

Reinforcement Learning for Charging Optimization of Inhomogeneous Dicke Quantum Batteries

Charging optimization is a key challenge to the implementation of quantum batteries, particularly under inhomogeneity and partial observability. This paper employs reinforcement learning to optimize piecewise-constant charging policies for an inhomogeneous Dicke battery. We systematically compare policies across four observability regimes, from full-state access to experimentally accessible observables (energies of individual two-level systems (TLSs), first-order averages, and second-order correlations). Simulation results demonstrate that full observability yields near-optimal ergotropy with low variability, while under partial observability, access to only single-TLS energies or energies plus first-order averages lags behind the fully observed baseline. However, augmenting partial observations with second-order correlations recovers most of the gap, reaching 94%–98% of the full-state baseline. The learned schedules are nonmyopic, trading temporary plateaus or declines for superior terminal outcomes. These findings highlight a practical route to effective fast-charging protocols under realistic information constraints.

Keywords: Inhomogeneous dicke quantum batteries; Charging optimization; Reinforcement learning; Partial observability

I. INTRODUCTION

The increasing demand for efficient and rapid energy storage solutions at the nanoscale has spurred significant interest in the development of quantum technologies for energy management. Quantum batteries, devices that store energy in quantum mechanical degrees of freedom, have emerged as a promising paradigm, offering the potential for faster charging time and higher energy densities compared to their classical counterparts [1–4]. The key advantage of quantum batteries lies in their ability to exploit quantum phenomena, such as entanglement and coherence, to achieve a quantum advantage in charging power, a phenomenon known as supercharging [5, 6].

The central challenge in harnessing the potential of quantum batteries lies in designing effective charging protocols [7, 8]. The charging process is fundamentally a problem of quantum control: how to manipulate external parameters, such as driving fields, to steer the battery from a low-energy (discharged) state to a high-energy (charged) state in the shortest possible time while maximizing the stored energy and minimizing energy dissipation [9].

Traditionally, this problem has been tackled using the tools of optimal control theory [10]. These methods seek to find a time-dependent control Hamiltonian that optimizes a specific objective function, such as the charging time or the final energy of the battery. Seminal works in this area have provided analytical solutions for simple quantum battery models, revealing fundamental limits on charging power and highlighting the role of quantum correlations in achieving charging speed-ups [11, 12]. However, these approaches are often limited to idealized models with a small number of components and simple interactions. In addi-

tion to classical optimal control, recent work in the broader control community has explored learning-based approximations of optimization-based controllers, primarily to reduce online computational overhead [13–15]. In particular, imitation learning has been used to distill expert controllers such as stochastic model predictive control into neural surrogates that can be evaluated efficiently at deployment. Such approaches, however, rely on access to an expert optimizer and on generating representative expert trajectories. For quantum-battery charging, the dynamics are governed by high-dimensional Hamiltonian evolution, and obtaining expert solutions via solver-in-the-loop optimal control can become computationally prohibitive as the dimension grows.

To address these limitations, researchers have begun to explore the use of machine learning, and specifically reinforcement learning (RL) [16–18], as a powerful alternative for discovering optimal charging protocols [19–21]. RL offers a model-free approach to control problems, where an agent learns an optimal policy by interacting with its environment and receiving feedback in the form of rewards. This data-driven approach is well-suited for the complexities of quantum systems, where it has been successfully applied to various tasks such as quantum state preparation and gate design [22–25]. It can potentially discover non-intuitive control strategies that are missed by conventional methods.

Recent work [19] applied reinforcement learning to optimize fast-charging protocols for Dicke quantum batteries with strong performance. However, the Dicke model considered there is homogeneous, and the agent has access to the full quantum state. These are idealizations that are rarely attainable in practice. Crucially, small imperfections in physical implementations would inevitably break the

uniformity of the atomic frequencies as exemplified by Doppler shifts from atomic motion in cavity QED systems [26, 27]. Consequently, real devices exhibit inhomogeneous parameters, and only limited, experimentally accessible observables are available for control.

Motivated by these gaps, this work explores reinforcement learning as a data-driven route to fast charging under such realistic conditions. We focus on an inhomogeneous Dicke quantum battery and cast charging as a discrete-time RL problem in which an agent learns a piecewise-constant control sequence for the charger–battery coupling within a bounded, continuous action space. The agent is guided by a dense, shaped reward that blends energy and ergotropy gains while the overall objective targets terminal ergotropy. To mirror experimental readout limits, we consider four observability regimes ranging from full-state access to accessible summaries built from single-TLS energies, first-order collective averages, and second-order correlation features. Policies are trained with Soft Actor–Critic [18, 28], leveraging entropy regularization and replay-based off-policy updates. Together, these elements lay the groundwork for intelligent charging protocols that remain effective under partial observability and realistic readout constraints.

The rest of this work is organized as follows. In Section II we introduce the Dicke quantum battery model. The reinforcement learning framework for charging optimization is stated in Section III. The performance of the framework is evaluated and discussed in Section IV, followed by conclusions in Section V.

II. DICKE QUANTUM BATTERY AND CHARGING DYNAMICS

A paradigmatic model for exploring collective effects in quantum batteries is the Dicke model[29]. In this framework, the quantum battery consists of N two-level systems (TLSs), or qubits, which act as the battery cells. These cells are collectively coupled to a single bosonic mode, typically a cavity field, which serves as the charger. The collective nature of the interaction is the key ingredient that allows for the supercharging effect[12].

The Hamiltonian of the total system (battery plus charger) in the standard uniform Dicke model is given by $\hat{H}(t) = \hat{H}_B + \hat{H}_C + \hat{H}_{\text{int}}(t)$. Here, \hat{H}_B describes an ensemble of N identical battery cells, each modeled as a TLS with ground state $|g\rangle$, excited state $|e\rangle$, and an energy splitting ω_0 . Specifically, the battery Hamiltonian takes the form (in

this paper we take $\hbar = 1$)

$$\hat{H}_B = \omega_0(\hat{J}_z + J), \quad (1)$$

where $\hat{J}_z = \frac{1}{2} \sum_{j=1}^N \hat{\sigma}_z^{(j)}$ is the collective pseudo-spin operator, with $\hat{\sigma}_z^{(j)}$ denoting the Pauli- Z operator for the j -th cell. The total spin is represented by $J = N/2$, and the constant J in Eq. (1) is introduced to ensure that the average energy remains positive, simplifying subsequent analysis. The charger subsystem is represented by a single-mode cavity with frequency ω_c , and its Hamiltonian is given by

$$\hat{H}_C = \omega_c \hat{a}^\dagger \hat{a}, \quad (2)$$

where \hat{a}^\dagger and \hat{a} denote the bosonic creation and annihilation operators of the cavity mode, respectively. Furthermore, we consider the resonant condition in which the cavity frequency matches the energy splitting of the TLSs, namely $\omega_c = \omega_0$. The energy exchange between the battery and the charger is governed by the interaction Hamiltonian \hat{H}_{int} . In the ideal Dicke model, all battery cells couple identically to the cavity mode, leading to the interaction Hamiltonian

$$\hat{H}_{\text{int}}(t) = 2\lambda_c(t)\hat{J}_x(\hat{a} + \hat{a}^\dagger), \quad (3)$$

where $\hat{J}_x = \frac{1}{2} \sum_{j=1}^N \hat{\sigma}_x^{(j)}$ with $\hat{\sigma}_x^{(j)}$ being the Pauli- X operator for the j -th cell, and $\lambda_c(t)$ denotes the uniform coupling strength, which serves as a classical external control parameter. It is noteworthy that we do not employ the rotating-wave approximation (RWA) here. This is because the RWA is typically valid only under weak coupling conditions, which may not hold for the arbitrary values of the control parameter $\lambda_c(t)$ considered in our study.

The dynamics of the ideal Dicke model is confined to an $(N+1)$ -dimensional subspace due to the conservation of the total spin operator $\hat{\mathbf{J}}^2$, a consequence of the system’s global $\text{SU}(2)$ symmetry, which is widely used in the study of Dicke quantum batteries [19]. However, the assumption of identical cells represents an idealization. In realistic physical implementations, imperfections and manufacturing variations lead to non-uniformity in the properties of the individual TLSs. To capture this effect, we extend the Dicke model by allowing each cell to have a distinct energy splitting ω_j . This non-uniformity breaks the collective symmetry of the standard Dicke model. The modified

battery and interaction Hamiltonians become:

$$\hat{H}'_B = \sum_{j=1}^N \frac{\omega_j}{2} (\hat{\sigma}_z^{(j)} + 1), \quad (4a)$$

$$\hat{H}'_{\text{int}}(t) = \lambda_c(t) \sum_{j=1}^N \hat{\sigma}_x^{(j)} (\hat{a} + \hat{a}^\dagger). \quad (4b)$$

Here, $\omega_j = \omega_0 + \delta_j$ denotes the transition frequency of the j -th TLS, where δ_j represents the disorder strength or deviation from the nominal frequency ω_0 . The total Hamiltonian of the non-uniform system is given by $\hat{H}'(t) = \hat{H}'_B + \hat{H}_C + \hat{H}'_{\text{int}}(t)$. The presence of disorder in both on-site energies and couplings renders the system's dynamics significantly more complex and generally intractable to analytical optimal control techniques. This complexity motivates the adoption of adaptive, model-free approaches such as reinforcement learning.

The charging process is initiated with the system in a well-defined initial state. The battery is assumed to be fully discharged, meaning all TLS are in their ground state. For the non-uniform model, this corresponds to the product state $|\psi(0)\rangle_B = |g_1, g_2, \dots, g_N\rangle$. The charger is prepared in a state with a large number of excitations, such as a Fock state $|n\rangle$ with $n = N$. The system then evolves according to the time-dependent Schrödinger equation:

$$i \frac{d}{dt} |\psi(t)\rangle = \hat{H}'(t) |\psi(t)\rangle, \quad (5)$$

where the initial state is $|\psi(0)\rangle = |\psi(0)\rangle_B \otimes |n\rangle$.

The key figure of merit is the instantaneous energy $E(t) = \sum_j E_j(t)$ stored in the battery at time t , where $E_j(t) = \langle \psi(t) | \frac{\omega_j}{2} (\hat{\sigma}_z^{(j)} + 1) | \psi(t) \rangle$ is the instantaneous energy of the j -th TLS. Another important quantity characterizing the stored energy is the ergotropy $\mathcal{E}(t)$ at time t , defined as

$$\mathcal{E}(t) = \text{Tr}[\hat{H}'_B \rho(t)] - \text{Tr}[\hat{H}'_B \tilde{\rho}(t)], \quad (6)$$

where $\rho(t) \equiv \text{Tr}_C[|\psi(t)\rangle\langle\psi(t)|]$ is the reduced density matrix of the battery, and $\tilde{\rho}(t)$ is the corresponding passive state. This passive state is given by $\tilde{\rho}(t) = \sum_{i=1}^N \lambda_i |\epsilon_i\rangle\langle\epsilon_i|$, where $\{\lambda_i\}$ are the eigenvalues of $\rho(t)$ sorted in decreasing order ($\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N$), and $\{|\epsilon_i\rangle\}$ are the eigenstates of \hat{H}'_B with energies sorted in increasing order ($\epsilon_1 \leq \epsilon_2 \leq \dots \leq \epsilon_N$). The ergotropy represents the maximum amount of energy that can be extracted from the quantum state via cyclic unitary operations[6]. The objective is to find a control protocol for the coupling strength $\lambda_c(t)$ that maximizes the final ergotropy $\mathcal{E}(\tau)$ at a given time τ .

III. REINFORCEMENT LEARNING FRAMEWORK FOR CHARGING OPTIMIZATION

To determine the optimal time-dependent control function $\lambda_c(t)$, we reformulate the charging process as a discrete-time reinforcement learning problem. An RL agent interacts with a simulated quantum battery system as its environment to acquire a control policy. At each time step k , the agent observes a state s_k and selects an action a_k according to a deterministic or stochastic policy $\pi(a_k|s_k)$. The objective is to maximize the expected return, i.e., the cumulative reward over the charging horizon.

A. Environment for the RL agent

The environment for the RL agent is the simulated or real inhomogeneous Dicke battery itself, composed of N TLSs and a charger. It emulates the quantum dynamics induced by a given charging protocol. The total charging time τ is divided into K discrete steps of duration $\Delta t = \tau/K$. At the beginning of each step k ($k = 0, 1, \dots, K-1$), the agent selects an action a_k , which sets the coupling strength λ_c to a constant value for the duration of that step. The environment then evolves the quantum state from $|\psi(t_k)\rangle$ to $|\psi(t_{k+1})\rangle$ by solving the Schrödinger equation with a piecewise-constant Hamiltonian:

$$|\psi(t_{k+1})\rangle = \exp(-i\hat{H}'(a_k)\Delta t) |\psi(t_k)\rangle \quad (7)$$

where $\hat{H}'(a_k)$ is the total system Hamiltonian with the coupling strength determined by a_k .

The state s_k should provide the agent with sufficient information about the system at time t_k ($t_k = k\Delta t$) and satisfy the Markov property (i.e., s_{k+1} must only depend on s_k and the applied action a_k). While the full quantum state vector $|\psi(t_k)\rangle$ is information-complete, it is generally not accessible in experiments and thus impractical as an RL state representation. To address this limitation, we replace s_k with experimentally accessible observables o_k and evaluate whether such observations suffice to learn high-quality charging protocols. In this paper we consider three families of measurements: the expected energy of each TLS, first-order averages, and second-order correlations. Moreover, as in [19], we also incorporate the last selected action a_{k-1} and the current time step t_k into the policy input for both the full-state and partial-observation settings.

Note that replacing the full state s_k by a restricted observation o_k does not guarantee that the

resulting process is Markovian in the observation space. Experimentally accessible observables are typically not information-complete, so the mapping $o_k = f(|\psi(t_k)\rangle)$ is generally many-to-one. Consequently, distinct quantum states may correspond to the same o_k , and the transition to o_{k+1} may also depend on additional, unobserved physical quantities beyond o_k and a_k , rendering the control problem effectively a partially observable Markov decision process (POMDP) [30]. In this work we treat this setting with a practical baseline by learning a memoryless policy $\pi(a_k | o_k)$ augmented with a_{k-1} and t_k , and we empirically assess the impact of partial observability by comparing a hierarchy of observation sets.

The action a_k specifies the coupling strength λ_c applied during step k . We define a bounded continuous action space:

$$\mathcal{A} = [\lambda_{c,\min}, \lambda_{c,\max}] \subset \mathbb{R} \quad (8)$$

where the bounds $\lambda_{c,\min}$ and $\lambda_{c,\max}$ are hyperparameters informed by physical knowledge. At each step k the agent selects $a_k \in \mathcal{A}$ and holds it fixed for the duration of that step (i.e., $[t_k, t_{k+1})$), thereby realizing a piecewise-constant charging protocol.

The reward function $R(s_k, a_k, s_{k+1})$ is crucial for guiding the agent towards the desired charging performance. To encourage efficient charging, we define the step- k reward as the increment in the battery’s ergotropy:

$$r_k = \mathcal{E}(t_{k+1}) - \mathcal{E}(t_k) \quad (9)$$

where $\mathcal{E}(t_k)$ denotes the ergotropy of the battery at time t_k . This dense, immediate feedback directly rewards actions that increase ergotropy. The agent’s objective is to maximize the undiscounted return $G = \sum_{k=0}^{K-1} r_k$, which telescopes to

$$G = \mathcal{E}(t_K) - \mathcal{E}(t_0) \quad (10)$$

with the final time $t_K = \tau$. Thus maximizing G is equivalent to maximizing the net increase in ergotropy. However, at early training stages the ergotropy-based reward is often zero, making it difficult to guide policy improvement. In practice, we therefore employ a weighted combination of energy gain and ergotropy gain as the reward. Following the weighting strategy in [19], the scheme emphasizes energy gain when the ergotropy increment is typically zero, and progressively increases the weight on ergotropy gain as the policy improves:

$$r_k = w(n_{\text{step}})(E(t_{k+1}) - E(t_k)) + (1 - w(n_{\text{step}}))(\mathcal{E}(t_{k+1}) - \mathcal{E}(t_k)), \quad (11)$$

where the scheduled weight is chosen as

$$w(n_{\text{step}}) = \left(1 + e^{(n_{\text{step}} - w_{\text{mean}})/w_{\text{width}}}\right)^{-1}. \quad (12)$$

Here n_{step} denotes the training step index, and w_{mean} and w_{width} are hyperparameters[19]. This schedule satisfies $w(n_{\text{step}}) \simeq 1$ at early training and $w(n_{\text{step}}) \rightarrow 0$ at late training, thereby shifting the reward from energy increments to ergotropy increments. While the default training uses the scheduled weight $w(n_{\text{step}})$, we assess sensitivity to the reward mixing in Appendix B by comparing several fixed mixing values $w \in [0, 1]$.

B. Soft Actor–Critic algorithm

We train the control policy with Soft Actor–Critic (SAC), an off-policy, entropy-regularized actor–critic method that alternates policy evaluation and policy improvement while maximizing a trade-off between return and action entropy [18]. Concretely, SAC optimizes

$$\arg \max_{\pi} \mathbb{E}_{\pi} \left[\sum_{k=0}^{K-1} \gamma^k (r_k + \alpha \mathcal{H}[\pi(\cdot | s_k)]) \mid s_0 = s_{\text{init}} \right] \quad (13)$$

where $\gamma \in [0, 1]$ is the discount factor and $\alpha \geq 0$ is the temperature controlling the exploration–exploitation. While our task is finite-horizon with an undiscounted terminal objective, we employ a discounted backup ($\gamma < 1$) during training to reduce target variance and stabilize off-policy updates.

a. Actor. The stochastic policy $\pi_{\theta}(a|s)$ outputs an unconstrained Gaussian sample $\tilde{a}(s) \sim \mathcal{N}(\mu_{\theta}(s), \sigma_{\theta}^2(s))$, which is pushed through a tanh map and then linearly rescaled to $[\lambda_{c,\min}, \lambda_{c,\max}]$ (i.e., the squashed Gaussian policy in [19]):

$$a(s) = \frac{\lambda_{c,\max} - \lambda_{c,\min}}{2} \tanh(\tilde{a}(s)) + \frac{\lambda_{c,\max} + \lambda_{c,\min}}{2}. \quad (14)$$

This yields differentiable, bounded actions compatible with the piecewise-constant control protocol of Sec. III A.

b. Critic and policy evaluation. Two state–action value networks Q_{ϕ_1} and Q_{ϕ_2} are learned to reduce over-estimation bias [18]. Given a replay buffer \mathcal{B} of one-step transitions (s, a, r, s') , the critics minimize

$$\mathcal{L}_Q(\phi_i) = \mathbb{E}_{(s,a,r,s') \sim \mathcal{B}} \left[(Q_{\phi_i}(s, a) - y)^2 \right], \quad i \in \{1, 2\}, \quad (15)$$

with a clipped double- Q target

$$y = r + \gamma \mathbb{E}_{a' \sim \pi_{\theta}(\cdot | s')} \left[\min_{j \in \{1, 2\}} Q_{\phi_j^{\text{targ}}}(s', a') - \alpha \log \pi_{\theta}(a' | s') \right], \quad (16)$$

and Polyak-averaged target networks $\phi_j^{\text{targ}} \leftarrow \rho \phi_j^{\text{targ}} + (1 - \rho) \phi_j$ with $\rho \in (0, 1)$ [18, 19].

c. Policy improvement. The actor is updated by minimizing [18, 19]:

$$\mathcal{L}_\pi(\theta) = \mathbb{E}_{s \sim \mathcal{B}, a \sim \pi_\theta(\cdot|s)} \left[\alpha \log \pi_\theta(a|s) - \min_{j \in \{1,2\}} Q_{\phi_j}(s, a) \right]. \quad (17)$$

The temperature α is adapted online to match a target entropy \mathcal{H} following [19, 28]:

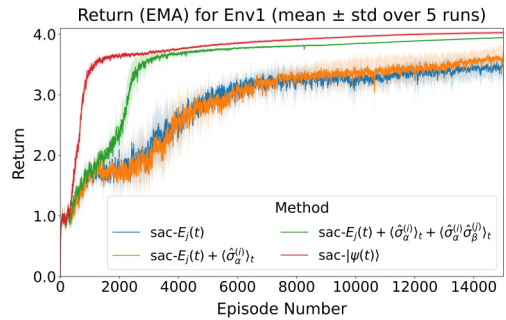
$$\mathcal{L}_{\text{temp}}(\alpha) = \mathbb{E}_{s \sim \mathcal{B}, a \sim \pi_\theta(\cdot|s)} [\alpha(-\log \pi_\theta(a|s) - \mathcal{H})]. \quad (18)$$

IV. SIMULATION RESULTS AND DISCUSSION


A. Simulation parameters

We simulate an inhomogeneous Dicke battery with $N = 4$ TLSs. Inhomogeneity is introduced as static disorder in the TLS frequencies: the frequency of each TLS ω_j ($j = 1, 2, 3, 4$) is randomly sampled from a uniform distribution $U[0.5, 1.5]$ independently, referenced to a central frequency $\omega_0 = 1$. The total charging duration is set to $\tau = 5.6$, discretized into $K = 28$ equal steps of length $\Delta t = 0.2$. The agent controls the coupling strength, chosen continuously from the range $[-0.3, 0.3]$. Each experiment is repeated with five independent random seeds to assess variability.


Note that the disorder realization $\omega_j \sim U[0.5, 1.5]$ corresponds to relative static detunings up to $\delta_{\text{max}} = |\omega_j - \omega_0|/\omega_0 \leq 0.5$. To contextualize the physical relevance of the inhomogeneity parameters used in our simulations, we note that the inhomogeneous Dicke model is directly applicable to several experimental platforms where an ensemble of two-level systems is collectively coupled to a single bosonic mode, such as cavity/circuit QED systems and solid-state spin ensembles coupled to microwave resonators [31–33]. In such systems, typical inhomogeneous broadening—arising from factors like fabrication spreads in superconducting qubits or local field variations in spin ensembles—often ranges from 0.1% to over 10% of the transition frequency [34, 35]. The broader disorder range examined here serves a dual purpose: it provides a conservative stress test for the reinforcement learning agent, verifying its ability to devise effective protocols even under significant detuning, and it accommodates potential additional detuning from environmental fluctuations or device aging, thereby demonstrating the robustness of the learned policies under realistically variable conditions.



(a)



(b)



(c)

Fig 1: Return (smoothed by an exponential moving average with $\gamma_G = 0.9$) versus training episodes under four settings of available information, averaged over five random seeds, for three inhomogeneous quantum batteries: (a) Env1, (b) Env2, (c) Env3. The shaded regions represent the standard deviation over five seeds.

B. Experimental scenarios

To assess the agent's ability to learn protocols under inhomogeneous conditions, we tested three disordered quantum battery systems with distinct qubit frequencies (independently sampled from $U[0.5, 1.5]$): (1) Env1: $[0.8745, 1.4507, 1.2320, 1.0987]$; (2) Env2: $[0.9170, 1.2203, 0.5001, 0.8023]$; (3) Env3:

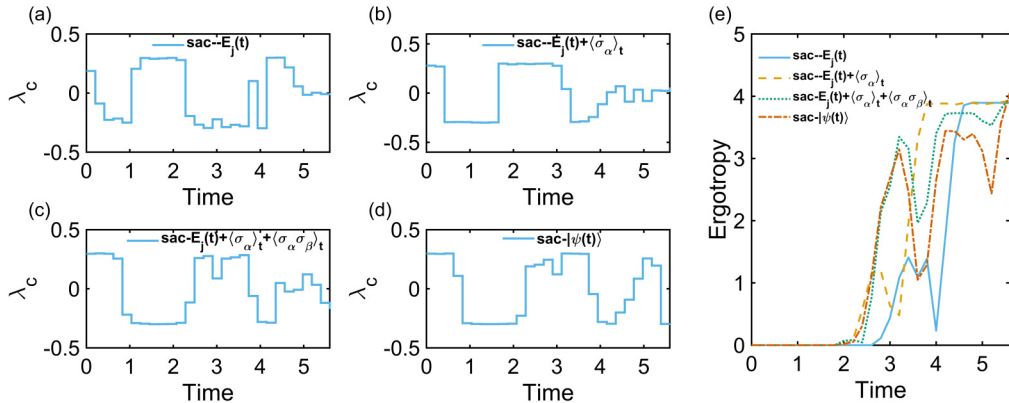


Fig 2: Best charging protocols among five seeds for four our observable regimes [(a): $E_j(t)$, (b): $E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t$, (c): $E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t + \langle \hat{\sigma}_\alpha^{(i)} \hat{\sigma}_\beta^{(j)} \rangle_t$, (d): $|\psi(t)\rangle$] and their ergotropy trajectories (e) in Env1.

[0.7220, 1.3707, 0.7067, 1.4186]. Within each system we vary observability by exposing the agent to one of four observable regimes based on physical quantities:

1. full state $|\psi(t)\rangle$;
2. individual TLS expected energies $E_j(t)$;
3. energies plus all possible first-order averages $\langle \hat{\sigma}_x^{(j)} \rangle_t$, $\langle \hat{\sigma}_y^{(j)} \rangle_t$ and $\langle \hat{\sigma}_z^{(j)} \rangle_t$ with $j \in \{1, 2, \dots, N\}$;
4. energies plus first-order averages and all possible second-order correlations $\langle \hat{\sigma}_\alpha^{(i)} \hat{\sigma}_\beta^{(j)} \rangle_t$ with $\alpha, \beta \in \{x, y, z\}$ and $i, j \in \{1, 2, \dots, N\}$.

During training we log the ergotropy achieved by all charging protocols explored by the agent and, for each setting, report the best value attained (i.e., the maximum over explored protocols).

C. Performance analysis

The numerical results are summarized in Table 1 and the training processes are visualized in Fig. 1, where the return is smoothed using an exponential moving average (EMA) with coefficient γ_G , defined as $\bar{G}_t = \gamma_G \bar{G}_{t-1} + (1 - \gamma_G) G_t$, where G_t is the return at episode t , \bar{G}_t is the EMA-smoothed return, and $\gamma_G \in [0, 1)$ controls the smoothing strength. Fig. 2 further displays, for Env1, the best charging protocol identified among the five seeds and the corresponding ergotropy trajectory. **Additional metrics across seeds, including distribution of final ergotropies and convergence rates, are provided in Appendix A.** From these results, we draw several key findings.

Under full observability, the agent attains the highest possible ergotropy, learning near-optimal charging protocols. It also exhibits the smallest variability across seeds, as evidenced by the narrowest shaded region in Fig. 1. This indicates that access to the complete state enables the agent to make consistent, high-quality control decisions.

Under more realistic, partially observable conditions, a clear performance hierarchy emerges with increasing information content of the observables: $[E_j(t)] \approx [E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t] < [E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t + \langle \hat{\sigma}_\alpha^{(i)} \hat{\sigma}_\beta^{(j)} \rangle_t]$. Observing only individual TLS energies is insufficient to learn an effective protocol, and adding first-order averages offers negligible—if any—improvement. However, including correlation features recovers a substantial fraction of the performance achieved under full observability: across Envs 1–3, the average best ergotropy (over 5 seeds) reaches 98.11%, 96.13%, and 94.64% of the fully observed baseline, respectively.

The performance gap across observability regimes is consistent with aliasing in POMDPs: restricted observables can map distinct quantum states to similar observations, resulting in suboptimal actions and increased variance across training runs. Incorporating second-order correlations enriches the observation with information about inter-spin correlations, which helps mitigate aliasing and improves both terminal performance and training stability.

From Fig. 2(c) and 2(d), enriching the observables with both first-order averages and second-order correlations yields a charging protocol that closely matches the one learned under full-state access. Moreover, Fig. 2(e) shows that an effective protocol does not maximize ergotropy at every instant; the agent selects the coupling strength to

TABLE 1: Mean and 95% confidence interval (CI, using Student- t interval) of the best terminal ergotropy attained during training across five seeds for each setting of available information in Env1–Env3.

Available Information	Ergotropy (mean \pm CI half-width)		
	Env1	Env2	Env3
$ \psi(t)\rangle$	4.0445 ± 0.0323	2.8357 ± 0.0235	3.3777 ± 0.0235
$E_j(t)$	3.8496 ± 0.0459	2.6757 ± 0.0902	3.0938 ± 0.1648
$E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t$	3.8984 ± 0.0326	2.6548 ± 0.0516	3.0585 ± 0.0519
$E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t + \langle \hat{\sigma}_\alpha^{(i)} \hat{\sigma}_\beta^{(j)} \rangle_t$	3.9683 ± 0.0139	2.7260 ± 0.0575	3.1966 ± 0.0397

maximize terminal ergotropy rather than immediate gains.

a. Robustness across disorder realizations. To further assess robustness of the learned RL agent beyond the training realization, we train the agent on Env1 and periodically evaluate it on Env1–Env3 without retraining. We select, for each seed, the checkpoint that maximizes the mean terminal ergotropy across Env1–Env3, and report the corresponding terminal ergotropy achieved on each environment. The resulting robustness evaluation averaged over five random seeds is summarized in Table 2.

Across all three disorder realizations, we observe a consistent ordering of robustness: full-state access yields the strongest cross-realization performance, while energy-only observations generalize the least. Augmenting the observation with first-order moments provides only a limited improvement, whereas including second-order correlations produces a clear robustness gain, consistent with the key role of correlation information under inhomogeneity. Nevertheless, a noticeable robustness gap to the full-state reference remains after adding correlations, especially in Env2, indicating that two-body information, while highly informative, is not sufficient to reach full-state-level robustness across realizations in the present setting.

D. Discussion

Our results confirm that reinforcement learning is an effective approach for learning high-quality charging protocols in inhomogeneous quantum batteries. When full state information (i.e., $|\psi(t)\rangle$) of the system is available, the agent consistently discovers ideal solutions and exhibits stable learning dynamics.

The more consequential finding concerns the practical scenario of partial observability. We show that the choice of observables is critical, with

second-order correlations playing an essential role in attaining high charging performance. In particular, correlation measurements largely compensate for the information gap between single-TLS energy readouts and full-state access, thereby enabling effective optimization.

In addition, an intriguing observation from our results is that first-order averages contribute little to estimation performance, whereas second-order correlations yield a substantial gain. This can be physically explained by a conserved discrete parity symmetry in our model. The symmetry is generated by the operator $\hat{P} = (\otimes_{j=1}^N i\hat{\sigma}_z^{(j)}) \otimes e^{i\pi\hat{a}^\dagger\hat{a}}$.

Under this parity, single-spin operators $\hat{\sigma}_x^{(j)}$ and $\hat{\sigma}_y^{(j)}$ are odd ($\hat{P}\hat{O}\hat{P}^\dagger = -\hat{O}$), while $\hat{\sigma}_z^{(j)}$ is even. Since the initial state is a parity eigenstate and the Hamiltonian commutes with \hat{P} , the expectation values of all odd single-spin operators remain exactly zero throughout the evolution. Consequently, their first-order averages provide no information beyond the energy, which depends solely on the even observable $\langle \hat{\sigma}_z^{(j)} \rangle$. In contrast, second-order correlators such as $\langle \hat{\sigma}_x^{(i)} \hat{\sigma}_x^{(j)} \rangle$ are products of two odd operators, resulting in an even composite observable that can take nonzero values. These correlators thus encode genuine multi-body quantum correlations that are fundamentally absent from the first-order data, accounting for the significant performance improvement when they are included.

The finding that second-order correlations are nearly sufficient for ergotropy optimization, while surprising from a quantum-state-tomography standpoint, has a physical basis. Ergotropy, as a macroscopic observable, does not require full state reconstruction. Its optimization depends predominantly on global energy distributions and collective order, which are typically governed by low-order correlations in physical systems. Consequently, the agent’s ability to achieve near-optimal charging using only these correlations is thus inherent to the physical task itself, not an artifact of the specific RL policy.

TABLE 2: Evaluation result of the learned RL agent (trained on Env1) across disorder realizations (Env1–Env3). We select, for each seed, the checkpoint that maximizes the mean terminal ergotropy across Env1–Env3, and report the corresponding terminal ergotropy achieved on each environment (Mean and 95% confidence interval).

Available Information	Ergotropy(mean \pm CI half-width)		
	Env1	Env2	Env3
$ \psi(t)\rangle$	4.0024 \pm 0.0246	2.6178 \pm 0.1222	3.1248 \pm 0.0905
$E_j(t)$	3.4844 \pm 0.3115	0.9452 \pm 0.6414	1.9187 \pm 0.7559
$E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t$	3.6658 \pm 0.1750	0.9623 \pm 0.6043	2.2375 \pm 0.7476
$E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t + \langle \hat{\sigma}_\alpha^{(i)} \hat{\sigma}_\beta^{(j)} \rangle_t$	3.8776 \pm 0.0537	1.7048 \pm 0.8825	2.7069 \pm 0.2357

Finally, the ergotropy trajectories indicate that effective charging is inherently nonmyopic. The learned agent does not chase instantaneous ergotropy gains; instead, it schedules coupling strengths to improve the terminal outcome, even at the cost of temporary plateaus or counterintuitive moves.

In the correlation observability regime, the agent is provided with the full set of two-body Pauli–Pauli expectation values $\{\langle \sigma_\alpha^{(i)} \sigma_\beta^{(j)} \rangle_t\}$, in addition to the single-body information. While a naive counting suggests up to $9N^2$ distinct correlators, physical constraints, including parity and exchange symmetry drastically reduce this number to just $5N(N-1)/2$ independent measurements, a reduction of approximately 72% from the naive scaling. For larger systems, this quadratic overhead can be further mitigated. Techniques from compressed sensing enable the reconstruction of all correlations from a number of random measurements scaling only linearly or sub-quadratically with N [36]. Alternatively, a recurrent neural network policy could be trained on a time-series of easily measurable local observables, implicitly estimating the correlations within a partially observable framework [30]. For the moderate system size studied here, and given the capability for parallel readout in platforms like superconducting qubits or trapped ions, the measurement set is experimentally manageable, ensuring the practical relevance of our approach.

V. CONCLUSIONS

In this work, we formulated fast charging of an inhomogeneous Dicke quantum battery as a discrete-time reinforcement learning problem and trained piecewise-constant control policies with Soft Actor–Critic algorithm. The agent optimizes terminal ergotropy using a dense, shaped reward that blends energy and ergotropy gains, and acts

within a bounded, continuous action space for the charger–battery coupling strength. To reflect realistic readout constraints, we examined four observability regimes—from full state access to experimentally accessible summaries built from single-TLS energies, first-order averages, and second-order correlation features.

Across three disordered four-TLS quantum batteries, SAC consistently discovers high-quality charging protocols. Under full observability, the agent attains near-optimal terminal ergotropy with low seed-to-seed variability. Under partial observability, a clear hierarchy emerges: access to only single-TLS energies ($E_j(t)$) or energies plus first-order averages ($E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t$) lags behind the fully observed baseline, whereas augmenting with second-order correlations ($E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t + \langle \hat{\sigma}_\alpha^{(i)} \hat{\sigma}_\beta^{(j)} \rangle_t$) recovers most of the gap, reaching 94%–98% of the full-state performance depending on the battery parameters. Analysis of the learned schedules further shows that effective charging is nonmyopic: the agent modulates the coupling to maximize terminal ergotropy rather than pursuing instantaneous gains, accepting temporary plateaus or declines to achieve superior final outcomes.

These results establish RL as a practical tool for optimizing fast-charging protocols in inhomogeneous quantum batteries under information constraints and point to several promising extensions: (i) **Adopting more principled POMDP approaches**—for example, equipping the agent with recurrent architectures that use observation histories—could help the policy form an internal belief over latent quantum states, thereby handling partial observability induced by compressed, experimentally accessible readouts. (ii) Seeking robust fast-charging protocols under per-episode fluctuations of TLS frequencies within specified ranges can improve reliability when device parameters drift between charges. Together, these extensions would enhance practicality by addressing both information limitations and parameter variability in ex-

perimental quantum batteries.

Appendix A: Additional metrics across seeds

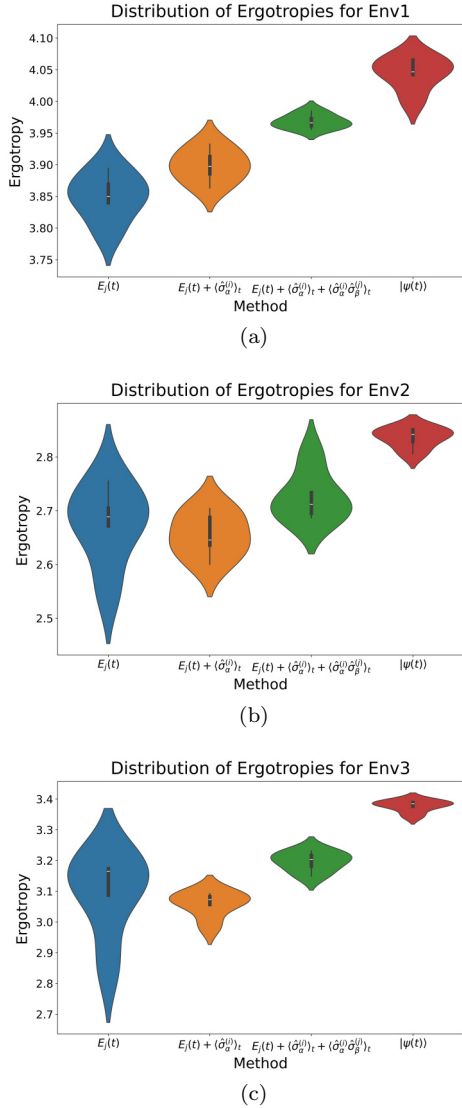


Fig S1: The distribution of the best terminal ergotropies attained during training under four settings of available information over five random seeds, for three inhomogeneous quantum batteries: (a) Env1, (b) Env2, (c) Env3.

Fig. S1 summarizes the distribution of the best terminal ergotropy attained during training under four observability settings for Env1–Env3. Across all disorder realizations, access to the full state consistently yields the highest terminal ergotropy and, at the same time, typically exhibits the most concentrated distributions, indicating both supe-

rior achievable performance and strong robustness to random initialization. In contrast, restricted observations lead to broader spreads and lower typical values. Notably, augmenting the observation with second-order correlations produces a clear and systematic improvement over energy-only and first-order information: the distributions shift toward higher ergotropy, in most cases, become more concentrated, reflecting a more reliable optimization outcome across seeds. Nevertheless, a non-negligible gap to the full-state reference remains, consistent with the partial-observability limitations discussed in the main text.

Table S1 quantifies the convergence rates. We define $T_{0.9}$ as the number of training episodes required for the best-so-far terminal ergotropy to first reach 90% of its final best value. As expected, full-state access generally leads to the fastest optimization (smallest $T_{0.9}$), reflecting more informative feedback for learning. By comparison, adding second-order correlations does not necessarily accelerate convergence relative to lower-order observations and can even lead to slower convergence in some environments; however, it yields substantially better terminal ergotropy as shown in Fig. S1 and Table 1. In our context, this trade-off is favorable: the primary objective is to maximize the extractable work at the end of charging, and the correlation-augmented setting achieves markedly improved final performance even when the convergence rate is not improved.

Appendix B: Sensitivity to the reward-mixing weight

The main text employs a shaped reward that mixes the energy and ergotropy increments $r_k = w \Delta E_k + (1 - w) \Delta \mathcal{E}_k$, where w is either fixed or scheduled during training according to Eq. (11) and (12), motivated by the observation that ergotropy increments can be sparse at early training stages. To assess how sensitive learning is to the reward-weighting choice, we perform an ablation study on the reward mixing under the same environment and observation setting used for our main partial-observability results. Specifically, all agents in this appendix are trained on Env1 using the correlation-augmented observation $E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t + \langle \hat{\sigma}_\alpha^{(i)} \hat{\sigma}_\beta^{(j)} \rangle_t$, and differ only in the reward-weighting scheme. The corresponding learning curves are shown in Fig. S2.

The comparison highlights the distinct roles of the energy and ergotropy terms at different stages of learning. At early training, increments in ergotropy are frequently zero (or extremely small), making $\Delta \mathcal{E}_k$ a sparse learning signal. Assigning

TABLE S1: Convergence rate across five seeds. $T_{0.9}$ denotes the number of training episodes required for the terminal ergotropy attained during training to first reach 90% of its final best value.

Available Information	$T_{0.9}$ (mean \pm CI half-width)		
	Env1	Env2	Env3
$ \psi(t)\rangle$	948.8 \pm 166.6	1474.8 \pm 374.0	3209.2 \pm 190.8
$E_j(t)$	1744.0 \pm 390.8	1631.8 \pm 418.5	3436.2 \pm 689.0
$E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t$	2039.8 \pm 499.0	1700.6 \pm 276.1	3653.4 \pm 1185.5
$E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t + \langle \hat{\sigma}_\alpha^{(i)} \hat{\sigma}_\beta^{(j)} \rangle_t$	1921.4 \pm 326.0	2097.0 \pm 610.6	4350.4 \pm 2530.3




Fig S2: Sensitivity to the reward-mixing weight. Comparison of scheduled w with fixed $w = 0.5$ (mixed reward) and $w = 0$ (ergotropy-only reward) in $r_k = w \Delta E_k + (1 - w) \Delta \mathcal{E}_k$. All agents are trained on Env1 with the correlation-augmented observation $E_j(t) + \langle \hat{\sigma}_\alpha^{(i)} \rangle_t + \langle \hat{\sigma}_\alpha^{(i)} \hat{\sigma}_\beta^{(j)} \rangle_t$. Curves show the ergotropy during training (EMA smoothing with $\gamma = 0.9$); shaded regions indicate the standard deviation over five seeds.

a nonzero weight to the energy increment therefore provides denser feedback and accelerates the initial growth of the achieved ergotropy. As exploration proceeds and nontrivial ergotropy gains become more common, increasing the relative weight on $\Delta \mathcal{E}_k$ guides the agent toward more goal-directed exploration and leads to faster improvement in ergotropy. Consistent with this interpretation, the scheduled-weight strategy—which interpolates between these regimes—achieves the best overall trade-off and attains the highest ergotropy at late training compared with either fixed weighting.

- [1] R. Alicki and M. Fannes. Entanglement boost for extractable work from ensembles of quantum batteries. *Phys. Rev. E*, 87:042123, 2013.
- [2] G. M. Andolina, M. Keck, A. Mari, V. Giovannetti, and M. Polini. Quantum versus classical many-body batteries. *Phys. Rev. B*, 99:205437, 2019.
- [3] F. Campaioli, S. Gherardini, J. Q. Quach, M. Polini, and G. M. Andolina. Colloquium: Quantum batteries. *Rev. Mod. Phys.*, 96:031001, 2024.
- [4] J. Q. Quach, K. E. McGhee, L. Ganzer, D. M. Rouse, B. W. Lovett, E. M. Gauger, J. Keeling, G. Cerullo, D. G. Lidzey, and T. Virgili. Superabsorption in an organic microcavity: Toward a quantum battery. *Sci. Adv.*, 8(2):eabk3160, 2022.
- [5] F. Campaioli, F. A. Pollock, F. C. Binder, L. C. Céleri, J. Goold, S. Vinjanampathy, and K. Modi. Enhancing the charging power of quantum batteries. *Phys. Rev. Lett.*, 118:150601, 2017.
- [6] F. C. Binder, S. Vinjanampathy, K. Modi, and J. Goold. Quantacell: powerful charging of quantum batteries. *New J. Phys.*, 17:075015, 2015.
- [7] Z.-G. Lu, G. Tian, X.-Y. Lü, and C. Shang. Topological quantum batteries. *Phys. Rev. Lett.*, 134:180401, 2025.
- [8] G. Gemme, G. M. Andolina, F. M. D. Pellegrino, M. Sasseti, and D. Ferraro. Off-resonant dicke quantum battery: Charging by virtual photons. *Batteries*, 9(4):197, 2023.
- [9] G. M. Andolina, M. Keck, A. Mari, M. Campisi, V. Giovannetti, and M. Polini. Extractable work, the role of correlations, and asymptotic freedom in quantum batteries. *Phys. Rev. Lett.*, 122:047702, 2019.
- [10] C. Brif, R. Chakrabarti, and H. Rabitz. Control

- of quantum phenomena: past, present, and future. *New J. Phys.*, 12:075008, 2010.
- [11] D. Rossini, G. M. Andolina, D. Rosa, M. Carrega, and M. Polini. Quantum advantage in the charging process of sachdev-ye-kitaev batteries. *Phys. Rev. Lett.*, 125:236402, 2020.
- [12] D. Ferraro, M. Campisi, G. M. Andolina, V. Pellegrini, and M. Polini. High-power collective charging of a solid-state quantum battery. *Phys. Rev. Lett.*, 120:117702, 2018.
- [13] A. Tagliabue, D.-K. Kim, M. Everett, and J. P. How. Efficient guided policy search via imitation of robust tube mpc. In *2022 Int. Conf. on Rob. and Autom.*, pages 462–468. IEEE, 2022.
- [14] A. Pozzi and D. Toti. Imitation learning for agnostic battery charging: A dagger-based approach. *IEEE Access*, 11:115190–115203, 2023.
- [15] A. Pozzi, A. Incremona, and D. Toti. Imitation learning-driven approximation of stochastic control models. *Appl. Intell.*, 55(12):1–19, 2025.
- [16] R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT press Cambridge, MA, 1998.
- [17] 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 preprint arXiv:1509.02971*, 2015.
- [18] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proc. Int. Conf. Mach. Learn.*, pages 1861–1870, 2018.
- [19] P. A. Erdman, G. M. Andolina, V. Giovannetti, and F. Noé. Reinforcement learning optimization of the charging of a dicke quantum battery. *Phys. Rev. Lett.*, 133:243602, 2024.
- [20] P.-Y. Sun, H. Zhou, and F.-Q. Dou. Cavity-heisenberg spin- j chain quantum battery and reinforcement learning optimization. *arXiv preprint arXiv:2412.01442*, 2024.
- [21] S. Zakavati, S. Salimi, and B. Arash. Optimizing the charging of open quantum batteries using long short-term memory-driven reinforcement learning. *arXiv preprint arXiv:2504.19840*, 2025.
- [22] M. Bukov. Reinforcement learning for autonomous preparation of floquet-engineered states: Inverting the quantum kapitza oscillator. *Phys. Rev. B*, 98:224305, 2018.
- [23] 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.
- [24] M. Y. Niu, S. Boixo, V. N. Smelyanskiy, and H. Neven. Universal quantum control through deep reinforcement learning. *npj Quantum Inf.*, 5:33, 2019.
- [25] S. Li, Y. Fan, X. Li, X. Ruan, Q. Zhao, Z. Peng, R.-B. Wu, J. Zhang, and P. Song. Robust quantum control using reinforcement learning from demonstration. *npj Quantum Inf.*, 11:124, 2025.
- [26] Z. Zhiqiang, C. H. Lee, R. Kumar, K. J. Arnold, S. J. Masson, A. S. Parkins, and M. D. Barrett. Nonequilibrium phase transition in a spin-1 dicke model. *Optica*, 4:424–429, 2017.
- [27] F. Dimer, B. Estienne, A. S. Parkins, and H. J. Carmichael. Proposed realization of the dicke-model quantum phase transition in an optical cavity qed system. *Phys. Rev. A*, 75:013804, 2007.
- [28] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, and S. Levine. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905*, 2019.
- [29] P. Kirton, M. M. Roses, J. Keeling, and E. G. Dalla Torre. Introduction to the dicke model: From equilibrium to nonequilibrium, and vice versa. *Adv. Quantum Technol.*, 2(1-2):1800043, 2019.
- [30] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artif. Intell.*, 101(1):99–134, 1998.
- [31] Kristian Baumann, Christine Guerlin, Ferdinand Brennecke, and Tilman Esslinger. Dicke quantum phase transition with a superfluid gas in an optical cavity. *Nature*, 464(7293):1301–1306, April 2010.
- [32] Nicolas Marquez Peraca, Xinwei Li, Jaime M. Moya, Kenji Hayashida, Dasom Kim, Xiaoxuan Ma, Kelly J. Neubauer, Diego Fallas Padilla, Chien-Lung Huang, Pengcheng Dai, Andriy H. Nevidomskyy, Han Pu, Emilia Morosan, Shixun Cao, Motoaki Bamba, and Junichiro Kono. Quantum simulation of an extended dicke model with a magnetic solid. *Communications Materials*, 5(1):42, 2024.
- [33] S. V. Remizov, A. A. Zhukov, W. V. Pogosov, and Yu. E. Lozovik. Analog-digital quantum simulation of the dicke model with superconducting circuits. *JETP Letters*, 108(11):748–753, December 2018.
- [34] J. Van Damme, S. Massar, R. Acharya, Ts. Ivanov, D. Perez Lozano, Y. Canvel, M. Demarets, D. Vangoidsenhoven, Y. Hermans, J. G. Lai, A. M. Vadiraj, M. Mongillo, D. Wan, J. De Boeck, A. Potočnik, and K. De Greve. Advanced cmos manufacturing of superconducting qubits on 300 mm wafers. *Nature*, 634(8032):74–79, October 2024.
- [35] Kento Sasaki, Ed E. Kleinsasser, Zhouyang Zhu, Wen-Di Li, Hideyuki Watanabe, Kai-Mei C. Fu, Kohei M. Itoh, and Eisuke Abe. Dynamic nuclear polarization enhanced magnetic field sensitivity and decoherence spectroscopy of an ensemble of near-surface nitrogen-vacancy centers in diamond. *Applied Physics Letters*, 110(19):192407, 05 2017.
- [36] David Gross, Yi-Kai Liu, Steven T. Flammia, Stephen Becker, and Jens Eisert. Quantum state tomography via compressed sensing. *Phys. Rev. Lett.*, 105:150401, Oct 2010.