decides to act) are not too drastic. This is accomplished by using a “clip” mechanism that keeps the ratio of the new policy to the old policy ($r_t(\theta)$) within a certain range, defined as

$$r_t(\theta) = \frac{\pi_\theta(a_t|s_t)}{\pi_{\theta_{\text{old}}}(a_t|s_t)}. \quad (\text{B7})$$

If the new policy is exactly the same as the old one, the ratio is one; Otherwise the ratio will deviate from one. The clipping keeps this ratio from going beyond the specified range,

$[1 - \epsilon, 1 + \epsilon]$, which helps slow down policy updates and makes learning more stable.

APPENDIX C: OPTIMIZED NONERGODIC CONTROL FLOW

The appearance of whether the optimal control flow is “wild” depends on the time scale. In our study, the time step size dt is constant, so the control flow is optimized as a piecewise constant control signal, as illustrated in Fig. 9. A recent work [91] and a previous related theoretical study [92] demonstrated that generating square waves on a 1 ns time scale is experimentally feasible in superconducting platforms, which are then viable candidates for simulating the 1D tilted Fermi-Hubbard model [58–60]. In our work, the time step is set to $dt = 0.005\tau$, where τ is determined by the nearest-neighbor coupling J as $\tau = \hbar/J$. Consequently, a 1 ns time step corresponds to coupling strength J on the order of 1 MHz, which is realistically achievable in superconducting qubits [93,94]. A reward function similar to the one used in our work was previously applied to the single-particle quantum control in systems with an upside-down potential [95] or the nonlinear double-well potential [96], demonstrating the versatility and effectiveness of our optimized reward function. We have also tested incorporating the control field into the reward function, but this resulted in significantly worse performance.

A fast Fourier transformation analysis indicates that the solutions for $\Delta(t)$ and $U(t)$ from Figs. 2(e) and 2(f) do not have any dominant Fourier mode but instead have a broad spectrum of frequencies. This complex frequency structure is consistent with the requirement of optimal nonergodic control, which involves intricate manipulation of the system dynamics. The DRL-aligned protocol described in Sec. IV for physical interpretation and insights was developed based on this observation.

- [1] P. W. Anderson, Absence of diffusion in certain random lattices, *Phys. Rev.* **109**, 1492 (1958).
- [2] D. Emin and C. F. Hart, Existence of Wannier-Stark localization, *Phys. Rev. B* **36**, 7353 (1987).
- [3] D. A. Abanin, E. Altman, I. Bloch, and M. Serbyn, Colloquium: Many-body localization, thermalization, and entanglement, *Rev. Mod. Phys.* **91**, 021001 (2019).
- [4] D. D. Vu, K. Huang, X. Li, and S. Das Sarma, Fermionic many-body localization for random and quasiperiodic systems in the presence of short- and long-range interactions, *Phys. Rev. Lett.* **128**, 146601 (2022).
- [5] T. Kohlert, S. Scherg, X. Li, H. P. Lüschen, S. Das Sarma, I. Bloch, and M. Aidelsburger, Observation of many-body localization in a one-dimensional system with a single-particle mobility edge, *Phys. Rev. Lett.* **122**, 170403 (2019).
- [6] D. S. Bhakuni, S. Dattagupta, and A. Sharma, Effect of noise on Bloch oscillations and Wannier-Stark localization, *Phys. Rev. B* **99**, 155149 (2019).
- [7] M. Serbyn, D. A. Abanin, and Z. Papić, Quantum many-body scars and weak breaking of ergodicity, *Nat. Phys.* **17**, 675 (2021).
- [8] R. Nandkishore and D. A. Huse, Many-body localization and thermalization in quantum statistical mechanics, *Annu. Rev. Condens. Matter Phys.* **6**, 15 (2015).
- [9] S. Parameswaran and R. Vasseur, Many-body localization, symmetry and topology, *Rep. Prog. Phys.* **81**, 082501 (2018).
- [10] D. Adler, D. Wei, M. Will, K. Srakaew, S. Agrawal, P. Weckesser, R. Moessner, F. Pollmann, I. Bloch, and J. Zeiher, Observation of Hilbert space fragmentation and fractonic excitations in 2D, *Nature* **636**, 80 (2024).
- [11] T. Kohlert, S. Scherg, P. Sala, F. Pollmann, B. Hebbe Madhusudhana, I. Bloch, and M. Aidelsburger, Exploring the regime of fragmentation in strongly tilted Fermi-Hubbard chains, *Phys. Rev. Lett.* **130**, 010201 (2023).
- [12] J.-Y. Desaules, A. Hudomal, C. J. Turner, and Z. Papić, Proposal for realizing quantum scars in the tilted 1D Fermi-Hubbard model, *Phys. Rev. Lett.* **126**, 210601 (2021).
- [13] W.-J. Huang, Y.-B. Wu, G.-C. Guo, W.-M. Liu, and X.-B. Zou, Strongly-tilted field induced Hamiltonian dimerization and nested quantum scars in the 1D spinless Fermi-Hubbard model, [arXiv:2402.18828](https://arxiv.org/abs/2402.18828).
- [14] S. Scherg, T. Kohlert, P. Sala, F. Pollmann, B. Hebbe Madhusudhana, I. Bloch, and M. Aidelsburger, Observing non-ergodicity due to kinetic constraints in tilted Fermi-Hubbard chains, *Nat. Commun.* **12**, 4490 (2021).
- [15] A. Hudomal, J.-Y. Desaules, B. Mukherjee, G.-X. Su, J. C. Halimeh, and Z. Papić, Driving quantum many-body scars in the PXP model, *Phys. Rev. B* **106**, 104302 (2022).
- [16] G.-X. Su, H. Sun, A. Hudomal, J.-Y. Desaules, Z.-Y. Zhou, B. Yang, J. C. Halimeh, Z.-S. Yuan, Z. Papić, and J.-W. Pan, Observation of many-body scarring in a Bose-Hubbard quantum simulator, *Phys. Rev. Res.* **5**, 023010 (2023).
- [17] D. Bluvstein, A. Omran, H. Levine, A. Keesling, G. Semeghini, S. Ebadi, T. T. Wang, A. A. Michailidis, N. Maskara, W. W. Ho *et al.*, Controlling quantum many-body dynamics in driven Rydberg atom arrays, *Science* **371**, 1355 (2021).
- [18] H. Zhao, J. Vovrosh, F. Mintert, and J. Knolle, Quantum many-body scars in optical lattices, *Phys. Rev. Lett.* **124**, 160604 (2020).
- [19] K. Mizuta, K. Takasan, and N. Kawakami, Exact Floquet quantum many-body scars under Rydberg blockade, *Phys. Rev. Res.* **2**, 033284 (2020).
- [20] S. Sugiura, T. Kuwahara, and K. Saito, Many-body scar state intrinsic to periodically driven system, *Phys. Rev. Res.* **3**, L012010 (2021).
- [21] K. Huang and X. Li, Engineering subharmonic responses beyond prethermalization via Floquet scar states, *Phys. Rev. B* **109**, 064306 (2024).
- [22] H. C. Po, L. Fidkowski, T. Morimoto, A. C. Potter, and A. Vishwanath, Chiral Floquet phases of many-body localized bosons, *Phys. Rev. X* **6**, 041070 (2016).
- [23] M. Sonner, M. Serbyn, Z. Papić, and D. A. Abanin, Thouless energy across the many-body localization transition in Floquet systems, *Phys. Rev. B* **104**, L081112 (2021).
- [24] P. Ponte, Z. Papić, F. Huveneers, and D. A. Abanin, Many-body localization in periodically driven systems, *Phys. Rev. Lett.* **114**, 140401 (2015).
- [25] S. D. Geraedts, R. Nandkishore, and N. Regnault, Many-body localization and thermalization: Insights from the entanglement spectrum, *Phys. Rev. B* **93**, 174202 (2016).
- [26] P. Sierant, M. Lewenstein, A. Scardicchio, and J. Zakrzewski, Stability of many-body localization in Floquet systems, *Phys. Rev. B* **107**, 115132 (2023).
- [27] N. Maskara, A. A. Michailidis, W. W. Ho, D. Bluvstein, S. Choi, M. D. Lukin, and M. Serbyn, Discrete time-crystalline order enabled by quantum many-body scars: Entanglement steering via periodic driving, *Phys. Rev. Lett.* **127**, 090602 (2021).
- [28] B. Mukherjee, S. Nandy, A. Sen, D. Sen, and K. Sengupta, Collapse and revival of quantum many-body scars via Floquet engineering, *Phys. Rev. B* **101**, 245107 (2020).
- [29] A. A. Diringer and T. Gulden, Impact of drive harmonics on the stability of Floquet many-body localization, *Phys. Rev. B* **103**, 214204 (2021).
- [30] P.-G. Rozon, M. J. Gullans, and K. Agarwal, Constructing quantum many-body scar Hamiltonians from Floquet automata, *Phys. Rev. B* **106**, 184304 (2022).
- [31] A. Chandran, T. Iadecola, V. Khemani, and R. Moessner, Quantum many-body scars: A quasiparticle perspective, *Annu. Rev. Condens. Matter Phys.* **14**, 443 (2023).
- [32] A. Haldar, D. Sen, R. Moessner, and A. Das, Dynamical freezing and scar points in strongly driven Floquet matter: Resonance vs emergent conservation laws, *Phys. Rev. X* **11**, 021008 (2021).
- [33] S. Pai and M. Pretko, Dynamical scar states in driven fracton systems, *Phys. Rev. Lett.* **123**, 136401 (2019).
- [34] M. Demirplak and S. A. Rice, Adiabatic population transfer with control fields, *J. Phys. Chem. A* **107**, 9937 (2003).
- [35] F. Petiziol, B. Dive, F. Mintert, and S. Wimberger, Fast adiabatic evolution by oscillating initial Hamiltonians, *Phys. Rev. A* **98**, 043436 (2018).
- [36] D. Sels and A. Polkovnikov, Minimizing irreversible losses in quantum systems by local counterdiabatic driving, *Proc. Natl. Acad. Sci.* **114**, E3909 (2017).
- [37] M. Ljubotina, B. Roos, D. A. Abanin, and M. Serbyn, Optimal steering of matrix product states and quantum many-body scars, *PRX Quantum* **3**, 030343 (2022).
- [38] M. Ljubotina, E. Petrova, N. Schuch, and M. Serbyn, Tangent space generators of matrix product states and exact Floquet quantum scars, *PRX Quantum* **5**, 040311 (2024).

- [39] T. Caneva, T. Calarco, and S. Montangero, Chopped random-basis quantum optimization, *Phys. Rev. A* **84**, 022326 (2011).
- [40] P. Doria, T. Calarco, and S. Montangero, Optimal control technique for many-body quantum dynamics, *Phys. Rev. Lett.* **106**, 190501 (2011).
- [41] M. M. Müller, R. S. Said, F. Jelezko, T. Calarco, and S. Montangero, One decade of quantum optimal control in the chopped random basis, *Rep. Prog. Phys.* **85**, 076001 (2022).
- [42] J. H. M. Jensen, F. S. Møller, J. J. Sørensen, and J. F. Sherson, Approximate dynamics leading to more optimal control: Efficient exact derivatives, *Phys. Rev. A* **103**, 062612 (2021).
- [43] J. H. M. Jensen, F. S. Møller, J. J. Sørensen, and J. F. Sherson, Achieving fast high-fidelity optimal control of many-body quantum dynamics, *Phys. Rev. A* **104**, 052210 (2021).
- [44] S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, Comparing, optimizing, and benchmarking quantum-control algorithms in a unifying programming framework, *Phys. Rev. A* **84**, 022305 (2011).
- [45] R. R. Agundez, C. D. Hill, L. C. L. Hollenberg, S. Rogge, and M. Blaauboer, Superadiabatic quantum state transfer in spin chains, *Phys. Rev. A* **95**, 012317 (2017).
- [46] 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).
- [47] F. Metz and M. Bukov, Self-correcting quantum many-body control using reinforcement learning with tensor networks, *Nat. Mach. Intell.* **5**, 780 (2023).
- [48] L.-L. Ye, C. Arenz, J. M. Lukens, and Y.-C. Lai, Entanglement engineering of optomechanical systems by reinforcement learning, *APL Mach. Learn.* **3**, 1 (2025).
- [49] L. P. Kaelbling, M. L. Littman, and A. W. Moore, Reinforcement learning: A survey, *J. Artif. Intell. Res.* **4**, 237 (1996).
- [50] H.-N. Wang, N. Liu, Y.-Y. Zhang, D.-W. Feng, F. Huang, D.-S. Li, and Y.-M. Zhang, Deep reinforcement learning: A survey, *Front. Inf. Technol. Electron. Eng.* **21**, 1726 (2020).
- [51] J. Ramírez, W. Yu, and A. Perrusquía, Model-free reinforcement learning from expert demonstrations: A survey, *Artif. Intell. Rev.* **55**, 3213 (2022).
- [52] L. Kaiser, M. Babaeizadeh, P. Milos, B. Osinski, R. H. Campbell, K. Czechowski, D. Erhan, C. Finn, P. Kozakowski, S. Levine *et al.*, Model-based reinforcement learning for atari, [arXiv:1903.00374](https://arxiv.org/abs/1903.00374).
- [53] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, Proximal policy optimization algorithms, [arXiv:1707.06347](https://arxiv.org/abs/1707.06347).
- [54] Y. Wang, H. He, and X. Tan, Truly proximal policy optimization, in *Proceedings of the 35th Uncertainty in Artificial Intelligence Conference*, edited by R. P. Adams and V. Gogate, Proc. Mach. Learn. Res. (PMLR, Tempe, AZ, 2020), Vol. 115, pp. 113–122.
- [55] Y. Wu, E. Mansimov, R. B. Grosse, S. Liao, and J. Ba, Scalable trust-region method for deep reinforcement learning using Kronecker-factored approximation, in *Advances in Neural Information Processing Systems*, edited by I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Curran Associates, Inc., 2017), Vol. 30.
- [56] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, Human-level control through deep reinforcement learning, *Nature (London)* **518**, 529 (2015).
- [57] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, Continuous control with deep reinforcement learning, [arXiv:1509.02971](https://arxiv.org/abs/1509.02971).
- [58] X.-Y. Guo, Z.-Y. Ge, H. Li, Z. Wang, Y.-R. Zhang, P. Song, Z. Xiang, X. Song, Y. Jin, L. Lu *et al.*, Observation of Bloch oscillations and Wannier-Stark localization on a superconducting quantum processor, *npj Quantum Inf.* **7**, 51 (2021).
- [59] J.-M. Reiner, M. Marthaler, J. Braumüller, M. Weides, and G. Schön, Emulating the one-dimensional Fermi-Hubbard model by a double chain of qubits, *Phys. Rev. A* **94**, 032338 (2016).
- [60] P. Zhang, H. Dong, Y. Gao, L. Zhao, J. Hao, J.-Y. Desaulles, Q. Guo, J. Chen, J. Deng, B. Liu *et al.*, Many-body Hilbert space scarring on a superconducting processor, *Nat. Phys.* **19**, 120 (2023).
- [61] E. Guardado-Sanchez, A. Morningstar, B. M. Spar, P. T. Brown, D. A. Huse, and W. S. Bakr, Subdiffusion and heat transport in a tilted two-dimensional Fermi-Hubbard system, *Phys. Rev. X* **10**, 011042 (2020).
- [62] P. Zechmann, A. Bastianello, and M. Knap, Tunable transport in the mass-imbalanced Fermi-Hubbard model, *Phys. Rev. B* **106**, 075115 (2022).
- [63] P. Zhang, Subdiffusion in strongly tilted lattice systems, *Phys. Rev. Res.* **2**, 033129 (2020).
- [64] N. Darkwah Oppong, G. Pasqualetti, O. Bettermann, P. Zechmann, M. Knap, I. Bloch, and S. Fölling, Probing transport and slow relaxation in the mass-imbalanced Fermi-Hubbard model, *Phys. Rev. X* **12**, 031026 (2022).
- [65] T. G. Kiely and E. J. Mueller, High-temperature transport in the one-dimensional mass-imbalanced Fermi-Hubbard model, *Phys. Rev. A* **109**, 063318 (2024).
- [66] J. Boesl, P. Zechmann, J. Feldmeier, and M. Knap, Deconfinement dynamics of fractons in tilted Bose-Hubbard chains, *Phys. Rev. Lett.* **132**, 143401 (2024).
- [67] E. Lake and T. Senthil, Non-fermi liquids from kinetic constraints in tilted optical lattices, *Phys. Rev. Lett.* **131**, 043403 (2023).
- [68] W. Morong, F. Liu, P. Becker, K. Collins, L. Feng, A. Kyprianidis, G. Pagano, T. You, A. Gorshkov, and C. Monroe, Observation of Stark many-body localization without disorder, *Nature (London)* **599**, 393 (2021).
- [69] Q. Guo, C. Cheng, H. Li, S. Xu, P. Zhang, Z. Wang, C. Song, W. Liu, W. Ren, H. Dong *et al.*, Stark many-body localization on a superconducting quantum processor, *Phys. Rev. Lett.* **127**, 240502 (2021).
- [70] M. Schreiber, S. S. Hodgman, P. Bordia, H. P. Lüschen, M. H. Fischer, R. Vosk, E. Altman, U. Schneider, and I. Bloch, Observation of many-body localization of interacting fermions in a quasirandom optical lattice, *Science* **349**, 842 (2015).
- [71] H. Gimpferlein, S. Wessel, J. Schmiedmayer, and L. Santos, Ultracold atoms in optical lattices with random on-site interactions, *Phys. Rev. Lett.* **95**, 170401 (2005).
- [72] S. Greschner, L. Santos, and D. Poletti, Exploring unconventional Hubbard models with doubly modulated lattice gases, *Phys. Rev. Lett.* **113**, 183002 (2014).
- [73] P. Sierant and J. Zakrzewski, Many-body localization of Bosons in optical lattices, *New J. Phys.* **20**, 043032 (2018).

- [74] H. F. Trotter, On the product of semigroups of operators, *Proc. Am. Math. Soc.* **10**, 545 (1959).
- [75] M. Suzuki, Generalized Trotter's formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems, *Commun. Math. Phys.* **51**, 183 (1976).
- [76] L.-L. Ye and Y.-C. Lai, Irregular Bloch-Zener oscillations in two-dimensional flat-band dirac materials, *Phys. Rev. B* **107**, 165422 (2023).
- [77] J. Šuntajs, T. Prosen, and L. Vidmar, Localization challenges quantum chaos in the finite two-dimensional Anderson model, *Phys. Rev. B* **107**, 064205 (2023).
- [78] C. Thompson, G. Vemuri, and G. S. Agarwal, Anderson localization with second quantized fields in a coupled array of waveguides, *Phys. Rev. A* **82**, 053805 (2010).
- [79] Y. Bar Lev, D. R. Reichman, and Y. Sagi, Many-body localization in system with a completely delocalized single-particle spectrum, *Phys. Rev. B* **94**, 201116(R) (2016).
- [80] P. Sierant, D. Delande, and J. Zakrzewski, Many-body localization due to random interactions, *Phys. Rev. A* **95**, 021601(R) (2017).
- [81] M. Schulz, C. A. Hooley, R. Moessner, and F. Pollmann, Stark many-body localization, *Phys. Rev. Lett.* **122**, 040606 (2019).
- [82] E. van Nieuwenburg, Y. Baum, and G. Refael, From Bloch oscillations to many-body localization in clean interacting systems, *Proc. Natl. Acad. Sci.* **116**, 9269 (2019).
- [83] J. M. Deutsch, Quantum statistical mechanics in a closed system, *Phys. Rev. A* **43**, 2046 (1991).
- [84] M. Srednicki, Chaos and quantum thermalization, *Phys. Rev. E* **50**, 888 (1994).
- [85] M. Rigol, V. Dunjko, and M. Olshanii, Thermalization and its mechanism for generic isolated quantum systems, *Nature (London)* **452**, 854 (2008).
- [86] S. Bravyi, D. P. DiVincenzo, and D. Loss, Schrieffer-Wolff transformation for quantum many-body systems, *Ann. Phys.* **326**, 2793 (2011).
- [87] Y. Lu, P. Shi, X.-H. Wang, J. Hu, and S.-J. Ran, Persistent ballistic entanglement spreading with optimal control in quantum spin chains, *Phys. Rev. Lett.* **133**, 070402 (2024).
- [88] <https://github.com/liliyequantum/nonergodicityRL>.
- [89] C. Yi and E. Crosson, Spectral analysis of product formulas for quantum simulation, *npj Quantum Inf.* **8**, 37 (2022).
- [90] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, High-dimensional continuous control using generalized advantage estimation, [arXiv:1506.02438](https://arxiv.org/abs/1506.02438).
- [91] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, Leakage reduction in fast superconducting qubit gates via optimal control, *npj Quant. Inf.* **7**, 14 (2021).
- [92] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple pulses for elimination of leakage in weakly nonlinear qubits, *Phys. Rev. Lett.* **103**, 110501 (2009).
- [93] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell *et al.*, Superconducting quantum circuits at the surface code threshold for fault tolerance, *Nature (London)* **508**, 500 (2014).
- [94] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro *et al.*, Qubit architecture with high coherence and fast tunable coupling, *Phys. Rev. Lett.* **113**, 220502 (2014).
- [95] Z. T. Wang, Y. Ashida, and M. Ueda, Deep reinforcement learning control of quantum cartpoles, *Phys. Rev. Lett.* **125**, 100401 (2020).
- [96] S. Borah, B. Sarma, M. Kewming, G. J. Milburn, and J. Twamley, Measurement-based feedback quantum control with deep reinforcement learning for a double-well nonlinear potential, *Phys. Rev. Lett.* **127**, 190403 (2021).