





Quantum reinforcement learning in the presence of thermal dissipation

María Laura Olivera-Atencio ¹, Lucas Lamata ^{2,3}, Manuel Morillo ¹ and Jesús Casado-Pascual ^{1,*}

¹*Física Teórica, Universidad de Sevilla, Apartado de Correos 1065, Sevilla 41080, Spain*

²*Departamento de Física Atómica, Molecular y Nuclear, Universidad de Sevilla, 41080 Sevilla, Spain*

³*Instituto Carlos I de Física Teórica y Computacional, 18071 Granada, Spain*



(Received 9 August 2022; revised 30 March 2023; accepted 30 June 2023; published 26 July 2023)

A study of the effect of thermal dissipation on quantum reinforcement learning is performed. For this purpose, a nondissipative quantum reinforcement learning protocol is adapted to the presence of thermal dissipation. Analytical calculations as well as numerical simulations are carried out, obtaining evidence that dissipation does not significantly degrade the performance of the quantum reinforcement learning protocol for sufficiently low temperatures, in some cases even being beneficial. Quantum reinforcement learning under realistic experimental conditions of thermal dissipation opens an avenue for the realization of quantum agents to be able to interact with a changing environment, as well as adapt to it, with many plausible applications inside quantum technologies and machine learning.

DOI: [10.1103/PhysRevE.108.014128](https://doi.org/10.1103/PhysRevE.108.014128)

I. INTRODUCTION

Quantum machine learning [1,2] aims at employing quantum technologies to achieve machine learning tasks more efficiently. Diverse quantum algorithms for machine learning have been proposed, including solvers for linear systems of equations [3], quantum principal component analysis [4], quantum support vector machines [5], quantum annealers [6], quantum variational eigensolvers [7], quantum reinforcement learning [8–15], quantum generative adversarial networks [16], and quantum kernels [17]. Experimental implementations showing, in some instances, speedups with respect to classical or other kinds of quantum algorithms have been carried out on noisy intermediate-scale quantum devices [16–21]. Many of these works study the advantages of quantum technologies over classical machine learning. However, the effect of thermal dissipation on quantum machine learning has not yet been analyzed in depth, although performing a real quantum experiment always involves working at nonzero temperature and therefore a certain amount of such dissipation. Some preliminary results on this issue can be found in Refs. [22,23].

In the field of quantum machine learning, as in the classical case, algorithms can be classified into three main groups: Supervised [24–32], unsupervised [33,34], and reinforcement quantum algorithms [8–15]. In this paper we focus on the third group, more specifically, on the effects of thermal dissipation on this type of algorithm. Reinforcement learning algorithms consist of successive interactions between a known agent and an unknown environment alternating with a reward function that improves a specific task performed by such an agent [35,36]. The goal is to learn from the environment. The information extraction from the environment and the channel used to communicate the information to the agent and the action of

the agent are established by the policy. The reward function defines the criterion used to punish or reward certain actions of the agent in order to improve its performance. Each reward reinforces the current strategy, while the punishment forces an adaptation of the strategy.

To study the effect of dissipation on quantum reinforcement learning, in this paper we adapt the nondissipative algorithm proposed in Ref. [15] to the presence of thermal dissipation. For this purpose, the environment will be considered to be an unknown quantum system whose dynamics, besides including a unitary part associated with a given Hamiltonian as in [15], also possesses a certain amount of thermal dissipation. The goal is to extract information or learn from the environment to obtain near-optimal knowledge of the eigenstates of the Hamiltonian. The agent corresponds to a known and manipulable quantum state that must conveniently adjust to the dynamics of the environment in order to approach the unknown eigenstate. It should be noted that the aim of this work is not to analyze the scalability or speedup of this type of protocols, which has been discussed previously [15,18], but to study the effect of thermal dissipation on it. For this reason, and for the sake of clarity, we will focus on the simplest case of a two-level quantum system.

The structure of the remainder of this paper is as follows. In Sec. II the nondissipative protocol proposed in Ref. [15] is adapted to the presence of thermal dissipation. Emphasis is placed on the peculiarities associated with the presence of such dissipation. In Sec. III we discuss the numerical implementation of the protocol and illustrate our results with numerical simulations. In Sec. IV we present a summary and conclusions for the main findings of our work.

II. PROTOCOL DESCRIPTION

In the nondissipative protocol proposed in Ref. [15], the agent A is considered to be a known manipulable quantum system described by a state vector $|\phi\rangle$. The environment E is

*jcasado@us.es

modeled as a black box that interacts with A for a time τ . The effect of this interaction on A is characterized by a unitary transformation $U \equiv e^{-i\tau H/\hbar}$ applied to the state vector $|\phi\rangle$, where H is an unknown interaction Hamiltonian whose eigenvectors are to be computed. For simplicity, we will henceforth restrict the analysis to the case of a single qubit with state basis vectors $\{|0\rangle, |1\rangle\}$ and unknown interaction Hamiltonian

$$H = \frac{\hbar\omega}{2}(|+\rangle\langle +| - |-\rangle\langle -|), \quad (1)$$

where ω is a positive constant with dimensions of frequency and $\{|+\rangle, |-\rangle\}$ are the eigenvectors to be computed.

Suppose now that, in addition to interacting with E , A is also in contact with a thermal bath B at a finite temperature T . We will also assume that the combined action of E and B on A is described by a Lindblad master equation of the form [37]

$$\dot{\rho}_t = -\frac{i}{\hbar}[H, \rho_t] + \sum_{j=\pm} \Gamma_j \left(\tilde{\sigma}_j^\dagger \rho_t \tilde{\sigma}_j - \frac{1}{2} \{ \tilde{\sigma}_j \tilde{\sigma}_j^\dagger, \rho_t \} \right), \quad (2)$$

where ρ_t is the density operator representing the state of A at time t , $\tilde{\sigma}_- = |-\rangle\langle +| = \tilde{\sigma}_+^\dagger$ is a Lindblad operator that induces dissipative decay from the excited state $|+\rangle$ to the ground state $|-\rangle$, and $\Gamma_\pm = \Gamma_0 e^{\pm\hbar\omega/(2k_B T)} \text{csch}[\hbar\omega/(2k_B T)]/2$, with Γ_0 the decay rate from the excited state to the ground state at zero temperature. The frequency ω in Eq. (1) must be redefined to include the frequency shift caused by the presence of the thermal bath. The first term on the right-hand side in Eq. (2) describes the coherent evolution of the system, while the second term gives rise to dissipation.

By solving Eq. (2), the density operator at an arbitrary time τ can be expressed in terms of the initial density operator ρ_0 in the form

$$\rho_\tau = \mathcal{E}(\rho_0) \equiv \sum_{j=0}^3 U E_j \rho_0 E_j^\dagger U^\dagger, \quad (3)$$

where $\{E_0, E_1, E_2, E_3\}$ are the Kraus operators for the generalized amplitude damping channel [38]. The explicit expressions for these operators are [38]

$$E_0 = \sqrt{p_+}(|+\rangle\langle +| + \sqrt{1-\gamma}|-\rangle\langle -|), \quad (4)$$

$$E_1 = \sqrt{p_+}\gamma\tilde{\sigma}_+, \quad (5)$$

$$E_2 = \sqrt{1-p_+}(\sqrt{1-\gamma}|+\rangle\langle +| + |-\rangle\langle -|), \quad (6)$$

$$E_3 = \sqrt{(1-p_+)\gamma}\tilde{\sigma}_-, \quad (7)$$

where $p_+ = e^{-\hbar\omega/(2k_B T)} \text{sech}[\hbar\omega/(2k_B T)]/2$ is the thermal equilibrium probability of the excited state and $\gamma = 1 - e^{-\Gamma\tau}$, with $\Gamma = \Gamma_+ + \Gamma_- = \Gamma_0 \coth[\hbar\omega/(2k_B T)]$. Note that in the absence of dissipation, i.e., for $\Gamma_0 = 0$, the parameter γ vanishes and therefore $E_0 = \sqrt{p_+}I$, $E_2 = \sqrt{1-p_+}I$, and $E_1 = E_3 = 0$, with I the identity operator. Thus, Eq. (3) reduces to $\rho_\tau = U\rho_0 U^\dagger$, which is the case considered in Ref. [15] with $\rho_0 = |\phi\rangle\langle\phi|$.

Bearing in mind the above results, the protocol proposed in Ref. [15] can be adapted to the dissipative case as detailed below. The procedure involves very many iterations, so the state of A in the k th iteration is denoted by $|\phi^{(k)}\rangle$, with $k \in \mathbb{N}$. We assume that, in the first iteration, A is prepared in one

of the basis states, for instance, in the state $|\phi^{(1)}\rangle = |0\rangle$. The states $|0\rangle$ and $|\phi^{(k)}\rangle$ are related by

$$|\phi^{(k)}\rangle = D^{(k)}|0\rangle, \quad (8)$$

where $D^{(k)}$ is a unitary operator constructed inductively, starting with $D^{(1)} = I$ and building $D^{(k+1)}$ out of $D^{(k)}$ as follows.

(i) From the initial density operator $\rho_0^{(k)} = |\phi^{(k)}\rangle\langle\phi^{(k)}|$, we let the system evolve according to the Lindblad equation (2) for a time τ . We call $\rho_\tau^{(k)} = \mathcal{E}(\rho_0^{(k)})$ the density at the end of that evolution obtained by Eq. (3).

(ii) We extract information from $\rho_\tau^{(k)}$ by measuring the observable $M^{(k)} = D^{(k)}|1\rangle\langle 1|D^{(k)\dagger}$. In order to always measure the same observable $M^{(1)} = |1\rangle\langle 1|$ for all iterations, first we apply the unitary transformation $D^{(k)\dagger}\rho_\tau^{(k)}D^{(k)}$ and then we measure $M^{(1)}$. After the measurement process, the state of A is $|m^{(k)}\rangle$, with $m^{(k)}$ the outcome of the measurement, which can be 0 or 1 with probabilities

$$P_0^{(k)} = \langle 0|D^{(k)\dagger}\rho_\tau^{(k)}D^{(k)}|0\rangle \quad (9)$$

and $P_1^{(k)} = 1 - P_0^{(k)}$, respectively. Once the measurement has been completed, the state $|\phi^{(k)}\rangle$ can be reconstructed from $|m^{(k)}\rangle$ by the unitary transformation $D^{(k)}[(1 - m^{(k)})I + m^{(k)}\sigma_x]|m^{(k)}\rangle$, with $\sigma_x = |0\rangle\langle 1| + |1\rangle\langle 0|$.

(iii) If the outcome of the measurement is $m^{(k)} = 1$, we generate three pseudorandom angles $\alpha_x^{(k)}$, $\alpha_y^{(k)}$, and $\alpha_z^{(k)}$, uniformly distributed in the exploration interval $[-w^{(k)}\pi, w^{(k)}\pi]$. The width of this interval is controlled by the exploration parameter $w^{(k)}$, which is computed inductively, starting with $w^{(1)} = 1$ (maximum width) and building $w^{(k+1)}$ out of $w^{(k)}$ using the rule

$$w^{(k+1)} = \min\{1, [(1 - m^{(k)})r + m^{(k)}p]w^{(k)}\}, \quad (10)$$

with $r < 1$ and $p > 1$ the reward and the punishment rates [15], respectively. In other words, the value of the exploration parameter is updated according to the outcome of the preceding measurement $m^{(k)}$ and the corresponding reward or punishment. Every time the measurement outcome is $m^{(k)} = 1$ a punishment is applied by increasing the value of the exploration parameter from $w^{(k)}$ to $w^{(k+1)} = \min(1, pw^{(k)})$, thus widening the exploration interval. The min function is required since the maximum value of the exploration parameter is 1. Conversely, when the measurement outcome is $m^{(k)} = 0$ a reward is granted by decreasing the value of the exploration parameter from $w^{(k)}$ to $w^{(k+1)} = rw^{(k)}$, thus narrowing the exploration interval. Once the three pseudorandom angles $\alpha_x^{(k)}$, $\alpha_y^{(k)}$, and $\alpha_z^{(k)}$ have been calculated, they are used to implement the pseudorandom rotation

$$R^{(k)} = e^{-i\alpha_y^{(k)}\sigma_y^{(k)}/2} e^{-i\alpha_z^{(k)}\sigma_z^{(k)}/2} e^{-i\alpha_x^{(k)}\sigma_x^{(k)}/2}, \quad (11)$$

where the operators $\sigma_x^{(k)}$, $\sigma_y^{(k)}$, and $\sigma_z^{(k)}$ are related to the Pauli operators $\sigma_x, \sigma_y = i(|1\rangle\langle 0| - |0\rangle\langle 1|)$, and $\sigma_z = |0\rangle\langle 0| - |1\rangle\langle 1|$ by the unitary transformation $\sigma_\alpha^{(k)} = D^{(k)}\sigma_\alpha D^{(k)\dagger}$, with $\alpha = x, y$, and z .

(iv) Finally, we construct $D^{(k+1)}$ from $D^{(k)}$ as

$$D^{(k+1)} = [(1 - m^{(k)})I + m^{(k)}R^{(k)}]D^{(k)}. \quad (12)$$

Therefore, the trade-off between exploration and exploitation, which is a characteristic of reinforcement learning [35], is

regulated using the measurement outcome $m^{(k)}$. If the measurement outcome is $m^{(k)} = 1$, the agent decides to explore and modifies its state from $|\phi^{(k)}\rangle$ to $|\phi^{(k+1)}\rangle = R^{(k)}|\phi^{(k)}\rangle$. In contrast, if the measurement outcome is $m^{(k)} = 0$, the agent decides to exploit and keeps its state invariant, i.e., $|\phi^{(k+1)}\rangle = |\phi^{(k)}\rangle$.

Since, according to Eq. (11), $R^{(k)}$ is related to the pseudo-random rotation

$$\check{R}^{(k)} = e^{-i\alpha_y^{(k)}\sigma_y/2} e^{-i\alpha_z^{(k)}\sigma_z/2} e^{-i\alpha_x^{(k)}\sigma_x/2} \quad (13)$$

by the unitary transformation $R^{(k)} = D^{(k)}\check{R}^{(k)}D^{(k)\dagger}$, Eq. (12) can also be expressed in the form

$$D^{(k+1)} = D^{(k)}[(1 - m^{(k)})I + m^{(k)}\check{R}^{(k)}]. \quad (14)$$

The recursive relations (12) and (14) can easily be solved to yield

$$D^{(k+1)} = \prod_{j=k}^1 [(1 - m^{(j)})I + m^{(j)}R^{(j)}] \quad (15)$$

and

$$D^{(k+1)} = \prod_{j=1}^k [(1 - m^{(j)})I + m^{(j)}\check{R}^{(j)}], \quad (16)$$

respectively, where we use the ordered product notation $\prod_{j=k}^1 A_j = A_k A_{k-1} \cdots A_1$ and $\prod_{j=1}^k A_j = A_1 A_2 \cdots A_k$.

The dissipative protocol just described, as well as the nondissipative protocol analyzed in [15], can be related to a Markov decision process (MDP) [39,40], as in the standard reinforcement learning setup. The state space of this MDP is represented by the set of all possible quantum states $|\phi\rangle$ in which the agent can be found. The action space is the set of all rotations R , including the identity I , acting on these quantum states. Since the actions or rotations act deterministically on the agent states, the MDP is deterministic [36] and consequently its transition model can be described in terms of a transition function rather than a transition probability. Specifically, the transition function is $\mathcal{T}(|\phi\rangle, R) = R|\phi\rangle$, i.e., as a result of the action R applied on state $|\phi\rangle$, the state changes to $|\phi'\rangle = R|\phi\rangle$. Finally, the reward function of the MDP is described by the function $\mathcal{R}(|\phi\rangle, R) = \mathcal{R}(R)$ that takes the value r if $R = I$ and p if $R \neq I$. Given the functions \mathcal{T} and \mathcal{R} , it is sufficient to know the current state $|\phi\rangle$ and the current action R to determine the next state $|\phi'\rangle$ and the corresponding reward, thus fulfilling the Markov property. Note, however, that the decision rule (iv) that determines which action to select, given the agent state at step k , is randomized and history dependent [39,40]. Its randomness stems both from the intrinsically random nature of the measurement outcomes $m^{(k)}$ and from the dependence of the rotations $R^{(k)}$ on the pseudo-random angles $\alpha_x^{(k)}$, $\alpha_y^{(k)}$, and $\alpha_z^{(k)}$. Its history dependence is due to the fact that, as described in (iii), to generate these pseudo-random angles, the value of the exploration parameter $w^{(k)}$ must first be determined, which implies solving the recurrence relation in Eq. (10) and thus knowing all the previous history.

III. NUMERICAL RESULTS

We have carried out numerical simulations implementing the protocol presented in the preceding section for a Hamiltonian of the form

$$H = \frac{\hbar\omega}{2}\sigma_x, \quad (17)$$

which corresponds to taking $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$ in Eq. (1). To deal with dimensionless quantities, we introduce the dimensionless parameters $\tilde{\tau} = \omega\tau$, $\tilde{\Gamma}_0 = \Gamma_0/\omega$, and $\tilde{T} = k_B T/(\hbar\omega)$. To simulate numerically the measurement process appearing in item (ii) of Sec. II, we first calculate the probability $P_0^{(k)}$ of obtaining $m^{(k)} = 0$ using Eq. (9). Then we draw a pseudorandom number $\xi^{(k)}$ uniformly distributed in the interval $[0, 1]$. If $\xi^{(k)} \leq P_0^{(k)}$, the outcome of the measurement is $m^{(k)} = 0$, whereas if $\xi^{(k)} > P_0^{(k)}$ the result will be $m^{(k)} = 1$.

In order to assess the accuracy of the protocol, we calculate the fidelity between the state $|\phi^{(k)}\rangle$ and the closest eigenvector of H as a function of the number of iterations k . Since it is not known *a priori* whether this eigenvector is $|+\rangle$ or $|-\rangle$, we consider the greater of the two values, i.e.,

$$f^{(k)} = \max(|\langle +|\phi^{(k)}\rangle|, |\langle -|\phi^{(k)}\rangle|). \quad (18)$$

The closer the value of $f^{(k)}$ is to 1 as k increases, the more accurate the computation of the corresponding eigenvector will be. In addition, at each iteration, the convergence of the protocol is quantified by the exploration parameter $w^{(k)}$ defined in item (iii) of Sec. II. The protocol is considered to converge if $w^{(k)}$ approaches zero as k increases. Moreover, the faster it approaches zero, the faster the convergence of the protocol is.

In the numerical simulations presented here, we repeat the protocol a large number N of realizations. Then, for each value of k , the mean fidelity

$$F^{(k)} = \frac{1}{N} \sum_{j=1}^N f_j^{(k)} \quad (19)$$

and the mean exploration parameter

$$W^{(k)} = \frac{1}{N} \sum_{j=1}^N w_j^{(k)} \quad (20)$$

are obtained by averaging over the N realizations, where the subscript j refers to the j th realization of the protocol. The number of realizations considered in this paper is $N = 1000$. In Fig. 1 we plot the mean fidelity $F^{(k)}$ [Figs. 1(a) and 1(b)] and the mean exploration parameter $W^{(k)}$ [Figs. 1(c) and 1(d)] versus the number of iterations k . We consider two different states for the first iteration, namely, $|\phi^{(1)}\rangle = (\sqrt{3}|0\rangle + |1\rangle)/2$ in Figs. 1(a) and 1(c) and $|\phi^{(1)}\rangle = (2\sqrt{2}|0\rangle + |1\rangle)/3$ in Figs. 1(b) and 1(d). With different types of lines we depict different values of the reward parameter r . When r is modified, the value of the punishment parameter p is also varied according to the relation $p = 2/r$. The parameters associated with thermal dissipation, i.e., the dimensionless decay rate at zero temperature and the dimensionless temperature, are set at $\tilde{\Gamma}_0 = 0.5$ and $\tilde{T} = 0.3$, respectively. As can be seen, the values of the reward and punishment parameters affect the accuracy and convergence of the protocol. The asymptotic values of




FIG. 1. Dependence of (a) and (b) the mean fidelity $F^{(k)}$ and (c) and (d) the mean exploration parameter $W^{(k)}$ on the number of iterations k for different values of the reward parameter r . The state in the first iteration is (a) and (c) $|\phi^{(1)}\rangle = (\sqrt{3}|0\rangle + |1\rangle)/2$ and (b) and (d) $|\phi^{(1)}\rangle = (2\sqrt{2}|0\rangle + |1\rangle)/3$. The remaining parameter values are $\tilde{\Gamma}_0 = 0.5$, $\tilde{T} = 0.3$, $\tilde{\tau} = 1$, $N = 1000$, and $p = 2/r$.

$F^{(k)}$ in the large iteration limit are less than 0.94 when $r = 0.7$, reach values close to 0.98 in the case of $r = 0.9$, and take intermediate values for $r = 0.8$. A more detailed analysis of how the nondissipative protocol is affected by varying the values of the reward and punishment parameters is provided in Ref. [15]. Note, however, that the particular state chosen for the first iteration does not seem to play a relevant role in the accuracy and convergence of the protocol, as can be seen by comparing Figs. 1(a) and 1(c) with Figs. 1(b) and 1(d). In order to focus on the effect of thermal dissipation on the performance of the protocol, we will henceforth set the reward and punishment parameters to $r = 0.9$ and $p = 2/r = \frac{20}{9}$, respectively, and the initial state to $|\phi^{(1)}\rangle = |0\rangle$.

In Fig. 2 we plot the mean fidelity $F^{(k)}$ [Figs. 2(a) and 2(b)] and the mean exploration parameter $W^{(k)}$ [Figs. 2(c) and 2(d)] versus k for the parameter values indicated in the figure caption. With different types of lines we depict different values of the dimensionless decay rate at zero temperature $\tilde{\Gamma}_0$. The figure illustrates the role played by the temperature in the accuracy and convergence of the protocol. For dimensionless temperature $\tilde{T} = 0.3$ [Figs. 2(a) and 2(c)], the differences between the nondissipative case $\tilde{\Gamma}_0 = 0$ (black solid line) and the dissipative cases $\tilde{\Gamma}_0 = 0.5$ (red dotted line) and $\tilde{\Gamma}_0 = 1$ (blue dashed line) are not significant. In the three cases, the asymptotic values of $F^{(k)}$ in the large iteration limit are very similar and quite close to 0.98 [see Fig. 2(a)] and $W^{(k)}$ converges to zero rather quickly [see Fig. 2(c)]. Therefore, contrary to what might be expected, temperatures much lower than $\hbar\omega/k_B$ are not required for the protocol to work. Nonetheless, for sufficiently high dimensionless temperatures, such as $\tilde{T} = 1.5$, the protocol fails in the dissipative cases, as shown in Figs. 2(b) and 2(d). On the one hand, the asymptotic values of $F^{(k)}$ in the large iteration limit are substantially less than in the absence of dissipation [see Fig. 2(b)]. On the other hand, $W^{(k)}$ does not approach zero as k increases [see Fig. 2(d)] and therefore the protocol does not converge.




FIG. 2. Dependence of (a) and (b) the mean fidelity $F^{(k)}$ and (c) and (d) the mean exploration parameter $W^{(k)}$ on the number of iterations, k . Black solid lines depict the nondissipative case $\tilde{\Gamma}_0 = 0$. The dissipative cases are indicated with red dotted lines ($\tilde{\Gamma}_0 = 0.5$) and blue dashed lines ($\tilde{\Gamma}_0 = 1$). The dimensionless temperature is (a) and (c) $\tilde{T} = 0.3$ and (b) and (d) $\tilde{T} = 1.5$. The remaining parameter values are $\tilde{\tau} = 1$, $N = 1000$, $r = 0.9$, and $p = \frac{20}{9}$.

In order to determine the range of temperatures in which the protocol works well, in Fig. 3 we plot the asymptotic values in the large iteration limit of the mean fidelity F_a and the mean exploration parameter W_a as a function of the dimensionless temperature \tilde{T} . As can be seen, there is hardly any difference between the dissipative and nondissipative cases up to $\tilde{T} \approx 0.4$. Above that temperature, the F_a values in the dissipative cases $\tilde{\Gamma}_0 = 0.5$ (red dotted line) and $\tilde{\Gamma}_0 = 1$ (blue dashed line) decrease abruptly and become substantially less than in the nondissipative case [see Fig. 3(a)]. Furthermore, there is a rapid increase in W_a to nonzero values [see Fig. 3(b)], which indicates that the protocol stops converging.




FIG. 3. Dependence of the asymptotic values in the large iteration limit of the mean fidelity, F_a , and the mean exploration parameter, W_a , on the dimensionless temperature \tilde{T} for different values of $\tilde{\Gamma}_0$. Black solid lines depict the nondissipative case $\tilde{\Gamma}_0 = 0$, while the dissipative cases are indicated with red dotted lines ($\tilde{\Gamma}_0 = 0.5$) and blue dashed lines ($\tilde{\Gamma}_0 = 1$). The remaining parameter values are $\tilde{\tau} = 1$, $N = 1000$, $r = 0.9$, and $p = \frac{20}{9}$.




FIG. 4. Dependence of the fidelities (a) $\langle +|\mathcal{E}(|+\rangle\langle +|)|+\rangle$ in Eq. (21) and (b) $\langle -|\mathcal{E}(|-\rangle\langle -|)|-\rangle$ in Eq. (22) on $\tilde{\Gamma}_0$ and \tilde{T} for $\tilde{\tau} = 1$, $N = 1000$, $r = 0.9$, and $p = \frac{20}{9}$.

To better understand the behavior observed in Fig. 3, it is worth recalling that, in the nondissipative case, the design of the protocol is based on the fact that the states $|+\rangle\langle +|$ and $|-\rangle\langle -|$ are fixed points of the quantum gate \mathcal{E} , i.e., $\mathcal{E}(|\pm\rangle\langle \pm|) = |\pm\rangle\langle \pm|$. However, in the presence of dissipation, $|+\rangle\langle +|$ and $|-\rangle\langle -|$ are no longer exact fixed points of \mathcal{E} , since the fidelities between the state $|+\rangle\langle +|$ and $\mathcal{E}(|+\rangle\langle +|)$, $\langle +|\mathcal{E}(|+\rangle\langle +|)|+\rangle$, and between the state $|-\rangle\langle -|$ and $\mathcal{E}(|-\rangle\langle -|)$, $\langle -|\mathcal{E}(|-\rangle\langle -|)|-\rangle$, are in general less than 1. Specifically, from Eq. (3), it can be shown that

$$\langle +|\mathcal{E}(|+\rangle\langle +|)|+\rangle = 1 - (1 - p_+)\gamma, \quad (21)$$

$$\langle -|\mathcal{E}(|-\rangle\langle -|)|-\rangle = 1 - p_+\gamma, \quad (22)$$

which are clearly less than 1 if $\tilde{\Gamma}_0 \neq 0$ (dissipative case). Despite this, there may be parameter values for which the fidelities $\langle +|\mathcal{E}(|+\rangle\langle +|)|+\rangle$ and/or $\langle -|\mathcal{E}(|-\rangle\langle -|)|-\rangle$ are quite close to 1. In that case, the corresponding states $|+\rangle\langle +|$ and/or $|-\rangle\langle -|$ would behave as approximate fixed points and the protocol would still be applicable. Figure 4 shows the dependence of the fidelities $\langle +|\mathcal{E}(|+\rangle\langle +|)|+\rangle$ and $\langle -|\mathcal{E}(|-\rangle\langle -|)|-\rangle$, obtained using Eqs. (21) and (22), on $\tilde{\Gamma}_0$ and \tilde{T} . As can be seen, there are regions where these fidelities are quite close to 1. In particular, Fig. 4(b) shows that $\langle -|\mathcal{E}(|-\rangle\langle -|)|-\rangle$ is quite close to 1 up to $\tilde{T} \approx 0.4$, which is the approximate value beyond which the protocol starts to fail in Fig. 3.

From Figs. 2 and 3 it may be concluded that the effect of dissipation is not very relevant for sufficiently low dimensionless temperatures. However, a more exhaustive analysis of the results reveals that it plays an important role in the protocol. Figure 5 depicts separately the mean fidelity $F_-^{(k)}$ between the state $|\phi^{(k)}\rangle$ and the ground state $|-\rangle$ [Figs. 5(a) and 5(b)] and the mean fidelity $F_+^{(k)}$ between the state $|\phi^{(k)}\rangle$ and the excited state $|+\rangle$ [Figs. 5(c) and 5(d)], defined as

$$F_{\mp}^{(k)} = \frac{1}{N} \sum_{j=1}^N |\langle \mp | \phi_j^{(k)} \rangle|. \quad (23)$$

While for $\tilde{T} = 0.3$ the differences between the dissipative and nondissipative cases are not significant when comparing




FIG. 5. Dependence of the state-specific mean fidelities (a) and (b) $F_-^{(k)}$ and (c) and (d) $F_+^{(k)}$ on the number of iterations k . Black solid lines depict the nondissipative case $\tilde{\Gamma}_0 = 0$. The dissipative cases are indicated with red dotted lines ($\tilde{\Gamma}_0 = 0.5$) and blue dashed lines ($\tilde{\Gamma}_0 = 1$). The dimensionless temperature is (a) and (c) $\tilde{T} = 0.3$ and (b) and (d) $\tilde{T} = 1.5$. The remaining parameter values are $\tilde{\tau} = 1$, $N = 1000$, $r = 0.9$, and $p = \frac{20}{9}$.

the mean fidelity $F^{(k)}$ [see Fig. 2(a)], they become rather noticeable when comparing the state-specific mean fidelities $F_-^{(k)}$ and $F_+^{(k)}$ [see Figs. 5(a) and 5(c)]. In the nondissipative case, $F_-^{(k)}$ and $F_+^{(k)}$ are almost indistinguishable [black solid lines in Figs. 5(a) and 5(c)], as is to be expected from the fact that $|+\rangle\langle +|$ and $|-\rangle\langle -|$ are equivalent fixed points of \mathcal{E} in the absence of dissipation. By contrast, the presence of dissipation substantially affects the state-specific fidelities, increasing $F_-^{(k)}$ and decreasing $F_+^{(k)}$ [red dotted lines and blue dashed lines in Figs. 5(a) and 5(c)]. This is because, for $\tilde{T} = 0.3$, $\tilde{\Gamma}_0 = 0.5$, and $\tilde{\Gamma}_0 = 1$, $|-\rangle\langle -|$ is an approximate fixed point of \mathcal{E} but $|+\rangle\langle +|$ is not (see Fig. 4). As the temperature increases, this difference becomes less appreciable, as can be seen in Figs. 5(b) and 5(d).

One important aspect not considered so far is the dependence of the protocol results on the evolution time τ . As mentioned in Sec. II, in the absence of dissipation, i.e., for $\Gamma_0 = 0$, the action of E on A is represented by the quantum channel $\mathcal{E}(\rho_0) = e^{-i\tau H/\hbar} \rho_0 e^{i\tau H/\hbar}$. For the case of the Hamiltonian H in Eq. (17), this quantum channel has the explicit form

$$\begin{aligned} \mathcal{E}(\rho_0) = & \frac{1}{2}(\rho_0 + \sigma_x \rho_0 \sigma_x) + \frac{\cos(\tilde{\tau})}{2}(\rho_0 - \sigma_x \rho_0 \sigma_x) \\ & + \frac{i \sin(\tilde{\tau})}{2}[\rho_0, \sigma_x], \end{aligned} \quad (24)$$

which is a periodic function of the dimensionless evolution time $\tilde{\tau}$ with period 2π . This periodicity is clearly visible in Fig. 6(a), where the dependence of the asymptotic values in the large iteration limit of the mean fidelity F_a on $\tilde{\tau}$ in the nondissipative case is depicted by a solid line. However, the periodicity is not evident in the behavior of W_a , which always remains close to 0 and indistinguishable from the abscissa axis, as shown in Fig. 6(b). Of particular interest is the case in which $\tilde{\tau}$ is an integer multiple of 2π . In this case, it

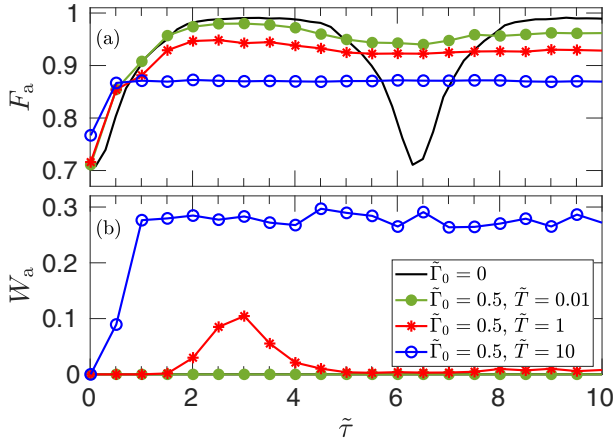


FIG. 6. Dependence of the asymptotic values in the large iteration limit of (a) the mean fidelity F_a and (b) the mean exploration parameter W_a on the dimensionless evolution time $\tilde{\tau}$ for different values of $\tilde{\Gamma}_0$ and \tilde{T} . Black solid lines depict the nondissipative case $\tilde{\Gamma}_0 = 0$. In (b) the black solid line is indistinguishable from the abscissa axis. The dissipative cases are indicated with green closed circles ($\tilde{\Gamma}_0 = 0.5$ and $\tilde{T} = 0.01$), red stars ($\tilde{\Gamma}_0 = 0.5$ and $\tilde{T} = 1$), and blue open circles ($\tilde{\Gamma}_0 = 0.5$ and $\tilde{T} = 10$). The remaining parameter values are $N = 1000$, $r = 0.9$, and $p = \frac{20}{9}$.

follows from Eq. (24) that E does not modify the state of A and therefore no learning is possible. Thus, for this case the mean fidelity $F^{(k)}$ is independent of the number of iterations k and consequently $F_a = F^{(1)} = 1/\sqrt{2} \approx 0.71$, as it can be observed in Fig. 6(a) (black solid line).

In the presence of dissipation, i.e., for $\Gamma_0 \neq 0$, the quantum channel $\mathcal{E}(\rho_0)$ in Eq. (3) ceases to be periodic in $\tilde{\tau}$ due to the exponential dependence of the parameter γ on the evolution time τ (see Sec. II). As a consequence, the oscillations observed in Fig. 6(a) for the nondissipative case (black solid line) are notably attenuated and, beyond a certain value of $\tilde{\tau}$, F_a becomes almost independent of $\tilde{\tau}$ [see green closed circles, red stars, and blue open circles in Fig. 6(a)]. The effect of temperature on the behavior of F_a is quite revealing. The long $\tilde{\tau}$ values of F_a for the dissipative case decrease as the temperature increases, as expected. Nonetheless, the F_a values are substantially larger than those for the nondissipative case for $\tilde{\tau}$ values around the integer multiples of 2π . Therefore, under these circumstances, dissipation plays a positive role in the protocol fidelity. Note that for not too high temperatures, the convergence of the protocol gauged by W_a is still guaranteed around the above-mentioned $\tilde{\tau}$ values [see green closed circles and red stars in Fig. 6(b)].

IV. CONCLUSION

In this study we investigated the impact of thermal dissipation on a protocol for quantum reinforcement learning.

To achieve this, we introduced a Lindblad dynamics for the density operator of a two-state system in contact with a thermal bath. First, we conducted a comprehensive theoretical analysis of the protocol and established the relevant quantities that characterize it. Subsequently, we performed a numerical analysis to demonstrate the influence of thermal dissipation on the protocol's performance. It is important to understand how thermal dissipation affects quantum reinforcement learning because machine learning algorithms are typically not implemented in machines that operate at absolute zero temperature and therefore thermal effects cannot be avoided in real-world situations.

The fundamental difference between the methodology presented in this study and that considered in Ref. [15] lies in the way in which the agent interacts with the environment to gather information from it. In Ref. [15] such an interaction is free from dissipation and is therefore described by the unitary transformation $\rho_\tau = U \rho_0 U^\dagger$. However, in this study, the interaction involves a certain degree of thermal dissipation which causes it to deviate from unitarity [see Eq. (3)]. This makes the description more realistic in our case as experiments always take place at nonzero temperatures and therefore in the presence of some degree of thermal dissipation.

The following are the main conclusions of this work.

(i) For sufficiently low temperatures, dissipation does not necessarily have a negative effect on the accuracy of the protocol gauged by the mean fidelity $F^{(k)}$. In fact, our analysis shows that there are ranges of parameter values for which the dissipative protocol performs better than the nondissipative one, as can be seen in Fig. 6. These results might be of interest for the experimental implementation of this type of protocol, since sometimes the presence of dissipation is unavoidable in an actual experiment.

(ii) Dissipation is particularly relevant when evaluating the state-specific mean fidelities $F_-^{(k)}$ and $F_+^{(k)}$. The influence of dissipation on these fidelities depends dramatically on which specific state is considered. While in the nondissipative case $F_-^{(k)}$ and $F_+^{(k)}$ are almost indistinguishable, in the dissipative case they can be quite different. Specifically, for sufficiently low temperatures, the mean fidelity with respect to the ground state, $F_-^{(k)}$, is much higher than the one obtained in the absence of dissipation, whereas the mean fidelity with respect to the excited state, $F_+^{(k)}$, is much lower. Therefore, if in an actual experiment one were interested in specifically computing the ground state, the presence of dissipation might be useful.

ACKNOWLEDGMENTS

We acknowledge funding from the Junta de Andalucía (Grants No. P20-00617 and No. US-1380840) and from the Spanish Ministry of Science, Innovation, and Universities under Grants No. PID2019-104002GB-C21 and No. PID2019-104002GB-C22.

[1] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, *Nature (London)* **549**, 195 (2017).
 [2] L. Lamata, *Mach. Learn.: Sci. Technol.* **1**, 033002 (2020).

[3] A. W. Harrow, A. Hassidim, and S. Lloyd, *Phys. Rev. Lett.* **103**, 150502 (2009).
 [4] S. Lloyd, M. Mohseni, and P. Rebentrost, *Nat. Phys.* **10**, 631 (2014).

- [5] P. Rebentrost, M. Mohseni, and S. Lloyd, *Phys. Rev. Lett.* **113**, 130503 (2014).
- [6] R. K. Nath, H. Thapliyal, and T. S. Humble, *SN Comput. Sci.* **2**, 365 (2021).
- [7] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *New J. Phys.* **18**, 023023 (2016).
- [8] D. Dong, C. Chen, H. Li, and T.-J. Tarn, *IEEE Trans. Syst. Man Cybern. B* **38**, 1207 (2008).
- [9] G. D. Paparo, V. Dunjko, A. Makmal, M. A. Martin-Delgado, and H. J. Briegel, *Phys. Rev. X* **4**, 031002 (2014).
- [10] V. Dunjko, J. M. Taylor, and H. J. Briegel, *Phys. Rev. Lett.* **117**, 130501 (2016).
- [11] M. Bukov, *Phys. Rev. B* **98**, 224305 (2018).
- [12] M. Bukov, A. G. R. Day, D. Sels, P. Weinberg, A. Polkovnikov, and P. Mehta, *Phys. Rev. X* **8**, 031086 (2018).
- [13] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, *Phys. Rev. X* **8**, 031084 (2018).
- [14] Y.-P. Liu, Q.-S. Jia, and X. Wang, *IFAC-PapersOnLine* **55**, 132 (2022).
- [15] F. Albarrán-Arriagada, J. C. Retamal, E. Solano, and L. Lamata, *Mach. Learn.: Sci. Technol.* **1**, 015002 (2020).
- [16] L. Hu, S.-H. Wu, W. Cai, Y. Ma, X. Mu, Y. Xu, H. Wang, Y. Song, D.-L. Deng, C.-L. Zou, and L. Sun, *Sci. Adv.* **5**, eaav2761 (2019).
- [17] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, *Nature (London)* **567**, 209 (2019).
- [18] S. Yu, F. Albarrán-Arriagada, J. C. Retamal, Y.-T. Wang, W. Liu, Z.-J. Ke, Y. Meng, Z.-P. Li, J.-S. Tang, E. Solano, L. Lamata, C.-F. Li, and G.-C. Guo, *Adv. Quantum Technol.* **2**, 1800074 (2019).
- [19] V. Saggio, B. E. Asenbeck, A. Hamann, T. Strömberg, P. Schiansky, V. Dunjko, N. Friis, N. C. Harris, M. Hochberg, D. Englund, S. Wölk, H. J. Briegel, and P. Walther, *Nature (London)* **591**, 229 (2021).
- [20] M. Spagnolo, J. Morris, S. Piacentini, M. Antesberger, F. Massa, A. Crespi, F. Ceccarelli, R. Osellame, and P. Walther, *Nat. Photon.* **16**, 318 (2022).
- [21] H.-Y. Huang, M. Broughton, J. Cotler, S. Chen, J. Li, M. Mohseni, H. Neven, R. Babbush, R. Kueng, J. Preskill, and J. R. McClean, *Science* **376**, 1182 (2022).
- [22] N. H. Nguyen, E. C. Behrman, and J. E. Steck, *Quantum Mach. Intell.* **2**, 1 (2020).
- [23] C. Lu, S. Kundu, A. Arunachalam, and K. Basu, *Proceedings of the IEEE 15th Dallas Circuit and Systems Conference* (IEEE, Piscataway, 2022), pp. 1–2.
- [24] L. Innocenti, L. Banchi, A. Ferraro, S. Bose, and M. Paternostro, *New J. Phys.* **22**, 065001 (2020).
- [25] A. Youssry, R. J. Chapman, A. Peruzzo, C. Ferrie, and M. Tomamichel, *Quantum Sci. Technol.* **5**, 025001 (2020).
- [26] S. Shrapnel, F. Costa, and G. Milburn, *Int. J. Quantum Inf.* **16**, 1840010 (2018).
- [27] I. A. Luchnikov, S. V. Vintskevich, D. A. Grigoriev, and S. N. Filippov, *Phys. Rev. Lett.* **124**, 140502 (2020).
- [28] A. A. Melnikov, L. E. Fedichkin, R.-K. Lee, and A. Alodjants, *Adv. Quantum Technol.* **3**, 1900115 (2020).
- [29] A. A. Melnikov, L. E. Fedichkin, and A. Alodjants, *New J. Phys.* **21**, 125002 (2019).
- [30] S. S. Kalantre, J. P. Zvolak, S. Ragole, X. Wu, N. M. Zimmerman, M. D. Stewart, and J. M. Taylor, *npj Quantum Inf.* **5**, 6 (2019).
- [31] G. Torlai, B. Timar, E. P. L. van Nieuwenburg, H. Levine, A. Omran, A. Keesling, H. Bernien, M. Greiner, V. Vuletić, M. D. Lukin, R. G. Melko, and M. Endres, *Phys. Rev. Lett.* **123**, 230504 (2019).
- [32] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, *Nat. Phys.* **14**, 447 (2018).
- [33] G. Liu, M. Chen, Y.-X. Liu, D. Layden, and P. Cappellaro, *Mach. Learn.: Sci. Technol.* **1**, 015003 (2020).
- [34] I. Agresti, N. Viggianiello, F. Flamini, N. Spagnolo, A. Crespi, R. Osellame, N. Wiebe, and F. Sciarrino, *Phys. Rev. X* **9**, 011013 (2019).
- [35] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. (MIT Press, Cambridge, 2018).
- [36] L. Busoniu, R. Babuska, B. D. Schutter, and D. Ernst, *Reinforcement Learning and Dynamic Programming Using Function Approximators*, 1st ed. (CRC, Boca Raton, 2010).
- [37] H.-P. Breuer and F. Petruccione, *Theory of Open Quantum Systems* (Oxford University Press, Oxford, 2003).
- [38] M. A. Nielsen and I. L. Chuang, *Quantum Computing and Quantum Information* (Cambridge University Press, Cambridge, 2000).
- [39] M. L. Puterman, in *Stochastic Models*, edited by D. P. Heyman and M. J. Sobel, Handbooks in Operations Research and Management Science Vol. 2 (Elsevier, Amsterdam, 1990), pp. 331–434.
- [40] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (Wiley, New York, 2014).