


Deep reinforcement learning for radiative heat transfer optimization problems

E. Ortiz-Mansilla^{1,2,†}, J.J. García-Esteban^{1,2,†}, J. Bravo-Abad^{1,2} and J.C. Cuevas^{1,2,*}

¹*Departamento de Física Teórica de la Materia Condensada, Universidad Autónoma de Madrid, Madrid 28049, Spain*

²*Condensed Matter Physics Center (IFIMAC), Universidad Autónoma de Madrid, Madrid 28049, Spain*

 (Received 28 August 2024; revised 11 October 2024; accepted 1 November 2024; published 25 November 2024)

Reinforcement learning is a subfield of machine learning that is having a huge impact in the different conventional disciplines, including physical sciences. Here, we show how reinforcement learning methods can be applied to solve optimization problems in the context of radiative heat transfer. We illustrate their use with the optimization of the near-field radiative heat transfer between multilayer hyperbolic metamaterials. Specifically, we show how this problem can be formulated in the language of reinforcement learning and tackled with a variety of algorithms. We show that these algorithms allow us to find solutions that outperform those obtained using physical intuition. Overall, our work shows the power and potential of reinforcement learning methods for the investigation of a wide variety of problems in the context of radiative heat transfer and related topics.

DOI: [10.1103/PhysRevApplied.22.054071](https://doi.org/10.1103/PhysRevApplied.22.054071)

I. INTRODUCTION

Thermal radiation is a ubiquitous physical phenomenon whose understanding is of critical importance for many different areas of science and engineering [1–3]. The field of radiative heat transfer is enjoying a revival due to various recent advances [4]. Maybe the most notable one is the demonstration that the near-field radiative heat transfer (NFRHT) between two closely placed bodies can largely overcome the blackbody limit set by the Stefan-Boltzmann law. This was predicted in the early 1970s [5] and it has been verified in recent years in a large variety of systems with the help of novel experimental techniques [4,6,7]. This effect originates from the fact that, when two objects are separated by a distance smaller than the thermal wavelength λ_{Th} (about 10 μm at room temperature), the radiative heat flux can be greatly enhanced by the additional contribution of evanescent waves—which is not considered in the Stefan-Boltzmann law. Near-field thermal radiation has opened new possibilities and holds the promise to have a notable impact in different technologies such as heat-assisted magnetic recording [8], scanning thermal microscopy [9–11], coherent thermal sources [12,13], near-field-based thermal management [4,7], or thermophotovoltaics [14].

NFRHT is by no means the only breakthrough in the field of thermal radiation in recent times. Thus, for instance, it has been shown that nanophotonic structures,

where at least one of the structural features is at sub-wavelength scale, can have thermal radiation properties that differ drastically from those of conventional thermal emitters [15]. This has led to the development and improvement of energy applications such as daytime passive radiative cooling [16,17], thermal radiative textiles [18,19], radiative cooling of solar cells [20], or thermophotovoltaic cells [21]. On a more fundamental level, another remarkable discovery has been the possibility of overcoming the far-field limits set by Planck’s law in the context of the thermal emission and the radiative heat transfer between subwavelength objects [22–24].

At this stage, the physical mechanisms of radiative heat transfer in the different regimes are relatively well understood and the interest is now shifting towards the optimization and design of novel thermal devices. This process is being mainly assisted by physical intuition and standard numerical optimization methods. Thus, for instance, in the context of NFRHT, many different analytical upper bounds have been put forward to establish the limits of near-field thermal radiation [7,25]. These bounds are extremely ingenious, but often lack the ability to guide in practice the fabrication of actual structures. On the other hand, conventional numerical optimization techniques, such as Bayesian or topology optimization [26], are also being routinely used in the field.

At the same time, the impressive achievements of machine learning techniques in different engineering areas have motivated many researchers to pursue a data-driven approach to investigate a plethora of problems in conventional science disciplines, including physical sciences

*Contact author: juancarlos.cuevas@uam.es

†These authors contributed equally to this work.

[27–29]. Radiative heat transfer is not an exception and in recent years different groups have applied various machine learning techniques to address key problems in this field. Most of the work thus far has been carried out with the help of artificial neural networks (ANNs) and deep learning algorithms. Thus, for instance, we have shown how ANNs can be used to tackle optimization and inverse design problems in the context of NFRHT, passive radiative cooling, and thermal emission of subwavelength objects [30]. There has also been tremendous activity in the context of deep learning–aided design and optimization of thermal metamaterials; for a recent review, see Ref. [31]. However, reinforcement learning, another subfield of machine learning, has barely been used in modern radiative heat transfer problems, with notable exceptions [32]. Reinforcement learning (RL) is much closer to the layman’s view of artificial intelligence and it deals with problems concerning sequential decision making [33]. In RL, an agent learns via the interaction with an environment from which it receives feedback to make good decisions towards a given goal, such as the optimization of a physical process or the inverse design of a device.

In this work we want to fill this gap and show how RL can be used to tackle optimization problems in the context of radiative heat transfer. To be precise, we illustrate the core ideas with a problem related to the optimization of NFRHT between multilayer hyperbolic metamaterials. We show how this type of problem can be framed in the language of RL and how different RL algorithms can be implemented to address them. In particular, we critically assess the advantages and disadvantages of the different methods to help new users of RL to select the most convenient algorithm for a given application. The methods presented in this work can be straightforwardly applied to a large variety of problems in thermal radiation science and related fields.

The rest of the manuscript is organized as follows. In Sec. II, we briefly introduce the topic of RL for nonexperts to make our contribution more self-contained. In Sec. III, we present the system and problem that we have chosen to illustrate the use of RL in the context of thermal radiation problems, namely, the optimization of NFRHT between multilayer hyperbolic metamaterials. Then, Sec. IV is devoted to the description of the main results of this work obtained with different RL algorithms. We have organized those results according to the RL algorithm employed and we also provide detailed descriptions of such algorithms. Finally, we present some additional discussions and summarize our main conclusions in Sec. V.

II. REINFORCEMENT LEARNING: A BRIEF REMINDER

In this section we provide a brief introduction to RL following Ref. [34]. This will allow us to set the language and




FIG. 1. The reinforcement learning control loop diagram.

make the manuscript more self-contained. Readers familiar with RL can safely skip this section.

RL is a subfield of machine learning that aims at solving sequential decision-making problems. Many problems can be formulated in this way, including those concerning the optimization of systems, devices, and processes in the physical sciences. To solve a problem within RL, we begin by defining a goal. Then, an algorithm takes actions and gets information about the external world on how well the goal is being achieved. To that end, we normally need to take many actions in a sequential fashion, where each action modifies the world around us. We observe the changes in the world and, with the help of the feedback we receive, we decide on the next action to take.

The RL formulation of the process described above is the following. RL problems are formulated as a system that comprises an *agent* and an *environment*, the world surrounding the agent. The environment produces information that allows us to describe the *state* of the system, while the agent interacts with the environment by observing the state and selecting an *action*. The environment accepts the action and transitions into a new state, which is then observed by the agent to select a new action. In doing so, the environment also returns a *reward* to the agent, which is used by the agent to select future actions. When the cycle of state \rightarrow action \rightarrow next state and reward is completed, we say that one *time step* has passed. This cycle is repeated until the environment terminates, for example, when the problem is solved. This process is summarized in the control loop diagram of Fig. 1.

Going deeper with the formulation of the process, a *policy* in RL is the agent’s action-producing function that maps states to actions. As previously indicated, actions change the environment and affect what an agent observes and does next, which can be viewed as a sequential decision-making process that evolves in time. This RL process is driven by an *objective*, which is defined as the sum of rewards received from the environment. The agent aims at maximizing the objective by selecting good actions and learns to do this by interacting with the environment according to an optimizing policy in a trial-and-error process, which uses rewards to reinforce good actions and penalize bad ones. Therefore, the signals exchanged between agent and environment are (s_t, a_t, r_t) , which stand

for state, action, and reward, respectively, and where t denotes the time step in which these signals occurred. The defined tuple (s_t, a_t, r_t) is referred to as an *experience*, which is the basic unit of information describing an RL system. The control loop is then repeated forever or terminated by reaching either a terminal state or a maximum time step $t = T$. The time frame from $t = 0$ to the terminal state [either a defined state(s) or a maximum state] is called an *episode*. In addition, the sequence of experiences over an episode, $\tau = (s_0, a_0, r_0), (s_1, a_1, r_1), \dots$, is known as a *trajectory*. An agent typically needs many episodes to learn a good policy.

In a more formal way, we can describe states, actions, and rewards as (i) $s_t \in \mathcal{S}$, where \mathcal{S} is the state space, (ii) $a_t \in \mathcal{A}$, where \mathcal{A} is the action space, and (iii) $r_t = \mathcal{R}(s_t, a_t, s_{t+1})$, where \mathcal{R} is the reward function. Here, the state space \mathcal{S} is the set of all possible states in an environment. It can be defined as integers, real numbers, vectors, etc. Similarly, the action space \mathcal{A} is the set of all possible actions. It is commonly defined as either a scalar or a vector. The reward function $\mathcal{R}(s_t, a_t, s_{t+1})$, for its part, assigns a real number (positive or negative) to each transition. The state space, action space, and reward function are determined by the environment.

Let us now consider how an environment transitions from one state to the next using the *transition function*. In RL, a transition function is formulated as a Markov decision process (MDP), which means in practice that one assumes that the transition to the next state s_{t+1} depends only on the previous state s_t and action a_t . This is known as the *Markov property* and can be mathematically formulated as

$$s_{t+1} \sim P(s_{t+1}|s_t, a_t), \quad (1)$$

which means that the next state s_{t+1} is sampled from a probability distribution $P(s_{t+1}|s_t, a_t)$.

With this new ingredient, we can now compile all MDP elements, i.e., \mathcal{S} , \mathcal{A} , $\mathcal{R}(\cdot)$, $P(\cdot)$, where we recall that \mathcal{S} is the set of states, \mathcal{A} is the set of actions, $P(s_{t+1}|s_t, a_t)$ is the transition function of the environment, and $\mathcal{R}(s_t, a_t, s_{t+1})$ is the reward function. Let us remark that RL algorithms tackled in this work are *model-free*, that is, the agents have access to neither the transition function, $P(s_{t+1}|s_t, a_t)$, nor the reward function, $\mathcal{R}(s_t, a_t, s_{t+1})$. The only way in which an agent gets information about these functions is through the states, actions, and rewards it actually experiences in the environment.

As previously indicated, to formulate an RL problem, it is necessary to formalize the objective that the agent is intended to maximize. For this purpose, we first define the *return* $G(\tau)$ using a trajectory from an episode,

$$G(\tau) = r_0 + \gamma r_1 + \gamma^2 r_2 + \dots + \gamma^T r_T = \sum_{t=0}^T \gamma^t r_t, \quad (2)$$

i.e., as a discounted sum of the rewards in a trajectory, where $\gamma \in [0, 1]$ is the *discount factor*. The discount factor is an important parameter that changes the way future rewards are considered. The smaller γ , the less weight is given to rewards in future time steps.

On the other hand, the *objective* $J(\pi)$ is simply defined as the expectation of the returns over many trajectories evaluated with a given policy π , i.e.,

$$J(\pi) = \mathbb{E}_{\tau \sim \pi}[G(\tau)] = \mathbb{E}_{\tau} \left[\sum_{t=0}^T \gamma^t r_t \right]. \quad (3)$$

The expectation accounts for stochasticity in the actions and the environment.

A key question in RL concerns what an agent should learn. There are three basic properties that can be useful to an agent: (i) a policy, (ii) a value function, and (iii) an environment model. First, if we recall, policy π is that which maps states to actions, which can be formalized with the notation $a \sim \pi(s)$. A policy can be stochastic and, therefore, we can write this as $\pi(a|s)$ to denote the probability of an action a given a state s .

The *value functions* provide information about the objective. They help an agent to understand how good the states and available actions are in terms of the expected future return, allowing one to determine a policy from this information. There are two types of value functions, defined as

$$V^\pi(s) = \mathbb{E}_{s_t=s, \tau \sim \pi} \left[\sum_{t=t'}^T \gamma^t r_t \right], \quad (4)$$

$$Q^\pi(s, a) = \mathbb{E}_{s_t=s, a_t=a, \tau \sim \pi} \left[\sum_{t=t'}^T \gamma^t r_t \right]. \quad (5)$$

The *state-value function* $V^\pi(s)$ in Eq. (4) evaluates the quality of a state. It measures the expected return from being in state s , assuming that the agent continues to act according to its current policy π . It is worth noting that the return $G(\tau) = \sum_{t=t'}^T \gamma^t r_t$ is measured from the current state to the end of an episode. The *action-value function* $Q^\pi(s, a)$ of Eq. (5) evaluates how good a state-action pair is. It measures the expected return from taking action a in state s , assuming that the agent continues to act according to its current policy, π .

Finally, an environment model is summarized in the transition function $P(s_{t+1}|s_t, a_t)$ that provides information about the environment. If an agent learns this function, it is able to predict the next state s_{t+1} that the environment will transition into after taking action a in state s . Often, good models of the environment are not available and in this work we will not make use of this type of function or the corresponding algorithms.

In RL, an agent learns a function of any of the previous properties to decide what actions to take with the goal of maximizing the objective. In most practical problems the different spaces—state, action, etc.—are so large that the key functions need to be approximated. Currently, the most popular methods to approximate these functions are based on deep neural networks, which gives rise to the concept of *deep reinforcement learning*. This is the method of choice in this work.

On the other hand, according to the three primary learnable functions in RL (see above), there are three major families of deep RL algorithms—*policy-based*, *value-based*, and *model-based* methods that learn policies, value functions, and models, respectively. In Sec. IV below, we present the main results of this work organized according to the corresponding RL algorithm employed and we also include a brief description of the main characteristics of every used algorithm.

III. OPTIMIZING NFRHT BETWEEN MULTILAYER HYPERBOLIC METAMATERIALS

A. Physical problem

In this section we describe the specific problem that we have selected to illustrate the use of RL in the context of radiative heat transfer, namely, the optimization of the near-field radiative heat transfer between multilayer hyperbolic metamaterials [30,35].

As discussed in the Introduction, a major breakthrough in recent years in the field of thermal radiation has been the confirmation of the possibility to overcome the Stefan-Boltzmann law for the radiative heat transfer between two bodies by bringing them sufficiently close [5]. This physical phenomenon is due to the fact that in the near-field regime, bodies can exchange radiative heat via evanescent waves. This type of contribution is not considered in the Stefan-Boltzmann law and dominates the NFRHT for sufficiently small separations [4,7,36]. Different strategies have recently been proposed to further enhance NFRHT. One of the most prominent ones makes use of multiple surface modes that appear in multilayer structures where dielectric and metallic layers are alternated to give rise to the so-called *hyperbolic metamaterials* [37–46]. The hybridization of surface modes in different metal-dielectric interfaces can lead to a great enhancement of the NFRHT, as compared to the case of two infinite parallel plates [44].

Following Ref. [44], here we consider the radiative heat transfer between two identical multilayer structures separated by a gap d_0 ; see Fig. 2(a). Each thermal reservoir contains N_l total layers alternating between a metallic layer with a permittivity ϵ_m and a lossless dielectric layer of permittivity ϵ_d . The thickness of layer i is denoted d_i . The dielectric layers are set to vacuum ($\epsilon_d = 1$) and the metallic layers are described by a permittivity given by a Drude

model: $\epsilon_m(\omega) = \epsilon_\infty - \omega_p^2/[\omega(\omega + i\gamma)]$ with ϵ_∞ the permittivity at infinite frequency, ω_p the plasma frequency, and γ the damping rate. From now on, we set $\epsilon_\infty = 1$, $\omega_p = 2.5 \times 10^{14}$ rad/s, and $\gamma = 1 \times 10^{12}$ rad/s. With this choice of parameters, the surface plasmon frequency is similar to the surface phonon-polariton frequency of the interface between SiC and a vacuum.

We describe the NFRHT between the hyperbolic metamaterials within the theory of fluctuational electrodynamics [47,48]. In this system, the NFRHT is dominated by TM- or p -polarized evanescent waves and the heat transfer coefficient (HTC) between the two bodies, i.e., the linear radiative thermal conductance per unit of area, is given by [49]

$$h = \frac{\partial}{\partial T} \int_0^\infty \frac{d\omega}{2\pi} \Theta(\omega, T) \int_{\omega/c}^\infty \frac{dk}{2\pi} k \tau_p(\omega, k), \quad (6)$$

where T is the temperature, $\Theta(\omega, T) = \hbar\omega/(e^{\hbar\omega/k_B T} - 1)$ is the mean thermal energy of a mode of frequency ω , k is the magnitude of the wave vector parallel to the surface planes, and $\tau_p(\omega, k)$ is the transmission (between 0 and 1) of the p -polarized evanescent modes given by

$$\tau_p(\omega, k) = \frac{4 [\text{Im} \{r_p(\omega, k)\}]^2 e^{-2q_0 d_0}}{|1 - r_p(\omega, k)|^2 e^{-2q_0 d_0}}. \quad (7)$$

Here, $r_p(\omega, k)$ is the Fresnel reflection coefficient of the p -polarized evanescent waves from the vacuum to one of the bodies and $q_0 = \sqrt{k^2 - \omega^2/c^2}$ ($\omega/c < k$) is the wave-vector component normal to the layers, in vacuum. The Fresnel coefficient needs to be computed numerically and we have done it by using the scattering matrix method described in Ref. [50]. In our numerical calculations of the HTC we also took into account the contribution of s -polarized modes, but it turns out to be negligible for the gap sizes explored in this work.

The interest in the NFRHT in these multilayer structures resides in the fact that the heat exchange in this regime is dominated by surface modes that can be tuned by playing with the layer thicknesses. In the case of two parallel plates made of a Drude metal, the NFRHT is dominated by the two cavity surface modes resulting from the hybridization of the surface plasmon polaritons (SPPs) of the two metal-vacuum interfaces [44]. These two cavity modes give rise to two near-unity lines in the transmission function $\tau_p(\omega, k)$. If we introduce more internal layers, we can have additional NFRHT contributions from surface states at multiple surfaces. This is illustrated in Fig. 2 for the case of $N = 16$ active layers with $d_i = 5$ nm and a gap size $d_0 = 10$ nm. Apart from the active layers, both reservoirs contain an additional 5-nm-thick metallic layer to ensure that the outer surfaces feature SPPs and to properly define the gap, as well as a semi-infinite metallic substrate on the

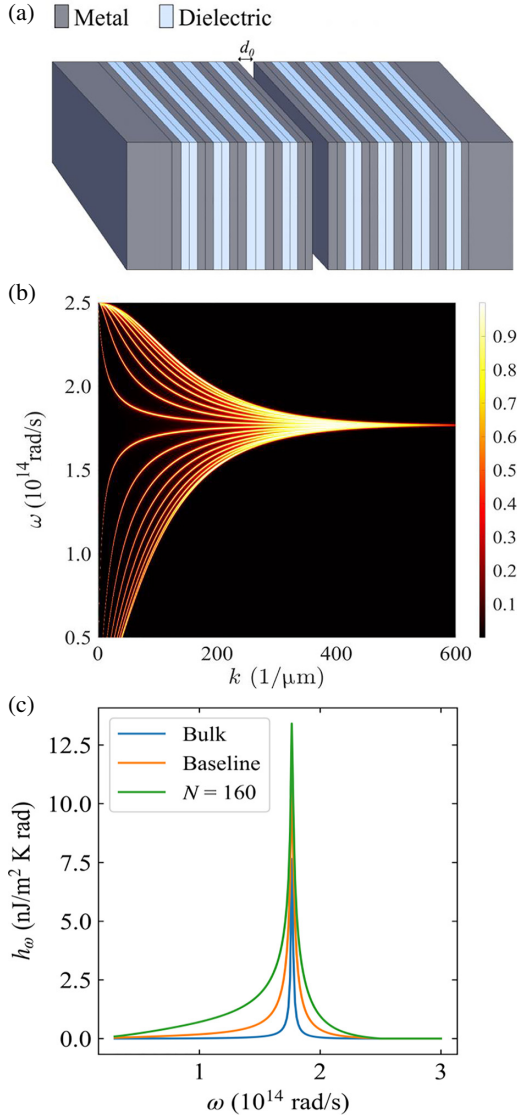


FIG. 2. (a) Schematic representation of the physical system under study. It features two identical hyperbolic metamaterials comprising alternating metallic (gray) and dielectric (blue) layers. Both reservoirs have infinitely extended layers and are separated by a distance $d_0 = 10$ nm. Each layer has a thickness of 5 nm and both subsystems are backed by a metallic substrate. (b) Transmission of evanescent waves as a function of the frequency (ω) and the parallel wave vector (k) for the periodic structure of panel (a) composed of 16 active layers per subsystem. (c) The corresponding spectral heat transfer coefficient h_ω at room temperature ($T = 300$ K) as a function of the frequency, labeled baseline in the legend. The result is compared to that of two metallic plates (bulk) with the same gap and a system like the baseline structure, but with 160 layers ($N = 160$), which mimics an infinite periodic system.

other side. This example, in which we have in practice a periodic structure with eight physical layers (four metallic and four dielectric layers) with thickness $d_i = 10$ nm [see Fig. 2(a)], exhibits multiple near-unity resonances in

the transmission function $\tau_p(\omega, k)$ [see Fig. 2(b)]. These contributions resulting from additional surface states originating from internal layers lead to a great enhancement of the NFRHT as compared to the bulk system (two parallel metallic plates) in a wide range of gap values [44]. This is illustrated in Fig. 2(c) where we show the spectral HTC, h_ω , defined as the HTC per unit of frequency: $h = \int_0^\infty h_\omega d\omega$ for both the system pictured in Fig. 2(a) (labeled baseline) and the bulk system with the same gap $d_0 = 10$ nm.

Our concrete goal is to maximize the HTC between these two hyperbolic metasurfaces by finding the optimal configuration of alternating dielectric and metallic layers (number and thickness). Theoretically, that maximum HTC is achieved when an infinite number of alternating layers is arranged so as to form ideal periodic one-dimensional photonic crystals, where the NFRHT is determined by the corresponding Bloch waves [44]. We illustrate this fact in Fig. 2(c) where we also show the result for the spectral HTC for a system like the baseline in panel (a), but for $N = 160$, which mimics the infinite periodic system. Obviously, in reality truly periodic systems cannot be fabricated and one has to deal with systems with a finite number of layers. It is then when physical intuition faces limitations and machine learning techniques become helpful to solve these optimization problems, as we show in this work. For concreteness, in our study we keep the gap size fixed at $d_0 = 10$ nm, the total active thickness of the multilayer areas, and assume room temperature ($T = 300$ K). We also assume the two multilayer systems to be identical since any asymmetry tends to reduce the HTC.

B. RL formulation of the optimization problem

We now describe how we tackle our optimization problem in the spirit of RL, which requires us to formulate it as a sequential decision-making problem. The physical system is composed of two identical layered structures that, unless we state otherwise, contain 16 active layers with a thickness of 5 nm. The two subsystems are separated by a gap $d_0 = 10$ nm, each having a semi-infinite metallic substrate.

Irrespective of the employed algorithm, we define the central RL concepts as follows.

(1) *Goal*: our goal is to maximize the HTC via the modification of the layer configuration.

(2) *State*: each state describes a layer configuration. We define the material by an integer label, 0 for the dielectric and 1 for the metal. Thus, a state is represented by a vector of 0s and 1s with 16 components, each representing one of the 5-nm-thick layers.

(3) *Action*: the action space is an ensemble of two decisions made concurrently, namely, which layer to study

and what material to consider for it. This includes the possibility for the configuration to remain unchanged.

(4) *Reward*: the reward is the HTC corresponding to the next state in units of 10^5 W/m²K. Thus, for this problem, $\mathcal{R} = \mathcal{R}(s_{t+1})$ only. For better performance, we consider as a baseline for the reward values the HTC of our physically intuitive “best-guess” configuration, which corresponds to the periodic photonic crystal shown in Fig. 2(a). Any positive reward implies that we have found a higher HTC value.

(5) *Episode termination*: we impose an episodic formulation by defining a fixed number of actions taken in a trajectory before resetting to the initial state, containing all 0s (all dielectric layers except the 5-nm-thick metallic layer defining the gap). This enables the network to perform much more training on known states, and to finish the optimization in an acceptable number of steps. We take the length of an episode to be twice or 4 times the number of layers of an state, so any existing state is potentially reached comfortably.

IV. RESULTS

In this section we describe the main results obtained for the optimization of the NFRHT between the multilayer hyperbolic structures described in the previous section. For didactic reasons, we present separately the results obtained with the different RL algorithms and in every subsection we describe the basics of the corresponding method alongside a discussion of the peculiarities concerning their application to our problem.

A. Value-based algorithms: SARSA algorithm, deep Q -learning, and extensions

As value-based algorithms are historically the most widely used and discussed in RL [33], we first address their formulation and usage. Value-based algorithms are based on two core ideas. The first one is *temporal difference* (TD) learning, which is an alternative to the use of Monte Carlo sampling for gathering experiences from an environment (see Sec. IV B below) to estimate state or state-action values. The key idea in TD learning is that state or state-action values are defined recursively, that is, their values in a given time step are defined in terms of the values in the next time step. This makes TD learning a useful method for backing up the reward information from later to earlier steps through time. As state-action-value functions represent an expectation over different trajectories, this leads to the display of a lower variance than Monte Carlo sampling.

The second idea has to do with the famous *exploration-exploitation trade-off* in RL. When the agent is learning an estimate of the state or state-action values, the usage of this estimate can lead to better returns (exploitation). However, if one always selects actions based on current values, which might be far from the optimal ones, this would

lead to deterministic behavior that can prevent the agent from discovering better unknown actions (exploration). This exploration-exploitation trade-off is a key challenge in RL and it can be addressed by employing *stochastic policies*, where the exploration can be distributed along all the training and, as the estimation gets better, it gradually shifts closer to a deterministic policy. An example of an stochastic policy, which will be used in this work, is the ε -greedy policy, where the agent explores with a probability of ε and exploits with a probability of $1 - \varepsilon$.

1. SARSA algorithm

The *state-action-reward-state-action* (SARSA) algorithm is one of the oldest RL value-based algorithms and, despite its limitations (see below), it is convenient to start by describing its use for our problem. This algorithm is based on the estimation of the action-value function or the Q function. It employs TD learning to produce the target state-action values or Q values, from now on denoted Q_{tar} . Therefore, it combines the reward given by the environment, r_t , with the Q -value estimates of the next state that approximate the remaining part of the expected return. This is summarized in the update rule

$$Q_{\text{tar}}^{\pi}(s_t, a_t) = r_t + \gamma Q^{\pi}(s_{t+1}, a_{t+1}). \quad (8)$$

Note that the Q -value estimate depends on the following action, as we base our estimates solely on state-action estimates. Over the course of many examples, the proportion of selected actions given a state will approximate the probability distribution over all actions.

In practice, we employ a neural network for the approximation of the Q function, the Q network, which returns the Q -value estimates of the selected state-action pairs. As a consequence, each update of the Q value is not complete, as neural networks learn gradually using gradient descent, moving partially towards the target value. With all that, the SARSA algorithm is summarized in Algorithm 1.

The workflow of the SARSA algorithm is similar to a supervised learning workflow, in which each estimate has a target value to reach and, with it, we can evaluate how well our neural network is performing and thus reduce the discrepancies between the values. In this sense, we use an iterative approach to improve the Q value, as we can explicitly see in line 3. Note from lines 5 and 7 that, as the SARSA algorithm is based on TD learning, only information from the next step is required to form the target of the current Q value, allowing us to update the Q function in a batched manner. Regarding sample efficiency, we can see from line 4 that the next action a_{t+1} is obtained with the same policy used to gather the previous action a_t , that is, the ε -greedy policy over the only Q network of the algorithm. This specific selection of the next action makes the SARSA algorithm an *on-policy* algorithm, that is, an algorithm in which information for improving the

ALGORITHM 1. SARSA pseudocode [33,34].

Input: a differentiable action-value function parametrization $\hat{q}: S \times A \times \mathbb{R}^d \rightarrow \mathbb{R}$

Input: algorithm parameters. Initialize learning rate $\alpha > 0$, epsilon $1 \geq \varepsilon > 0$, discount rate $1 \geq \gamma \geq 0$.

Output: optimized \hat{q}

- 1 Initialize arbitrarily \hat{q} weights, $\vec{w} \in \mathbb{R}^d$
- 2 Initialize $s_0 \neq \text{terminal}$
- 3 **for** training step **do**
- 4 Generate a new batch of episodic experiences $s_t, a_t, r_t, s_{t+1}, a_{t+1}$ with ε -greedy policy wrt. $\hat{q}(s_t, \cdot)$
- 5 **for** experience in batch **do**
- 6 Store estimation(\vec{w}): $\hat{q}(s_t, a_t)$
- 7 Store target: $r_t + \gamma \hat{q}(s_{t+1}, a_{t+1})$
- 8 $\vec{w} \leftarrow \vec{w} + \alpha \nabla_{\vec{w}} \text{loss}(\text{estimation}(\vec{w}), \text{target})$
- 9 Decay ε
- 10 Decay learning rate

policy (a_{t+1}) depends on the policy to gather data (a_t). Because of this on-policy behavior, each training iteration can only use experiences obtained following the current policy, so each time the Q -network parameters are updated, all experiences must be discarded and new experiences have to be collected, as reflected again with line 4 and its position within the training loop. Finally, something that arises from the use of currently collected experiences for the estimation of the target Q value is the high correlation between experiences, as the data used to update the network are often from a single episode, which can lead to high variance in different parameter updates.

In Fig. 3 we present a summary of the results obtained with the SARSA algorithm for our hyperbolic multilayer system with 16 layers, which includes 40 independent runs, represented with their means and standard deviations. These results were obtained with the hyperparameters specified in Table I in Appendix A. Figure 3(a) shows the largest HTC obtained as a function of the number of found states. To gauge the quality of our method, we compare SARSA results in this panel with an algorithm in which different states are randomly selected and the maximum HTC is recorded as the algorithm progresses. This random algorithm is particularly efficient for relatively small state spaces, as it is forced to always find new states. Therefore, its results constitute a good test for the different RL algorithms. To ensure reliable statistics, its mean value and deviation were obtained with 1000 independent runs in all cases.

In Fig. 3(b) we present the evolution of the loss function of the Q network as the training proceeds. Finally, Fig. 3(c) displays the return in a greedy simulation of an episode with the Q function obtained each training step. Recalling the basis of RL (see Sec. II), the usual goal of RL is to maximize the expected return $G(\tau)$. In this sense, we can see in Fig. 3(c) that the return increases along the exploratory

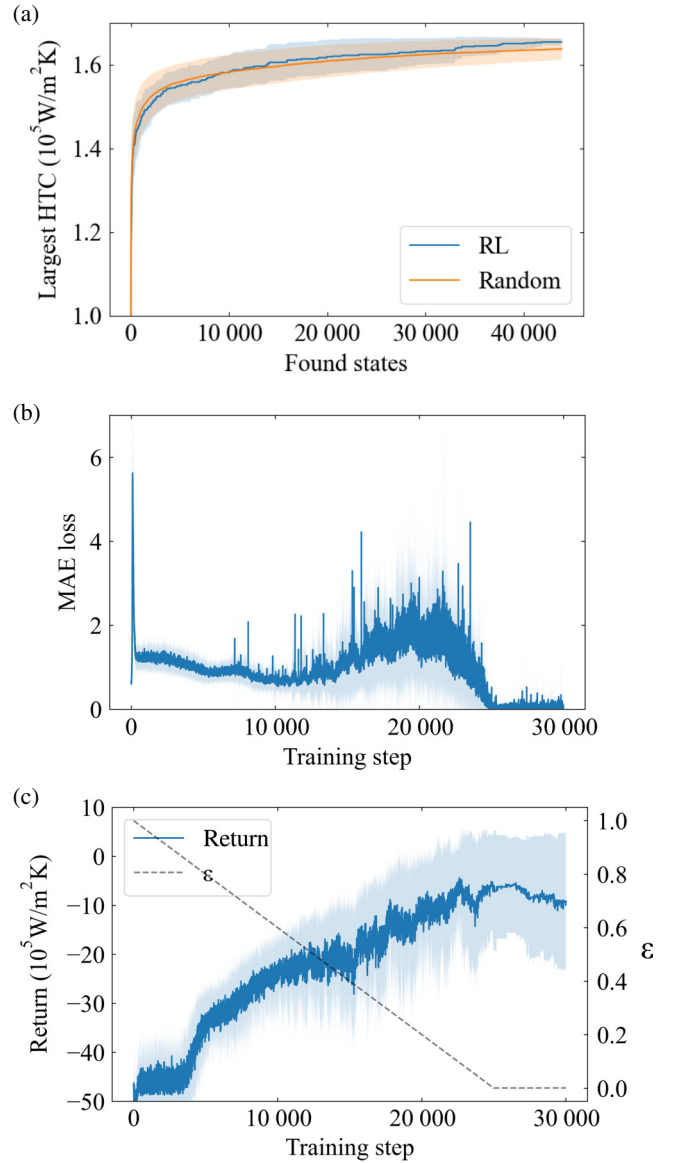


FIG. 3. Training of the SARSA algorithm for our physical problem of interest. (a) Largest HTC discovered as a function of the number of found states in the problem with 16 layers obtained with the SARSA algorithm. We also present the results obtained with the random algorithm. (b) The evolution of the corresponding loss curve of the SARSA algorithm. (c) Return obtained in a simulation of an episode with the Q network of the SARSA algorithm at each training step. The dashed line corresponds to the value of ε (right scale). In all panels the solid lines correspond to mean values and the shaded areas to standard deviations, as obtained in 40 independent runs of the SARSA algorithm.

phase, so the SARSA algorithm is achieving better policies as the training proceeds, as expected. Regarding the loss function, Fig. 3(b) shows that its general tendency is indeed to decrease, indicating that the Q network tends to converge, as desired. However, some noise during the

training can be noticed. When ε is close to 0.5, the average loss starts increasing, and when ε reaches its minimum value, the loss rapidly adopts its lowest value. Apart from the variance arising just from the SARSA formulation due to the high correlation between the experiences of each batch, the noise seems to arise from the greedy behavior of the method in intersection with the shape of the state-action space. In our case, the noise could indicate that the Q value of the state actions visited through the greedy policy is not that close to the value of the previously seen state-action pairs. With that, an overfitting of the Q network is thought to be made to the state actions of the greedy behavior, which increases the mean loss when those states are not that frequently visited and gradually decreases the loss as the greedy behavior is more prominent.

On the other hand, Fig. 3(a) shows that the SARSA algorithm is capable of finding the optimal configuration for our problem, but it performs similarly to the random algorithm with the exception of the end of the curve, where SARSA runs exhibit a smaller variance. Thus, the main conclusion from this analysis is that although the SARSA algorithm can learn an improved policy, it is not sample efficient enough to achieve our objective for the selected environment. It is important to emphasize that, as in any machine learning problem, we cannot rule out that, with a better selection of hyperparameters, the SARSA algorithm could clearly beat the random algorithm, especially in problems of higher dimensionality. In any case, as we will see in the next subsection, sample efficiency can be notably improved using other types of value-based algorithms, so we will not dwell too much here with the SARSA algorithm.

2. Q learning

As in the SARSA algorithm, Q learning is based on TD learning in order to obtain the target value. In this case, the update rule for the Q values reads

$$Q_{\text{tar}}^{\pi}(s, a) = r_t + \gamma \max_{a_{t+1}} Q^{\pi}(s_{t+1}, a_{t+1}). \quad (9)$$

Note that the selection of the following action a_{t+1} is made with a max operator, which indicates that we are taking the action that maximizes the Q value of the next state. This might be seen as a small change with respect to the SARSA algorithm, but it has important consequences, overcoming some of the SARSA algorithm's limitations. With it, in Eq. (9) we are learning the optimal Q function instead of the Q function of the current policy as in the SARSA algorithm, and so improving the stability and speed of learning. In addition, this makes Q learning an *off-policy* algorithm, as the information used to learn the Q value is independent of the policy used for gathering data. Therefore, off-policy behavior allows us to learn from experiences gathered by any policy. It allows us to reuse

ALGORITHM 2. Q -learning pseudocode [33,34].

Input: a differentiable action-value function parametrization $\hat{q}: S \times A \times \mathbb{R}^d \rightarrow \mathbb{R}$
Input: algorithm parameters. Initialize learning rate $\alpha > 0$, epsilon $1 \geq \varepsilon > 0$, discount rate $1 \geq \gamma \geq 0$.
Output: optimized \hat{q}

- 1 Initialize arbitrarily \hat{q} weights, $\vec{w} \in \mathbb{R}^d$
- 2 Initialize $s_0 \neq \text{terminal}$
- 3 **for** *training step* **do**
- 4 Generate a batch of episodic experiences s_t, a_t, r_t, s_{t+1} with ε -greedy behavior policy wrt. $\hat{q}(s_t, \cdot)$
- 5 **for** *experience in batch* **do**
- 6 Store *estimation*(\vec{w}): $\hat{q}(s_t, a_t)$
- 7 Choose a_{t+1} with *greedy* target policy wrt. $\hat{q}(s_{t+1}, \cdot)$
- 8 Store *target*: $r_t + \gamma \hat{q}(s_{t+1}, a_{t+1})$
- 9 $\vec{w} \leftarrow \vec{w} + \alpha \nabla_{\vec{w}} \text{loss}(\text{estimation}(\vec{w}), \text{target})$
- 10 Decay ε
- 11 Decay *learning rate*

and decorrelate experiences, reducing the variance of each update and improving the sample efficiency with respect to the SARSA algorithm.

It is interesting to underline that, for our application, the reuse of experiences is of great importance as our goal is to obtain the maximum HTC with as few explored states as possible. In addition, the usage of neural networks makes this aspect even more relevant as they rely on gradient descent, for which each parameter update must be small because the gradient only returns meaningful information near employed parameters. This makes the possibility of reusing experiences important as the network's parameters may need to be updated multiple times.

We summarize the Q -learning algorithm in Algorithm 2. Comparing lines 4 and 7, one sees that Q learning gathers experiences selecting action a_t with ε -greedy behavior, while it estimates the target Q value with a greedy selection of a_{t+1} , as previously seen in the update rule of Eq. (9). Although those lines are the only difference with the SARSA algorithm, they enable us to further implement all previously mentioned advantages like experience reuse with the extensions described below. In this work we present only the Q -learning results obtained with the most refined version of the Q -learning algorithm. Therefore, we postpone discussion of the Q -learning results and introduce in what follows several sophistications for the naive version of this algorithm.

3. Q -learning extensions

As mentioned above, Q learning is potentially better than the SARSA algorithm to achieve our goal because of its off-policy nature [33]. Building upon this nature, here we make use of several modifications of Q learning that

have been proposed to enhance its sample efficiency and stability.

Experience replay. Introduced by Lin [51], this idea consists of storing experiences in a memory in order to reuse them even if they were taken with old policies, allowing for more efficient learning from a reduced number of gathered experiences. In practice, an *experience replay memory* stores the agent's most recent experiences up to a given memory size, large enough to contain many episodes, replacing the oldest experiences by the newest ones once this size is reached. With that, every time an agent needs batches to be trained, it retrieves them from the replay memory in a random-uniform manner. Then, each one of the batches is used to update the training network. In this way, in addition to introducing a higher sample efficiency, we ensure that we have decorrelated experiences for training as they are likely to be from different policies and episodes, contrary to the SARSA algorithm, stabilizing the training as we reduce the variance of parameter updates. Finally, to set a widely used notation, the combination of Q learning with the usage of Q networks and the presented memory replay technique is referred to as the *deep Q networks* (DQN), set by Mnih *et al.* [52]. We use this notation throughout the rest of the paper.

Target network. Introduced by Mnih *et al.* [53], it focuses on reducing the changes in the target value by means of a second network, called the *target network*. The idea is to use a lagged copy of the training network whose update frequency is less than that of the training network. Then, this secondary network is used to generate the state-action estimate for the target value, $\max_{a_{t+1}} Q^\pi(s_{t+1}, a_{t+1})$, stopping its value from moving. This idea addresses the issue of the target constantly changing because of network updates, stabilizing the training and making divergences less likely as it avoids ambiguity regarding the values the network must approach.

Double Q learning. Introduced by van Hasselt *et al.* [54,55], the basic idea is to use two different networks trained with different experiences for the estimation of the next Q value used for obtaining the target value. This double estimation is computed using a network for retrieving the maximizing action, and the remaining network for producing the Q value with the selected action:

$$Q_{\text{tar}}^\pi(s, a) = r_t + \gamma Q_2^\pi\left(s_{t+1}, \underset{a_{t+1}}{\operatorname{argmax}} Q_1^\pi(s_{t+1}, a_{t+1})\right). \quad (10)$$

It mitigates the systematic overestimation of the state-action values by the deep Q -learning algorithm. This effect arises from the use of an approximated algorithm. As it does not return a perfect estimation, if Q estimations contain any errors, maximum state actions are likely to be positively biased, resulting in an overestimation of the Q values as Hasselt *et al.* [54] showed in their paper. With

that, if we introduce the usage of a second network trained with different experiences, we can remove the positive bias in the estimation. As with the introduction of the target network, we already have a second network and we want to avoid sampling more experiences; a common practice is the usage of it as the secondary network. Although it is just a lagged copy of the training network, if the update frequency of the target network is sufficiently low, it is considered to be different enough from the training network to function as a different one.

Taking into consideration all these modifications, the final double DQN algorithm we use for our application is summarized in Algorithm 3. Let us emphasize the main differences with respect to the Q -learning algorithm in Algorithm 2. Lines 4, 6, and 8 describe the usage of memory replay, first initializing it to gradually add more experiences at each training step. For training, B batches of experiences are sampled, further leveraging the reuse of experiences compared to not using several batches. Next, lines 2, 12, and 15 refer to the usage of a target network, which is updated to the weights of the training network with a frequency F . Finally, double estimation is reflected in lines 11 and 12, where a_{t+1} is taken with the training network \hat{q} , while the next state-action value is obtained through the target network \tilde{q} .

Making use of this final *double DQN algorithm* and the hyperparameters listed in Table II in Appendix A, we obtained the results summarized in Fig. 4 for our multilayer system, again with 40 independent runs of the double DQN algorithm. In Fig. 4(a) we display the maximum HTC for the double DQN and random algorithms as a function of the number of found states. The double DQN algorithm surpasses the random algorithm for a relatively small amount of explored states, about 7000 versus about a possible 65 000 states of the system. To emphasize the quality of these results, let us say that, for the 40 runs of the double DQN algorithm, the top five best HTC values for this system are found in 77.5% of the runs and the best possible state in 35.0% of the runs, while for the 1000 runs of the random algorithm, these values are found in 44.3% and 12.9% of the runs, respectively.

To gain some insight into the training process of our algorithm, we display in Fig. 4(b) the evolution of the loss of the Q network. This loss decreases almost monotonically, which indicates that our algorithm is training. However, some irregularities can be appreciated again. First, we can see some regular peaks. These just correspond to the update of the target network, which produces the sudden change of the target. In addition, some noise appears at the end of the curve, which we believe could arise from two effects. First, from the same effect as in the SARSA's loss, i.e., the overfitting of the Q network. Second, it could be due to the target network not being that frequently updated, leading to a target with less information. Therefore, we can

ALGORITHM 3. Double Q learning with memory replay and target network pseudocode [33,34].

Input: differentiable action-value function parametrizations $\hat{q}, \tilde{q} : S \times A \times \mathbb{R}^d \rightarrow \mathbb{R}$

Input: algorithm parameters. Initialize learning rate $\alpha > 0$, epsilon $1 \geq \varepsilon > 0$, discount rate $1 \geq \gamma \geq 0$, new experiences per episode $h > 0$, batches per training step $B > 0$, target network update frequency $F > 0$.

Output: optimized \hat{q}

- 1 Initialize arbitrarily \hat{q} weights, $\vec{w} \in \mathbb{R}^d$
- 2 Equal target network \tilde{q} weights to \hat{q} weights, $\vec{\varphi} \in \mathbb{R}^d = \vec{w} \in \mathbb{R}^d$
- 3 Initialize $s_0 \neq \text{terminal}$
- 4 Initialize memory replay with a batch of episodic experiences s_t, a_t, r_t, s_{t+1} with ε -greedy behavior policy wrt. $\hat{q}(s_t, \cdot)$
- 5 **for** training step **do**
- 6 Store in memory h episodic experiences s_t, a_t, r_t, s_{t+1} with ε -greedy behavior policy wrt. $\hat{q}(s_t, \cdot)$
- 7 **for** batch B **do**
- 8 Sample a batch of experiences from memory
- 9 **for** experience in batch **do**
- 10 Store estimation(\vec{w}): $\hat{q}(s_t, a_t)$
- 11 Choose a_{t+1} with greedy target policy wrt. $\hat{q}(s_{t+1}, \cdot)$
- 12 Store target: $r_t + \gamma \hat{q}(s_{t+1}, a_{t+1})$
- 13 $\vec{w} \leftarrow \vec{w} + \alpha \nabla_{\vec{w}} \text{loss}(\text{estimation}(\vec{w}), \text{target})$
- 14 **if** training step $\propto F$ frequency **then**
- 15 Update target network \tilde{q} weights, $\vec{\varphi} = \vec{w}$
- 16 Decay ε
- 17 Decay learning rate

again end with a less generic Q network, which can lead to higher values of loss.

Finally, as previously discussed, an important metric for the performance of the algorithm is the return in a greedy simulation with the Q function obtained each training step; see Fig. 4(c). Note that the return first increases, but ends up decreasing, which could appear to be upsetting. Let us recall that the return is the main result for RL applications with the standard algorithm's objective: to obtain a policy that maximizes the return of the system. However, our final objective here is slightly different: to explore the optimal state with as few explored states as possible. This is why the observed decay in the return is not worrisome in our case, although it means that we end up having a nonoptimal policy.

In Fig. 4(c) we also show as a vertical green line with a shaded area both the mean and standard deviation of the training step at which the highest HTC is discovered in the different runs. Two things are worth noting: in this region ε still has a sizeable value, so the algorithm still has chances of discovering better states, and the return is still growing and is higher than the return of a completely

exploratory policy, so we have a policy with some learning. This supports the fact that the decrease of the return is not something to worry about: with selected hyperparameters, our algorithm uses the policy learned at early stages to explore better states.

Although the decay of the return is not something critical in our case, it is important to understand why it occurs. A possible explanation is the lack of convergence to the optimal Q function. As mentioned in Sec. III B, our reward is the difference of the HTC of the periodic multilayer system and that of the next state. In Fig. 4(c), we can see that the return does not surpass the zero value, so we are not close to the optimal policy and, therefore, to the optimal Q function. The reason why a below-zero value is not close to a simulation of an episode of the optimal policy is the following. Knowing the optimal state (see Sec. V below) and given a large episode length of 64 steps, an intuitively good policy would imply transitioning with as few steps as possible from the start state to the optimal one. It would result in transitioning to the optimal state in just seven steps. Therefore, this policy would output seven unknown rewards and a reward of 0.28×10^5 W/m²K during 57 steps. In addition to the best reward, we know the worst possible one, which has a value of -1.37×10^5 W/m²K. With it, the seven unknown rewards must be equal or higher than the worst possible reward. Therefore, the good policy we have imagined has a return $G(\tau) \geq 6.37 \times 10^5$ W/m²K, which has a positive value. By definition, the optimal policy is such that its expected return is greater than or equal to any of the remaining existent policies for all states [33]. With this example, we have found a policy whose expected return from the initial state is higher than zero, which is over the return displayed at Fig. 4(c). Thus, this demonstrates that we have not reached the optimal policy in that figure.

As we have not reached the optimal policy and therefore the true value of the Q function, we only have an imperfect estimation of it. This estimation can help us reach the optimal policy during training. However, if we exploit it instead of using it to continue looking for the optimal policy, we can end up overtraining our network with experiences that a good policy is not likely to visit. This puts our policy farther from the optimal Q function, losing the part of the estimation that was towards the good policy and, finally, turning the policy into a worst one because we are following nonoptimal state actions.

An interesting issue at this stage concerns the impact of having a better estimation of the Q values with the usage of more experiences and exploiting this knowledge to obtain the optimal state. To elucidate this issue, in Fig. 5 we present the results obtained for the double DQN algorithm using now 25 experiences stored in the memory replay per training step, rather than four, as presented in Fig. 4. In Fig. 5(b), the loss function of the neural network still decays, as expected with already known irregularities,

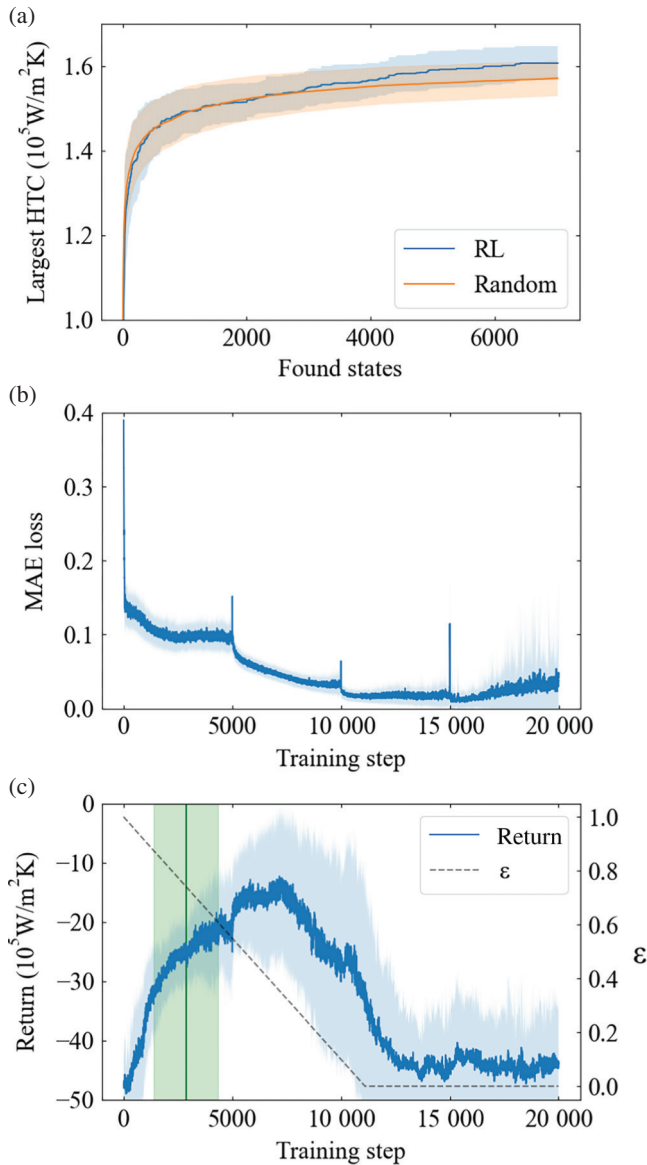


FIG. 4. Training of the double DQN algorithm. (a) Largest HTC discovered as a function of the number of found states in the problem with 16 layers obtained with the double DQN algorithm. We also present the results obtained with the random algorithm. (b) Evolution of the corresponding loss curve of the double DQN algorithm. (c) Return obtained in a simulation of an episode with the Q network of the double DQN algorithm at each training step. The dashed line corresponds to the value of ϵ (right scale). The vertical line and the green shaded area correspond to the training steps at which the highest HTC of the runs is found. In all panels the solid lines correspond to mean values and the shaded areas to standard deviations, as obtained in 40 independent runs. In all cases, four experiences were stored per training step.

while in Fig. 5(c), the return increases up to a maximum value and stays there, as desired in regular applications of RL. This can mean that a sufficiently good policy is reached, so following it does not put us further from the state actions an optimal or suboptimal policy would

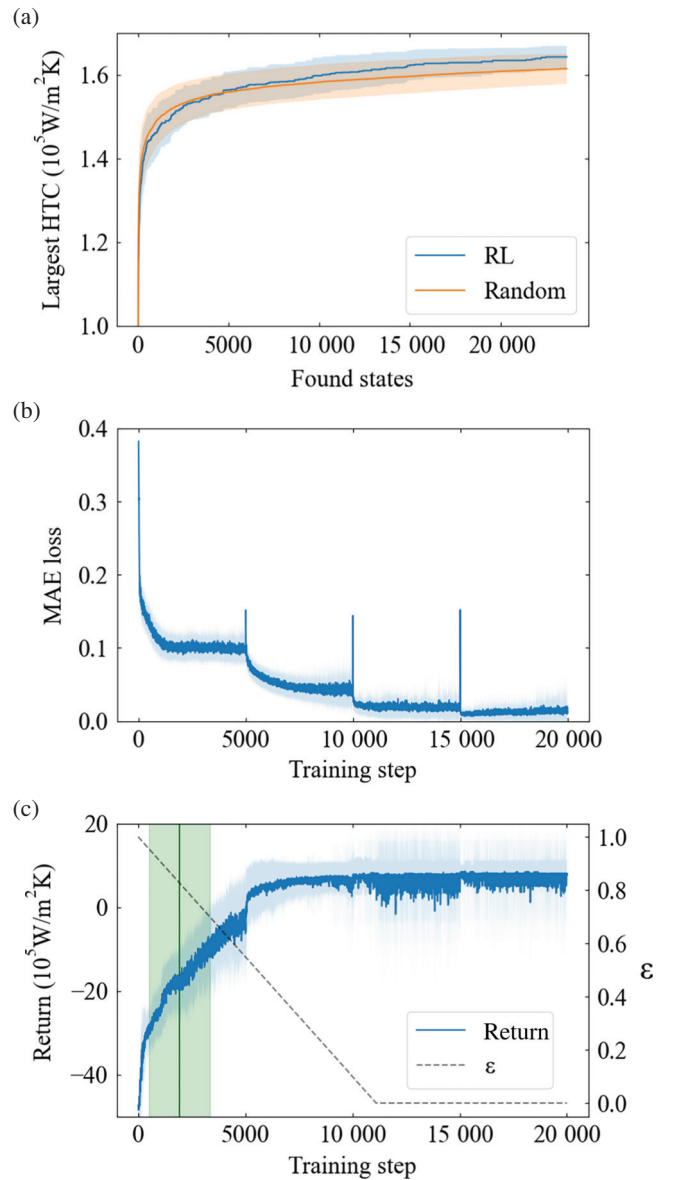


FIG. 5. Same as in Fig. 4, but with 25 experiences stored per training step.

follow. However, something must be noted from Fig. 5(a), namely, higher HTC values with respect to the random algorithm are now discovered when more states have been explored, distancing us from our true objective: obtaining the state that gives us the maximum HTC with as few explored states as possible. This leads to the conclusion that, although we can miss the opportunity of learning a decent estimation of relevant state-action values, it is worth gathering states more slowly during the training as, just employing them, we can discover the optimal states too. A similar behavior will also be seen with policy-based algorithms.

Finally, Fig. 5(c) shows that the highest HTC values are discovered early during the training. Again, this means that

there is no need to reach a good policy in order to find the good states of our application, which suggests training the Q network with fewer states during more training time.

B. Policy-based algorithms: REINFORCE

Now we focus on the analysis of the results obtained with REINFORCE [56], which is the most widely used policy-based RL algorithm. In this type of algorithm, the agent learns a policy function π , which in turn is used to produce actions and generate trajectories τ that maximize the objective $J(\tau)$. REINFORCE needs three components: (i) a parametrized policy, (ii) an objective to be maximized, like any other RL algorithm, and (iii) a method for updating the policy parameters. Concerning the parametrized policy, this is obtained with the help of deep neural networks that learn a set of parameters θ . We denote the policy network as π_θ to emphasize that is parametrized by θ . Framed in this way, the process of learning a good policy is equivalent to searching for a good set of values for θ .

The objective that is maximized by an agent in REINFORCE is the expected return over all complete trajectories generated by an agent:

$$J(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta}[G(\tau)] = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^T \gamma^t r_t \right]. \quad (11)$$

Note that the expectation is calculated over many trajectories sampled from a policy, that is, $\tau \sim \pi_\theta$. This expectation approaches the true value as more samples are collected, and it is specific to policy π_θ used.

The final component of the algorithm is the policy gradient, which formally solves the problem

$$\max_{\theta} J(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta}[G(\tau)]. \quad (12)$$

To maximize the objective, we perform gradient ascent on the policy parameters θ :

$$\theta \leftarrow \theta + \alpha \nabla_{\theta} J(\pi_\theta). \quad (13)$$

Here α is the *learning rate*, which controls the size of the parameter update. The term $\nabla_{\theta} J(\pi_\theta)$ is known as the *policy gradient* and, thanks to the policy gradient theorem, [33,34] can be expressed as

$$\nabla_{\theta} J(\pi_\theta) = \mathbb{E}_{\tau \sim \pi_\theta} \left[\sum_{t=0}^T G_t(\tau) \nabla_{\theta} \log \pi_{\theta}(a_t | s_t) \right]. \quad (14)$$

Here, the term $\pi_{\theta}(a_t | s_t)$ is the probability of the action taken by the agent at time step t . The action is sampled from the policy, $a_t \sim \pi_{\theta}(s_t)$.

ALGORITHM 4. REINFORCE pseudocode [34,56].

Input: a differentiable policy parametrization
 $\pi_{\theta} : S \times \mathbb{R}^d \rightarrow A$

Input: algorithm parameters. Initialize learning rate
 $\alpha > 0$ and discount rate $1 \geq \gamma \geq 0$

Output: optimized policy π_{θ}

- 1 Initialize arbitrarily π_{θ} weights, $\vec{\theta} \in \mathbb{R}^d$
- 2 **for** *episode* **do**
- 3 Initialize $s_0 \neq \text{terminal}$
- 4 Initialize $\nabla_{\theta} J(\pi_{\theta}) \leftarrow 0$
- 5 **for** *episode step* **do**
- 6 Take a_t following $\pi_{\theta}(a_t | s_t)$, observe r_t, s_{t+1}
- 7 $s_t \leftarrow s_{t+1}$
- 8 **for** $t = 0 \dots T$ **do**
- 9 $G_t(\tau) \leftarrow \sum_{t'=t}^T \gamma^{t'-t} r_{t'}$
- 10 $\nabla_{\theta} J(\pi_{\theta}) \leftarrow \nabla_{\theta} J(\pi_{\theta}) + G_t(\tau) \nabla_{\theta} \log \pi_{\theta}(a_t | s_t)$
- 11 $\vec{\theta} \leftarrow \vec{\theta} + \alpha \nabla_{\theta} J(\pi_{\theta})$

In practice, the REINFORCE algorithm numerically estimates the policy gradient using Monte Carlo sampling. Instead of sampling many trajectories per policy, one samples just one:

$$\nabla_{\theta} J(\pi_{\theta}) \approx \sum_{t=0}^T G_t(\tau) \nabla_{\theta} \log \pi_{\theta}(a_t | s_t). \quad (15)$$

The REINFORCE algorithm is summarized in Algorithm 4. Let us emphasize how this is an episodic, *on-policy* algorithm: every episode we start from the starting state s_0 in line 3 and collect a full trajectory $\tau = (s_0, a_0, r_0), \dots, (s_T, a_T, r_T)$ for an episode in lines 6 and 7. Then, in line 9 the return $G_t(\tau)$ is computed for each time step t in the current trajectory, which is later used in line 10 to estimate the policy gradient along with this policy's action probabilities $\pi_{\theta}(a_t | s_t)$. In line 11, the sum of the policy gradients for all time steps calculated in line 10 is used to update the policy network parameters θ , and the experiences are discarded.

It is known that the fact that the policy gradient is estimated, sampling with a single trajectory typically leads to a high variance. One way to reduce this variance is to modify the returns by subtracting a suitable action-independent baseline as

$$\nabla_{\theta} J(\pi_{\theta}) \approx \sum_{t=0}^T [G_t(\tau) - b(s_t)] \nabla_{\theta} \log \pi_{\theta}(a_t | s_t). \quad (16)$$

In our case we chose as a baseline the mean return over the trajectory, namely, $b = (1/T) \sum_{t=0}^T G_t(\tau)$. Our baseline is state independent, being a constant value for each trajectory, and centers the returns around 0. This enables faster learning, correcting any imbalance between positive and negative returns. Note that this REINFORCE-specific

baseline is independent of the baseline considered for the HTC values in the rewards.

Let us now describe our application of REINFORCE to our optimization problem. We first consider hyperparameters given in Table III in Appendix A. The corresponding evolution of the training for 10^5 full episodes is shown in Fig. 6(a), where we present the return of each episode as training progresses. These results correspond to 40 independent runs. Note that the return increases until it reaches a plateau, when the policy algorithm converges. The vertical line and the green shaded area correspond to the mean value and standard deviation of the training step at which the highest HTC of the run is found. We can see that it was found in a much smaller number of steps compared to the total number of training steps needed for the convergence of the policy. This is a very similar result to what we found using the DQN algorithm. In Fig. 6(b) we show the comparison between the best HTC found by REINFORCE and by a random search, also for 40 different runs of the REINFORCE algorithm and 1000 runs of the random algorithm. The random search algorithm is the same as in previous algorithms, so as to ensure a fair comparison. We can see how REINFORCE seems to behave in a very similar fashion to random search, being unable to significantly outperform it.

To improve these results, we explored the same idea as in the double DQN algorithm, namely, to increase the learning rate so it can learn with a reduced number of explored states. In Fig. 6(c) we show the same as in Fig. 6(a), but with a value of $\alpha = 3 \times 10^{-4}$, 10 times larger than before. As one can see, the training of the policy fails, with the return initially increasing, but eventually decreasing abruptly, with a much larger standard deviation than before. In a regular RL problem, we would say that training has failed, and the learning rate is too high. However, if we look at Fig. 6(d), the same as in Fig. 6(b) but with the new value for the learning rate, REINFORCE shows the same effect as in the double DQN algorithm: it is able to slightly outperform the random search algorithm in both the mean and the area covered by the standard deviation. Moreover, the green shaded area in Fig. 6(c) clearly demonstrates how the best examples are found with a really small number of training steps, even lower than in Fig. 6(a), and the decay of the return is inconsequential to our goal. This is also analogous to our results for the double DQN algorithm.

We can further identify this improvement through the statistics of the runs, both from REINFORCE and the random search algorithm. Considering the 10^4 new examples found in Fig. 6(b), the random search algorithm is able to find one of the five configurations with the highest HTC 57.0% of the runs, with 17.3% finding the actual maximum. On the other hand, REINFORCE is able to find one of the five highest HTC configurations 60% of the runs, which is only slightly better than the random search, with 12.5% finding the actual maximum, less than the random

search. With this, we would conclude that the algorithm is worse than random search. But using the improved version in Fig. 6(d), with a new restriction of 5000 new examples, which is half the amount of new states compared to Fig. 6(b), the random search can only find one of the five configurations with the highest HTC 34.9% of the runs, with 9.3% finding the actual maximum, and REINFORCE can find one of the five configurations with the highest HTC 42.5% of the runs, with 15% finding the actual maximum, being able to surpass the random search results. However, the difference is not significant enough, and as such, REINFORCE is not the best option for solving our RL problem. The Monte Carlo sampling and the sample inefficiency of the algorithm prevent it from obtaining comparable or even better results than other algorithms. Still, its convergence is more easily guaranteed, albeit noisy, and its implementation is more straightforward than value-based algorithms, so it can still be a candidate for other possible applications.

C. Combined methods: actor-critic and proximal policy optimization algorithms

So far, we have discussed both policy-based and value-based algorithms. Now, we consider *combined methods* that learn two or more of the primary RL functions. To be precise, we discuss both *advantage actor-critic* (A2C) [57] and *proximal policy optimization* (PPO) [58] algorithms, which are among the most widely used RL algorithms for a wide variety of applications. Actor-critic algorithms receive their name from the two elements that compose them: an actor, which learns a parameterized policy like in REINFORCE; and a critic, which learns a value function to evaluate state-action pairs, becoming a learned reinforcement signal. In simple words, the foundations of actor-critic algorithms involve trying to learn a value function to give the policy a more informative metric than just the rewards. When the learned reinforcement signal is the advantage function, the algorithm is called the advantage actor-critic algorithm. The advantage function is defined as

$$A^\pi(s_t, a_t) = Q^\pi(s_t, a_t) - V^\pi(s_t), \quad (17)$$

thus describing how preferable an action would be compared to the average weighted by the policy in a particular state. This allows us to rescale their values for all states and actions, similarly to a state-dependent baseline, and presents other useful properties such as $\mathbb{E}_{a \in A}[A^\pi(s_t, a)] = 0$. The actor uses this signal in place of the estimate of the return from the REINFORCE algorithm, performing the same gradient ascent technique. That is, the actor performs the following policy optimization (setting $\pi_\theta \rightarrow \pi$):

$$\nabla_\theta J(\pi) = \mathbb{E}_t[A_t^\pi \nabla_\theta \log \pi(a_t|s_t)]. \quad (18)$$

The critic is tasked with estimating this advantage function for all states and actions. In principle, this would imply

being able to estimate both $Q^\pi(s, a)$ and $V^\pi(s)$, but there are methods that allow the estimation of $Q^\pi(s, a)$ through $V^\pi(s)$ over the trajectory, allowing us to only need to learn the latter. There are various ways to estimate the advantage function using these value functions. The first one is called n -step returns, in which we expand the definition of $Q^\pi(s, a)$ using the rewards obtained in the current trajectory for n steps, and then take the V value of the following one. That is, we expand

$$\begin{aligned} Q^\pi(s_t, a_t) &= \mathbb{E}_{\tau \sim \pi} [r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots + \gamma^n r_{t+n}] \\ &\quad + \gamma^{n+1} V^\pi(s_{t+n+1}) \\ &\approx r_t + \gamma r_{t+1} + \gamma^2 r_{t+2} + \dots + \gamma^n r_{t+n} \\ &\quad + \gamma^{n+1} V^\pi(s_{t+n+1}), \end{aligned} \quad (19)$$

which assumes accurate $V^\pi(s)$ estimations. This leaves our bias-variance trade-off explicit: the rewards from the

trajectory have high variance, while $V^\pi(s)$ is a biased estimation. Higher values of n present higher variance, so n should be chosen to balance these two effects. Another way to estimate the advantage is called *generalized advantage estimation* (GAE) [59]. GAE calculates an exponentially weighted average of all n -step advantages, intending to reduce the variance of the estimation while keeping the bias low. The expression is

$$A_{\text{GAE}}^\pi(s_t, a_t) = \sum_{k=0}^{\infty} (\gamma \lambda)^k \delta_{t+k}, \quad (20)$$

where $\delta_t = r_t + \gamma V^\pi(s_{t+1}) - V^\pi(s_t)$ with $\lambda \in [0, 1]$ controlling the decay rate. A higher value introduces higher variance, up to a value of 1 that represents the Monte Carlo estimate. A value of 0, on the other hand, computes the TD estimate of the returns. GAE is our estimation of choice for the advantage function, which we use for our physical problem. Finally, we need to obtain a way to estimate

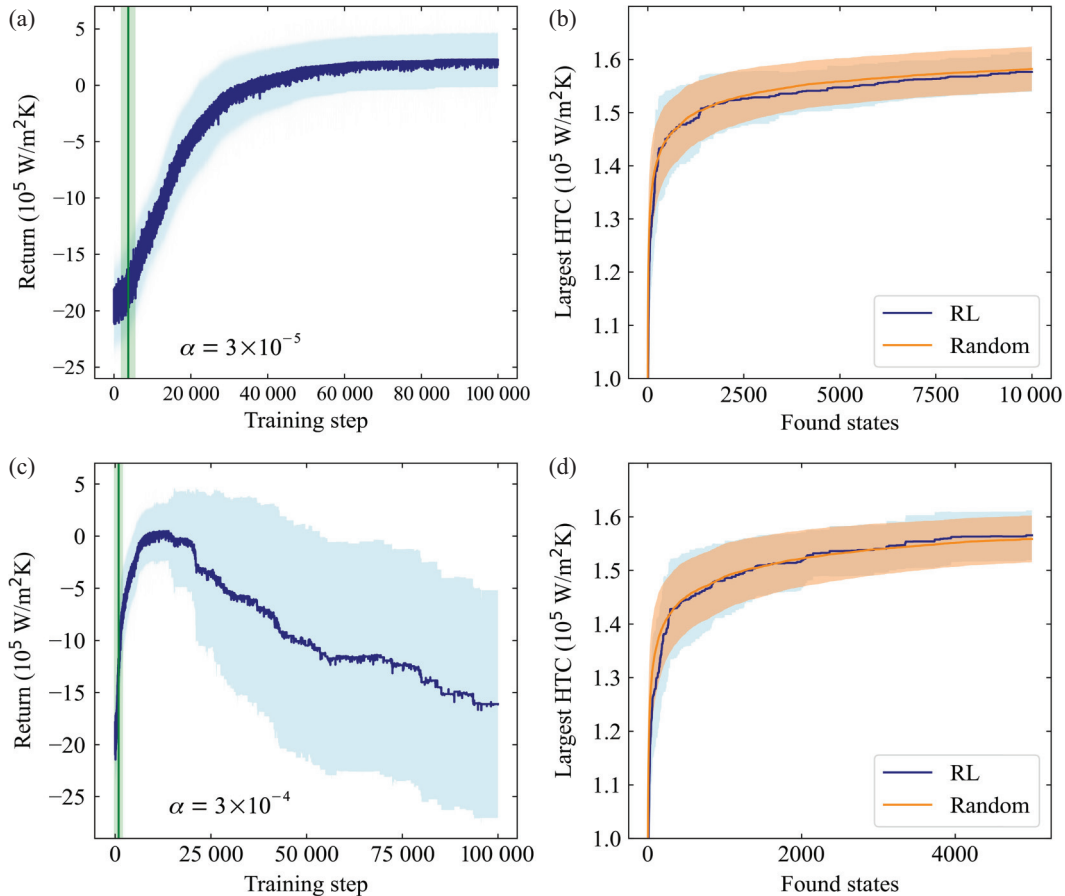


FIG. 6. Training of a REINFORCE algorithm. (a) Evolution of the return for the trajectories as training progresses. The vertical line and the green shaded area correspond to the training steps at which the highest HTC of the runs is found. Learning rate chosen as $\alpha = 3 \times 10^{-5}$. (b) Largest HTC discovered as a function of the number of found states for the evolution in (a). REINFORCE in blue versus a random search in orange. (c),(d) Same as in (a) and (b), but with a learning rate $\alpha = 3 \times 10^{-4}$. In all panels the solid lines correspond to mean values and the shaded areas to standard deviations, as obtained in 40 independent runs of the REINFORCE algorithm.

ALGORITHM 5. A2C pseudocode with GAE and mean square error (MSE) [34].

Input: a differentiable policy parametrization $\pi : S \times \mathbb{R}^{d_1} \rightarrow A$

Input: a differentiable state-value function parametrization $\hat{V} : S \times \mathbb{R}^{d_2} \rightarrow \mathbb{R}$

Input: algorithm parameters. Initialize learning rate $\alpha_1 > 0$ and $\alpha_2 > 0$, GAE exponential weight $1 \geq \lambda \geq 0$ and discount rate $1 \geq \gamma \geq 0$

Output: optimized policy π and improved \hat{V}

- 1 Initialize arbitrarily π weights, $\vec{\theta} \in \mathbb{R}^{d_1}$
- 2 Initialize arbitrarily \hat{V} weights, $\vec{\omega} \in \mathbb{R}^{d_2}$
- 3 **for** *episode* **do**
- 4 Initialize $s_0 \neq \text{terminal}$
- 5 **for** *episode step* **do**
- 6 Take a_t following $\pi(a_t|s_t)$, observe r_t, s_{t+1}
- 7 $s_t \leftarrow s_{t+1}$
- 8 **for** $t = 0 \dots T$ **do**
- 9 Obtain $\hat{V}(s_t)$
- 10 Calculate $\hat{A}_{\text{GAE}}^\pi(s_t, a_t) = \sum_{k=0}^{T-t} (\gamma\lambda)^k \delta_{t+k}$
- 11 Calculate $\hat{V}_{\text{tar}}^\pi(s_t) = \hat{A}_{\text{GAE}}^\pi(s_t, a_t) + \hat{V}^\pi(s_t)$
- 12 Obtain actor loss: $L_{\text{actor}} = -\frac{1}{T} \sum_{t=0}^T \hat{A}_{\text{GAE}}^\pi(s_t, a_t) \log \pi(a_t|s_t)$
- 13 Obtain critic MSE: $L_{\text{critic}} = \frac{1}{T} \sum_{t=0}^T (\hat{V}_{\text{tar}}^\pi(s_t) - \hat{V}^\pi(s_t))^2$
- 14 Update actor parameters: $\vec{\theta} \leftarrow \vec{\theta} + \alpha_1 \nabla_{\vec{\theta}} L_{\text{actor}}$
- 15 Update critic parameters: $\vec{\omega} \leftarrow \vec{\omega} + \alpha_2 \nabla_{\vec{\omega}} L_{\text{critic}}$

the $V^\pi(s)$ values for each state. Following the structure of value-based algorithms and the definition of advantage, we set the target of the critic network as

$$V_{\text{tar}}^\pi(s_t) = A_{\text{GAE}}^\pi(s_t, a_t) + V^\pi(s_t), \quad (21)$$

which we obtain in a similar manner to the SARSA and double DQN algorithms [34]. Thus, the full A2C algorithm can be summarized as in Algorithm 5.

Actor-critic algorithms present a series of issues too, some of the most common including performance collapse and sample inefficiency from being on-policy algorithms. PPO is one of the most popular algorithms designed to solve them, using a surrogate objective that ensures monotonic improvements and allows one to reuse off-policy data samples. This new PPO objective replaces the original A2C objective, and could also be applied to REINFORCE.

To understand this algorithm, we first need to consider that we are performing the search of optimal policies in the parameter space of Θ , while the policies are sampled from the policy space Π . Thus, regular steps in the parameter space do not translate to regular steps in the policy space, where the optimal step size may vary depending on the local geometry, and might result in too big or too small policy steps. This is what eventually causes performance collapse.

To avoid this issue, we consider a constraint to the change in the policy space. We define the distance in the objective between policies as [58]

$$J(\pi') - J(\pi) = \mathbb{E}_{\tau \sim \pi'} \left[\sum_{t=0}^T \gamma^t A^\pi(s_t, a_t) \right], \quad (22)$$

where π is the original policy, which we used to calculate A^π , and π' is the policy that we would obtain after the parameter update. This is a measure of the performance of the new policy, and our goal would be to maximize it. This new maximization problem ensures that there is always a monotonic positive improvement, since the worst possible result would be to let $\pi' = \pi$, without any modification to the policy.

Still, we cannot properly use this function as an objective function, because the expectation is performed by sampling from the new policy, but the new policy is only available after the update that would require said new policy. To solve this issue, we perform the sampling using the old policy, but including *importance sampling* terms, the ratio between new and old policies. Thus, the surrogate objective [renamed $J(\pi')$] becomes an expectation over the current policy π , as desired:

$$J(\pi') = \mathbb{E}_{\tau \sim \pi} \left[\sum_{t=0}^T \gamma^t A^\pi(s_t, a_t) \frac{\pi'(a_t|s_t)}{\pi(a_t|s_t)} \right]. \quad (23)$$

Optimization using this objective is still gradient ascent, so this can become the new objective for the policy gradient function. Lastly, we need to check that the error of the estimation given by the approximation with importance sampling is not big enough to no longer fulfill the condition of always having a positive distance in the objective between policies. We know that, for sufficiently close policies, we can bind their error by their KL divergence [60]. We only need to ensure that the policy improvement is bigger than this limit to accept a change. The application of this constraint can be quite straightforward: we simply need to constrain this KL divergence to be smaller than a given value, δ . Thus, the problem becomes

$$\max_{\theta} \mathbb{E} \left[\frac{\pi'(s_t, a_t)}{\pi(s_t, a_t)} A^\pi(s_t, a_t) \right] \quad (24)$$

ensuring that $\mathbb{E}_t \{KL(\pi'(s_t, a_t) || \pi(s_t, a_t))\} \leq \delta$.

We solve this problem using PPO with a clipped surrogate objective, a much simpler implementation doing away with the need to compute the KL divergence. For this, we define a hyperparameter ϵ to constrain the importance sampling terms. Thus, we constrain the objective between $(1 - \epsilon)A_t^\pi$ and $(1 + \epsilon)A_t^\pi$. The implementation is fairly

simple, only needing to change the objective function from A2C in line 12 of Algorithm 5 by this new function:

$$J^{\text{CLIP}}(\theta) = \mathbb{E}_t \left\{ \min \left[\frac{\pi'(s_t, a_t)}{\pi(s_t, a_t)} A^\pi(s_t, a_t), \text{clip} \left(\frac{\pi'(s_t, a_t)}{\pi(s_t, a_t)}, 1 - \epsilon, 1 + \epsilon \right) A^\pi(s_t, a_t) \right] \right\}. \quad (25)$$

Let us now describe our application of A2C and PPO to our optimization problem. Since these algorithms involve a larger amount of hyperparameters, we made use of external packages to a more optimized implementation. The RL problems themselves were implemented with the help of the library Stable-Baselines3 [61], which features a set of reliable RL algorithms using PyTorch. For the optimization of the hyperparameters, we used the library Optuna [62], an automatic hyperparameter optimization software framework. Still, fundamentally, there is no change in the algorithms themselves or their application.

The hyperparameters used in the implemented codes for A2C and PPO are detailed in Tables IV and V in Appendix A, respectively. Most of them were chosen and kept fixed from the beginning, with only the last four entries of their respective tables being optimized by the Optuna hyperparameter search. The hyperparameters of the said search are available in Table VI in Appendix A. The training results for both A2C and PPO are shown in Fig. 7, where we again compare the best state found by the algorithms and a random search; see Fig. 7(a) for A2C and Fig. 7(b) for PPO. Both algorithms were trained for only 4000 full episodes, since training with 10^5 episodes like in REINFORCE yielded the exact same results. Following the format from the rest of the paper, the algorithm is presented in blue, the average in dark blue, and the standard deviation in light blue, while the random search is presented in orange. The random search algorithm results are always obtained using 1000 runs, while for the RL algorithms, we only considered those that were able to perform a reasonable exploration of new found states, resulting in 11 for A2C and 14 for PPO. Note that this is similar to having considered a “pruning” procedure over exploration, and does not consider how efficient those runs were in finding good states, only new ones.

Figure 7(a) compares the A2C and random search algorithms for a total number of 3500 states found, much smaller than other algorithms. The A2C algorithm is able to outperform the random algorithm slightly, in a similar fashion to REINFORCE, but it is not decidedly better. Figure 7(b) compares the PPO and random search algorithms for a total number of 1500 states found, an even smaller number. The PPO algorithm is able to outperform both the random search and A2C algorithms with much

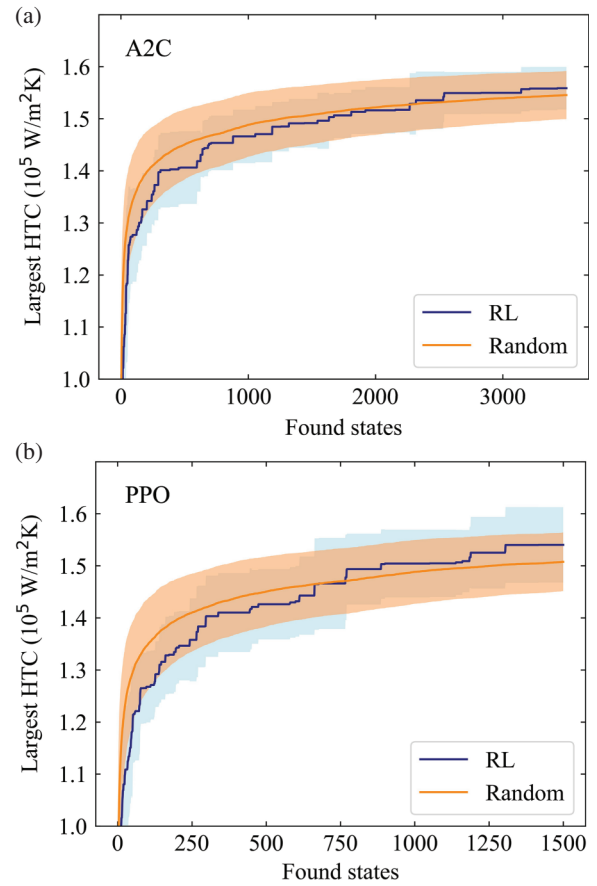


FIG. 7. Training of both the A2C and PPO algorithms for our physical problem of interest. (a) Evolution of the largest HTC found as new states are explored for the A2C algorithm. The A2C algorithm in blue versus a random search in orange. (b) Same as in (a), but with the PPO algorithm instead. In both panels the solid lines correspond to mean values and the shaded areas to standard deviations, as obtained in 11 and 14 independent runs of the A2C and PPO algorithms, respectively.

smaller exploration. However, it also presents a wider standard deviation compared to the A2C algorithm.

We can more clearly describe these results by considering the statistics from the runs presented. For the 3500 found states in the study of A2C, the algorithm is able to find one of the five configurations with highest HTC 36% of the runs, finding the best configuration 9.1% of the runs. In comparison, random search is only able to find one of the best five configurations in 26.2% of the runs, and the best one in 6.8% of them. For PPO, the results are even better than for A2C with a much lower number of states. PPO is able to find one of the best five configurations 29% of the runs, with the best configuration 21% of the runs. Conversely, random search can only find one of the top five configurations 12.3% of the runs, and 3.2% the configuration with highest HTC. This makes PPO the best combined method for our purposes, and a strong candidate for an RL algorithm in these kinds of problem.

V. DISCUSSION AND CONCLUSIONS

All algorithms employed in this work were able to identify, with differing success, the best configuration among the possibilities explored. This optimal 16-layer configuration is shown in Fig. 8(a). Note that it differs from the baseline configuration of Fig. 2(a), which corresponds to the periodic structure that one would propose based on physical intuition. The best state found by our algorithms does not seem to have any particular symmetry, although one can identify a periodic pattern forming a “supercell”; see Fig. 8(a). This pattern is repeated 4

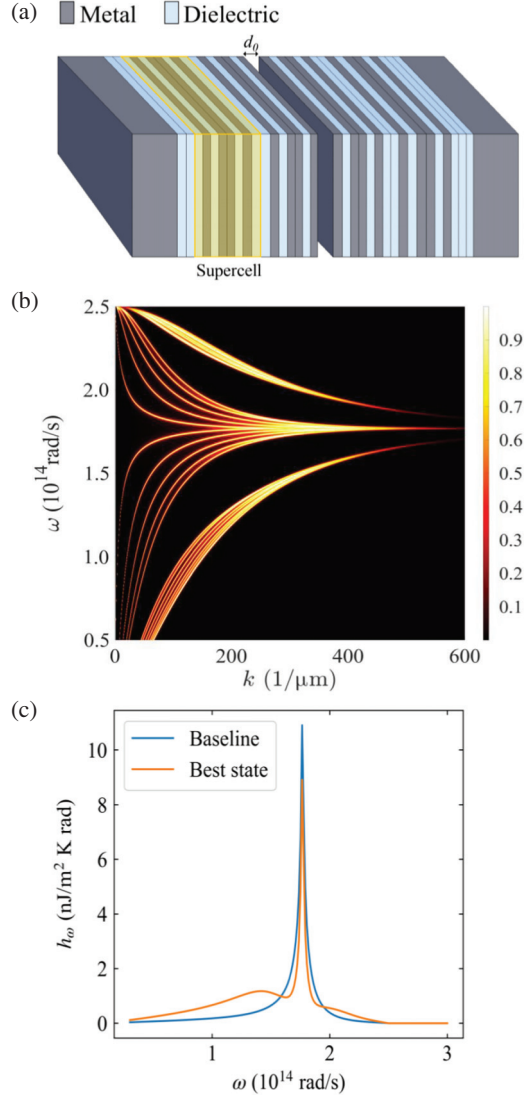


FIG. 8. Description of the best configuration found for the 16-layer problem. (a) Schematic structure of the best state found. A supercell structure is identified as a yellow shaded region on the left subsystem. (b) Transmission of evanescent waves as a function of the frequency (ω) and the parallel wave vector (k) for the optimal structure shown in panel (a). (c) Comparison between the spectral HTC between the best state and the baseline of Fig. 2(b) at room temperature ($T = 300$ K).

times in the whole system when taking into account the gap—which is two dielectric layers wide—and both subsystems. The baseline configuration exhibits an HTC of 1.37×10^5 W/m²K, while the best state found has an HTC of 1.66×10^5 W/m²K, 21% higher.

Figure 8(b) shows the transmission of evanescent waves as a function of the frequency ω and the parallel wave vector k for the optimal configuration. The transmission pattern shows a series of narrow lines of values close to unity, which, as explained in Sec. III, result from the hybridization of the SPPs that are formed in the interfaces between the metallic and dielectric layers, similar to the baseline case in Fig. 2(b). We can notice some differences when comparing the two: while in the baseline case the lines all joined to form the same structure, in the optimal configuration we find a sizeable gap between some of the lines, likely stemming from the supercell structure mentioned above. This implies that their radiative heat transfer behavior is clearly different. In Fig. 8(c) we compare the spectral HTC h_ω of both the baseline [also shown in Fig. 2(c)] and the best state found. We can see that both present the main central resonant peak, but that of the baseline is higher while the best state has a higher value in many other frequencies, with a notable secondary maximum at lower frequencies. The integration of these spectra over frequency yields the total HTC values mentioned at the end of the previous paragraph.

So far, we have focused our study on the case of 16 active layers. This case is complex enough for the illustration of RL techniques, while still being small enough that we only have about 65 000 possible configurations. Thus, we can directly find out which state has the highest HTC by simply analyzing all of them. To show that the RL methods reported here can also be applied to problems in which we do not know the solution beforehand, we consider now the case with 24 active layers, where the total number of states is about 17×10^6 . We used the exact same techniques as before, but taking as the baseline the perfectly periodic case with 24 layers instead of 16. We solved this problem using the double DQN algorithm, the value-based algorithm that showed the highest efficiency in the previous section, for 20 independent runs and 2500 training steps. The results for the largest HTC are shown in Fig. 9. They were obtained using the optimal hyperparameters found for the 16-layer problem. Compared to the results discussed above, the RL algorithm surpasses the random algorithm after a few thousand explored states. In addition, note that the random algorithm reaches a mean value of 1.617×10^5 W/m²K, while the RL algorithm yields 1.673×10^5 W/m²K, which in turn overcomes the value of the baseline state (1.623×10^5 W/m²K). This shows again that we can beat the physical intuition with the proper use of RL, reaching a state with an HTC value of 1.807×10^5 W/m²K. Therefore, this analysis shows that the double DQN algorithm can be successfully applied to

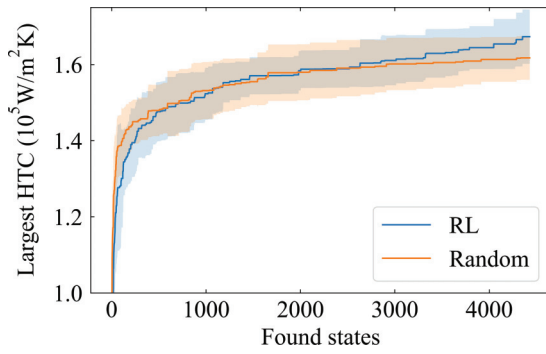


FIG. 9. Largest HTC discovered as a function of the number of found states in the problem with 24 active layers obtained with the double DQN algorithm. We also present the results obtained with the random algorithm. The solid lines correspond to mean values and the shaded area to standard deviations, as obtained in 20 independent runs of both algorithms.

a state space for which the computation of the HTC of all its states is not feasible in a reasonable amount of time.

So, in summary, we have shown in this work how RL can be used to tackle optimization problems in the context of radiative heat transfer. As an illustration, we have addressed the maximization of the NFRHT between hyperbolic metamaterials made of a combination of metallic and dielectric layers. This problem is quite generic and contains the basic ingredients of most optimization problems in the field of thermal radiation. Our work demonstrates that these problems can be naturally formulated as a sequential decision-making problem and, therefore, they are susceptible to be tackled with RL methods. It is worth remarking that one could address inverse design problems in the same way by simply redefining the objective function.

In the physical problem studied in this work, we have shown that essentially all RL algorithms are able to find near optimal solutions, albeit with different efficiencies. In our case, we have found that the double DQN algorithm

is the most efficient, with the PPO algorithm also providing high-quality results. While the PPO algorithm finds the top five HTC values less frequently than the double DQN algorithm, it explores substantially fewer states and is still able to find the best ones much more reliably than the random algorithm. Therefore, both algorithms present their pros and cons in the application to our particular problem. In any case, we have provided a comprehensive guide on how to utilize in practice most of the key RL algorithms. Thus, we hope that our work will help other researchers to employ RL techniques as part of their toolkit for the investigation of optimization and inverse design problems in the context of radiative heat transfer and related topics.

The codes for the different RL algorithms used in this work are available from GitHub [63].

ACKNOWLEDGMENTS

E.O.-M. thanks the Komorebi AI Technologies team, with special mention to its founders for their broad coding insights and discussions regarding some of the reported results. Furthermore, she acknowledges financial support from the Comunidad de Madrid and Komorebi AI Technologies under the Program “Doctorados Industriales” (reference IND2022/IND-23536). J.J.G.-E. was supported by the Spanish Ministry of Science and Innovation (Grant No. FPU19/05281). J.B.-A. acknowledges funding from the Spanish Ministry of Science and Innovation (PID2022-139995NB-I00). J.C.C. thanks the Spanish Ministry of Science and Innovation for financial support through Grant No. PID2020-114880GB-I00 and the “María de Maeztu” Programme for Units of Excellence in R&D (CEX2023-001316-M).

APPENDIX: HYPERPARAMETERS

In Tables I–V we summarize the hyperparameters used in the different algorithms employed in this work. In particular, we describe their meanings and roles.

TABLE I. Hyperparameters of the SARSA algorithm. SELU stands for scaled exponential linear unit and MAE for mean absolute error.

Variable name	Value	Description
hidden layers	4	Number of hidden layers of the neural network
hidden neurons	64	Number of neurons of each one of the hidden layers
activation function	SELU	Activation function of the neurons
loss function	MAE	Loss function to evaluate the performance of the neural network
optimizer	Adam	Optimizer during training
γ	0.99	Discount rate parameter
episode length	64	Number of steps per episode
ϵ	1	Initial value of the epsilon parameter
ϵ decay	4×10^{-5}	Value to reduce epsilon each training step
ϵ minimum	10^{-3}	Minimum reachable epsilon value
α	10^{-3}	Value of the learning rate parameter
α decay	0	Value to reduce the learning rate each training step
batch size	32	Number of experiences to process each training step

TABLE II. Hyperparameters of the double DQN algorithm.

Variable name	Value	Description
hidden layers	4	Number of hidden layers of the neural network
hidden neurons	64	Number of neurons of each one of the hidden layers
activation function	SELU	Activation function of the neurons
loss function	MAE	Loss function to evaluate the performance of the neural network
optimizer	Adam	Optimizer during training
γ	0.99	Discount rate parameter
episode length	64	Number of steps per episode
ε	1	Initial value of the epsilon parameter
ε decay	9×10^{-5}	Epsilon decay each training step
ε minimum	10^{-3}	Minimum reachable epsilon value
α	10^{-4}	Value of the learning rate parameter
α decay	0	Value to reduce the learning rate each training step
batch size	32	Number of experiences to process each training step
B	4	Number of batches processed each training step
K	10^4	Memory size, maximum number of experiences stored in memory
h	4	Number of new experiences added each training step
F	5×10^3	Number of train steps to update the target network

TABLE III. Hyperparameters of the REINFORCE algorithm.

Variable name	Value	Description
hidden layers	4	Number of hidden layers of the neural network
hidden neurons	64	Number of neurons of each one of the hidden layers
activation function	SELU	Activation function of the neurons
optimizer	Adam	Optimizer during training
α	3×10^{-5}	Value of the learning rate parameter
γ	0.99	Discount rate parameter
episode length	32	Number of experiences in an episode
n episodes	10^5	Number of episodes in training

TABLE IV. Hyperparameters of the A2C algorithm.

Variable name	Value	Description
hidden layers	4	Number of hidden layers of the neural network; same for the actor and critic networks
hidden neurons	100	Number of neurons of each one of the hidden layers; same for the actor and critic networks
activation function	SELU	Activation function of the neurons
loss function	MAE	Loss function to evaluate the performance of the critic network
optimizer	Adam	Optimizer during training
episode length	32	Number of experiences in an episode
n episodes	10^5	Number of episodes in training
gradient clip	True	Use of gradient clipping
vf coef	1/2	Constant to weight the value function loss
use rms prop	False	Use RMSProp instead of Adam
γ	0.946 45	Discount rate parameter
α	3.53×10^{-4}	Value of the learning rate parameter
λ	0.997	Value of the GAE exponential factor
max grad	1.285	Maximum gradient in the optimization step, gradient clipping

TABLE V. Hyperparameters of the PPO algorithm.

Variable name	Value	Description
hidden layers	4	Number of hidden layers of the neural network; same for the actor and critic networks
hidden neurons	100	Number of neurons of each one of the hidden layers; same for the actor and critic networks
activation function	SELU	Activation function of the neurons
loss function	MAE	Loss function to evaluate the performance of the critic network
optimizer	Adam	Optimizer during training
episode length	32	Number of experiences in an episode
n episodes	10^5	Number of episodes in training
ϵ	0.2	Clipping range for the importance sampling terms
batch size	32	Batch size for the PPO algorithm training
gradient clip	True	Use of gradient clipping
vf coef	1/2	Constant to weight the value function loss
γ	0.99988	Discount rate parameter
α	6.24×10^{-5}	Value of the learning rate parameter
λ	0.871	Value of the GAE exponential factor
max grad	4.042	Maximum gradient in the optimization step, gradient clipping

TABLE VI. Hyperparameters of the Optuna search algorithm for A2C and PPO.

Variable name	Value	Description
n configs	100	Number of hyperparameter configurations studied
n steps	10^5	Number of experiences explored each configuration
sampler	TPESampler	Sampler chosen to choose the next configuration
pruner	Median	Method of pruning a configuration choice
n startup	5	Number of configurations before starting the pruner
n warmup	3.33×10^4	Experiences before checking for pruning
γ range	[0.9, 0.9999]	Range of exploration values for the γ parameter
α range	$[10^{-6}, 10^{-3}]$	Range of exploration values for the α parameter
λ range	[0.0, 1.0]	Range of exploration values for the λ parameter
max grad range	[0.5, 5.0]	Range of exploration values for the max grad parameter

[1] M. F. Modest, *Radiative Heat Transfer* (Elsevier, Academic Press, 2013).

[2] J. R. Howell, M. P. Mengüç, K. Daun, and R. Siegel, *Thermal Radiation Heat Transfer* (CRC press, Boca Raton, 2020).

[3] Z. M. Zhang, *Nano/Microscale Heat Transfer* (Springer, Cham (Switzerland), 2020).

[4] J. C. Cuevas and F. J. García-Vidal, Radiative heat transfer, *ACS Photonics* **5**, 3896 (2018).

[5] D. Polder and M. Van Hove, Theory of radiative heat transfer between closely spaced bodies, *Phys. Rev. B* **4**, 3303 (1971).

[6] B. Song, A. Fiorino, E. Meyhofer, and P. Reddy, Near-field radiative thermal transport: from theory to experiment, *AIP Adv.* **5**, 053503 (2015).

[7] S.-A. Biehs, R. Messina, P. S. Venkataram, A. W. Rodriguez, J. C. Cuevas, and P. Ben-Abdallah, Near-field radiative heat transfer in many-body systems, *Rev. Mod. Phys.* **93**, 025009 (2021).

[8] W. Challener, C. Peng, A. Itagi, D. Karns, W. Peng, Y. Peng, X. Yang, X. Zhu, N. Gokemeijer, and Y.-T. Hsia *et al.*, Heat-assisted magnetic recording by a near-field transducer

with efficient optical energy transfer, *Nat. Photonics* **3**, 220 (2009).

[9] Y. De Wilde, F. Formanek, R. Carminati, B. Gralak, P.-A. Lemoine, K. Joulain, J.-P. Mulet, Y. Chen, and J.-J. Greffet, Thermal radiation scanning tunnelling microscopy, *Nature* **444**, 740 (2006).

[10] A. Kittel, U. F. Wischnath, J. Welker, O. Huth, F. Rueting, and S.-A. Biehs, Near-field thermal imaging of nanostructured surfaces, *Appl. Phys. Lett.* **93**, 193109 (2008).

[11] A. C. Jones, B. T. O’Callahan, H. U. Yang, and M. B. Raschke, The thermal near-field: Coherence, spectroscopy, heat-transfer, and optical forces, *Prog. Surf. Sci.* **88**, 349 (2013).

[12] R. Carminati and J.-J. Greffet, Near-field effects in spatial coherence of thermal sources, *Phys. Rev. Lett.* **82**, 1660 (1999).

[13] J.-J. Greffet, R. Carminati, K. Joulain, J.-P. Mulet, S. Mainy, and Y. Chen, Coherent emission of light by thermal sources, *Nature* **416**, 61 (2002).

[14] R. Mittapally, A. Majumder, P. Reddy, and E. Meyhofer, Near-field thermophotovoltaic energy conversion: Progress and opportunities, *Phys. Rev. Appl.* **19**, 037002 (2023).

[15] W. Li and S. Fan, Nanophotonic control of thermal radiation for energy applications, *Opt. Express* **26**, 15995 (2018).

- [16] E. Rephaeli, A. Raman, and S. Fan, Ultrabroadband photonic structures to achieve high-performance daytime radiative cooling, *Nano Lett.* **13**, 1457 (2013).
- [17] A. P. Raman, M. A. Anoma, L. Zhu, E. Rephaeli, and S. Fan, Passive radiative cooling below ambient air temperature under direct sunlight, *Nature* **515**, 540 (2014).
- [18] J. K. Tong, X. Huang, S. V. Boriskina, J. Loomis, Y. Xu, and G. Chen, Infrared-transparent visible-opaque fabrics for wearable personal thermal management, *ACS Photonics* **2**, 769 (2015).
- [19] P.-C. Hsu, A. Y. Song, P. B. Catrysse, C. Liu, Y. Peng, J. Xie, S. Fan, and Y. Cui, Radiative human body cooling by nanoporous polyethylene textile, *Science* **353**, 1019 (2016).
- [20] W. Li, Y. Shi, K. Chen, L. Zhu, and S. Fan, A comprehensive photonic approach for solar cell cooling, *ACS Photonics* **4**, 774 (2017).
- [21] A. Lenert, D. M. Bierman, Y. Nam, W. R. Chan, I. Celanović, M. Soljačić, and E. N. Wang, A nanophotonic solar thermophotovoltaic device, *Nat. Nanotechnol.* **9**, 126 (2014).
- [22] S.-A. Biehs and P. Ben-Abdallah, Revisiting super-Planckian thermal emission in the far-field regime, *Phys. Rev. B* **93**, 165405 (2016).
- [23] V. Fernández-Hurtado, A. I. Fernández-Domínguez, J. Feist, F. J. García-Vidal, and J. C. Cuevas, Super-Planckian far-field radiative heat transfer, *Phys. Rev. B* **97**, 045408 (2018).
- [24] D. Thompson, L. Zhu, R. Mittapally, S. Sadat, Z. Xing, P. McArdle, M. M. Qazilbash, P. Reddy, and E. Meyhofer, Hundred-fold enhancement in far-field radiative heat transfer over the blackbody limit, *Nature* **561**, 216 (2018).
- [25] P. Chao, B. Strekha, R. Kuate Defo, S. Molesky, and A. W. Rodriguez, Physical limits in electromagnetism, *Nat. Rev. Phys.* **4**, 543 (2022).
- [26] S. Molesky, Z. Lin, A. Y. Piggott, W. Jin, J. Vucković, and A. W. Rodriguez, Inverse design in nanophotonics, *Nat. Photonics* **12**, 659 (2018).
- [27] P. Mehta, M. Bukov, C.-H. Wang, A. G. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, A high-bias, low-variance introduction to machine learning for physicists, *Phys. Rep.* **810**, 1 (2019).
- [28] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, Machine learning and the physical sciences, *Rev. Mod. Phys.* **91**, 045002 (2019).
- [29] F. Marquardt, Machine learning and quantum devices, *SciPost Phys. Lect. Notes*, 29 (2021).
- [30] J. J. García-Esteban, J. Bravo-Abad, and J. C. Cuevas, Deep learning for the modeling and inverse design of radiative heat transfer, *Phys. Rev. Appl.* **16**, 064006 (2021).
- [31] C. Zhu, E. A. Bamidele, X. Shen, G. Zhu, and B. Li, Machine learning aided design and optimization of thermal metamaterials, *Chem. Rev.* **124**, 4258 (2024).
- [32] S. Yu, P. Zhou, W. Xi, Z. Chen, Y. Deng, X. Luo, W. Li, J. Shiomi, and R. Hu, General deep learning framework for emissivity engineering, *Light: Sci. Appl.* **12**, 291 (2023).
- [33] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT press, Cambridge, MA, 2018).
- [34] L. Graesser and W. L. Keng, *Foundations of Deep Reinforcement Learning: Theory and Practice in Python* (Addison-Wesley Professional, Boston, 2019).
- [35] J. J. García-Esteban, J. C. Cuevas, and J. Bravo-Abad, Generative adversarial networks for data-scarce radiative heat transfer applications, *Mach. Learn.: Sci. Technol.* **5**, 015060 (2024).
- [36] B. Song, A. Fiorino, E. Meyhofer, and P. Reddy, Near-field radiative thermal transport: from theory to experiment, *AIP Adv.* **5**, 053503 (2015).
- [37] Y. Guo, C. L. Cortes, S. Molesky, and Z. Jacob, Broadband super-Planckian thermal emission from hyperbolic metamaterials, *Appl. Phys. Lett.* **101**, 131106 (2012).
- [38] S.-A. Biehs, M. Tschikin, and P. Ben-Abdallah, Hyperbolic metamaterials as an analog of a blackbody in the near field, *Phys. Rev. Lett.* **109**, 104301 (2012).
- [39] Y. Guo and Z. Jacob, Thermal hyperbolic metamaterials, *Opt. Express* **21**, 15014 (2013).
- [40] S.-A. Biehs, M. Tschikin, R. Messina, and P. Ben-Abdallah, Super-Planckian near-field thermal emission with phonon-polaritonic hyperbolic metamaterials, *Appl. Phys. Lett.* **102**, 131106 (2013).
- [41] T. J. Bright, X. L. Liu, and Z. M. Zhang, Energy streamlines in near-field radiative heat transfer between hyperbolic metamaterials, *Opt. Express* **22**, A1112 (2014).
- [42] O. D. Miller, S. G. Johnson, and A. W. Rodriguez, Effectiveness of thin films in lieu of hyperbolic metamaterials in the near field, *Phys. Rev. Lett.* **112**, 157402 (2014).
- [43] S.-A. Biehs and P. Ben-Abdallah, Near-field heat transfer between multilayer hyperbolic metamaterials, *Z. Naturforsch., A* **72**, 115 (2017).
- [44] H. Iizuka and S. Fan, Significant enhancement of near-field electromagnetic heat transfer in a multilayer structure through multiple surface-states coupling, *Phys. Rev. Lett.* **120**, 063901 (2018).
- [45] J. Song, Q. Cheng, L. Lu, B. Li, K. Zhou, B. Zhang, Z. Luo, and X. Zhou, Magnetically tunable near-field radiative heat transfer in hyperbolic metamaterials, *Phys. Rev. Appl.* **13**, 024054 (2020).
- [46] E. Moncada-Villa and J. C. Cuevas, Near-field radiative heat transfer between one-dimensional magnetophotonic crystals, *Phys. Rev. B* **103**, 075432 (2021).
- [47] S. Rylov, *Theory of Electric Fluctuations and Thermal Radiation* (Air Force Cambridge Research Center, Bedford, MA, 1953).
- [48] S. Rylov, Y. Kravtsov, and V. Tatarskii, *Principles of Statistical Radiophysics* (Springer-Verlag, Berlin, 1989).
- [49] S. Basu, Z. M. Zhang, and C. J. Fu, Review of near-field thermal radiation and its application to energy conversion, *Int. J. Energy Res.* **33**, 1203 (2009).
- [50] B. Caballero, A. García-Martín, and J. C. Cuevas, Generalized scattering-matrix approach for magneto-optics in periodically patterned multilayer systems, *Phys. Rev. B* **85**, 245103 (2012).
- [51] L.-J. Lin, Self-improving reactive agents based on reinforcement learning, planning and teaching, *Mach. Learn.* **8**, 293 (1992).
- [52] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, Playing Atari with deep reinforcement learning, [arXiv:1312.5602](https://arxiv.org/abs/1312.5602).
- [53] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A.

- K. Fidjeland, and G. Ostrovski *et al.*, Human-level control through deep reinforcement learning, *Nature* **518**, 529 (2015).
- [54] H. van Hasselt, A. Guez, and D. Silver, Deep reinforcement learning with double Q-learning, [arXiv:1509.06461](https://arxiv.org/abs/1509.06461).
- [55] H. Hasselt, in *Advances in Neural Information Processing Systems*, edited by J. Lafferty, C. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta (Curran Associates, Inc., Red Hook, NY, 2010), Vol. 23, https://proceedings.neurips.cc/paper_files/paper/2010/file/091d584fced301b442654dd8c23b3fc9-Paper.pdf.
- [56] R. J. Williams, Simple statistical gradient-following algorithms for connectionist reinforcement learning, *Mach. Learn.* **8**, 229 (1992).
- [57] V. Konda and J. Tsitsiklis, in *Advances in Neural Information Processing Systems*, edited by S. Solla, T. Leen, and K. Müller (MIT Press, Rochester, NY, 1999), Vol. 12, https://proceedings.neurips.cc/paper_files/paper/1999/file/6449f44a102fde848669bdd9eb6b76fa-Paper.pdf.
- [58] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, Proximal policy optimization algorithms, [arXiv:1707.06347](https://arxiv.org/abs/1707.06347).
- [59] 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).
- [60] J. Achiam, D. Held, A. Tamar, and P. Abbeel, Constrained policy optimization, [arXiv:1705.10528](https://arxiv.org/abs/1705.10528).
- [61] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, Stable-baselines3: Reliable reinforcement learning implementations, *J. Mach. Learn. Res.* **22**, 1 (2021).
- [62] T. Akiba, S. Sano, T. Yanase, T. Ohta, and M. Koyama, Optuna: A next-generation hyperparameter optimization framework, [arXiv:1907.10902](https://arxiv.org/abs/1907.10902).
- [63] https://github.com/jjgarciae/Deep_RL_RHT.