

Lawrence Berkeley National Laboratory

LBL Publications

Title

Optimizing thermodynamic trajectories using evolutionary and gradient-based reinforcement learning

Permalink

<https://escholarship.org/uc/item/7cp034r7>

Journal

Physical Review E, 104(6)

ISSN

2470-0045

Authors

Beeler, Chris
Yahorau, Uladzimir
Coles, Rory
et al.

Publication Date

2021-12-01

DOI

10.1103/physreve.104.064128

Peer reviewed

Optimizing thermodynamic trajectories using evolutionary and gradient-based reinforcement learning

Chris Beeler^{1,2,*}, Uladzimir Yahorau³, Rory Coles⁴, Kyle Mills^{3,5}, Stephen Whitlam⁶, and Isaac Tamblyn^{1,3,5†}

¹*University of Ottawa, Ottawa, ON, Canada*

²*National Research Council of Canada, Ottawa, ON, Canada*

³*University of Ontario Institute of Technology,
Oshawa, ON, Canada*

⁴*University of Victoria, Victoria, BC, Canada*

⁵*Vector Institute for Artificial Intelligence,
Toronto, Ontario, Canada*

⁶*Molecular Foundry,*

Lawrence Berkeley National Laboratory, Berkeley, CA, USA

(Dated: November 22, 2021)

Using a model heat engine, we show that neural network-based reinforcement learning can identify thermodynamic trajectories of maximal efficiency. We consider both gradient and gradient-free reinforcement learning. We use an evolutionary learning algorithm to evolve a population of neural networks, subject to a directive to maximize the efficiency of a trajectory composed of a set of elementary thermodynamic processes; the resulting networks learn to carry out the maximally-efficient Carnot, Stirling, or Otto cycles. When given an additional irreversible process, this evolutionary scheme learns a previously unknown thermodynamic cycle. Gradient-based reinforcement learning is able to learn the Stirling cycle, whereas an evolutionary approach achieves the optimal Carnot cycle. Our results show how the reinforcement learning strategies developed for game playing can be applied to solve physical problems conditioned upon path-extensive order parameters.

I. INTRODUCTION

Games, whether played on a board, such as chess or Go, or played on the computer, are a major component of human culture [1]. In the language of physics, instances of a game are *trajectories*, time-ordered sequences of elementary steps. The outcome of a game is a path-extensive order parameter determined by the entire history of the trajectory. Playing games was once the preserve of human beings, but machine learning methods, more specifically reinforcement learning, now outperform the most talented humans in all the aforementioned examples [2–19]. While there are more scientific examples of reinforcement learning [20–25], we choose to compare with games as they are the easiest applications to understand. Motivated by the correspondence between games and trajectories, it is natural to ask how the machine-learning and optimal control methods that have mastered game-playing might be applied to understand physical processes whose outcomes are path-extensive quantities.

There are many examples of such processes. For instance, the success or failure of molecular self-assembly is determined by a time history of elementary dynamical processes, including the binding and unbinding of particles [26–29]. Dynamical systems, such as chemical networks and molecular machines [30–33], are characterized by time-extensive observables, such as work or entropy production [34–40]. In none of these cases do we possess a complete theoretical or practical understanding of how

to build an arbitrary structure or maximize the efficiency of an arbitrary machine. Traditional methods of inquiry in physics focus on applying physical intuition and the manipulation and simulation of equations; perhaps machine learning can provide us with further insight into physical problems of a path-extensive nature.

Motivated by this speculation, we show here that neural network-based reinforcement learning can maximize the efficiency of the simplest type of physical trajectories, the deterministic, quasi-static ones of classical thermodynamics. We introduce a model heat engine characterized by a set of thermodynamic state variables. A neural network takes as input the current observation of the engine and chooses one of a set of basic thermodynamic processes to produce a new observation; this change comprises one step of a trajectory. We chose this simple and well-known system for pedagogical purposes. Using an easy-to-understand system allows us to focus on how reinforcement learning methods can be applied to physics problems. We generate a set of trajectories of fixed length using a set of networks whose parameters are initially randomly chosen and optimized using two different methods. In the first method (gradient-free), we retain and mutate only those networks whose trajectories show the greatest thermal efficiency. Repeating this evolutionary process many times results in networks whose trajectories reproduce the maximally efficient Carnot, Stirling, or Otto cycles, depending upon which basic thermodynamic processes are allowed. This evolutionary procedure can also learn previously unknown thermodynamic cycles if new processes are allowed. In the second method (gradient-based), we update the network parameters using gradient-based reinforcement learning. In this study,

* christopher.beeler@uottawa.ca

† isaac.tamblyn@uottawa.ca

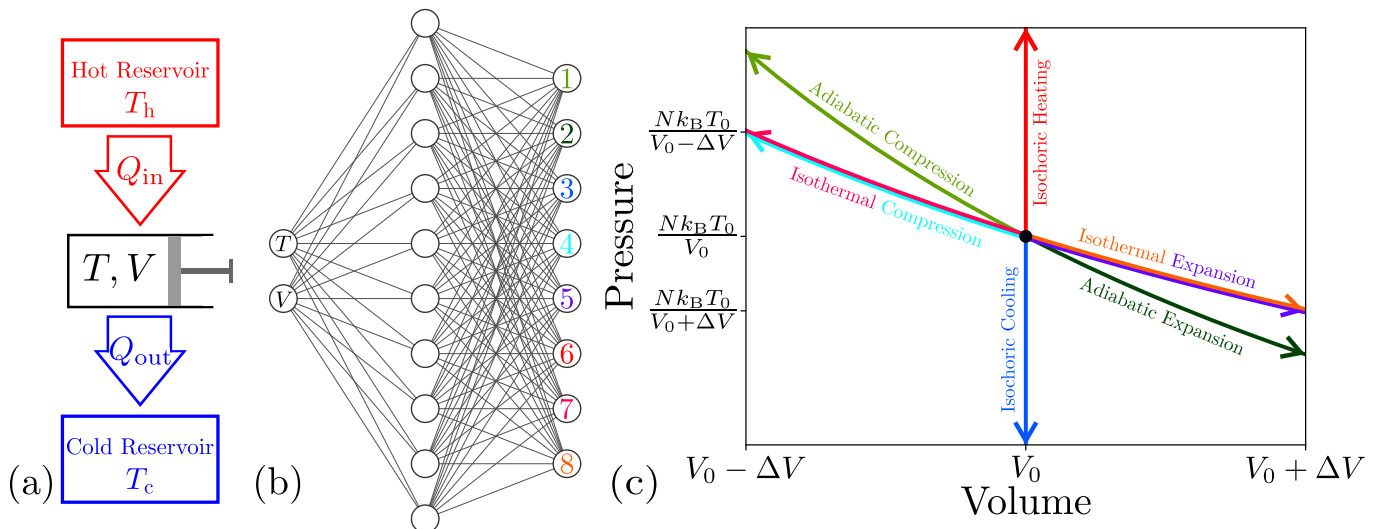


FIG. 1. (a) Model heat engine and (b) the neural network that controls it. Note that the diagram of the neural network has two additional input nodes for the gradient-based reinforcement learning method. (c) A summary of the actions in P - V space available to the network; see Tab. I.

we explore how a reinforcement learning problem must be framed and the types of solutions acquired by these two methods. We also compare the advantages and disadvantages of each in finding a solution to this problem.

II. MODEL HEAT ENGINE AND THERMODYNAMIC TRAJECTORIES

In Fig. 1(a) we show a model heat engine, a device able to transform thermal energy into work [41, 42]. The engine consists of a working substance, which we assume to be a monatomic ideal gas, housed within a frictionless container of variable volume V , whose minimum and maximum values are V_{\min} and V_{\max} , respectively. The working substance may be connected to a hot or cold reservoir held at temperature, $T_h = 500$ K and $T_c = 300$ K, respectively. For the gradient-free method, the instantaneous observation s_t of the system at time t is then specified by the volume-temperature vector $s_t = (V, T)$, with the pressure of the system fixed by the ideal-gas equation $PV = Nk_B T$ [43]. Further on, we will discuss the reasons this formulation of the problem cannot be used for gradient-based reinforcement learning methods.

To evolve the heat engine we use the neural network shown in Fig. 1(b). The network is a nonlinear function that takes as input the current observation s_t of the system, and outputs the probabilities $\pi_\theta(a_t|s_t)$ of moving to any new observation s_{t+1} through a thermodynamic process $a_t \in \{a_1, a_2, \dots, a_M\}$ (in the language of reinforcement learning this mapping is called a *policy* [44]). The symbol θ denotes the internal parameters of the network, discussed shortly. Here we consider deterministic evolution through configuration space, with $\pi_\theta(a_t^*|s_t)$ equal to 1 for a chosen process a_t^* , and equal to zero

otherwise. Enacting the chosen process corresponds to one step of a trajectory. Given an initial observation s_0 , K applications of the network produces a trajectory $\omega = s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_K$ of K steps through configuration space. We focus on trajectories of fixed length ($K = 200$).

The elementary actions available to the network correspond to the basic thermodynamic processes shown in Tab. I, summarized graphically in Fig. 1(c). These processes include reversible compression and expansion, along isotherms or adiabats, and reversible temperature changes along isochores. Implicitly this means infinitely many heat baths spaced between T_h and T_c are available, a condition, which we shall prove in Appendix A, does not undermine Carnot's Theorem of maximum efficiency. All compression and expansion processes are performed using a fixed change in volume. If an isothermal process at T_h (or T_c) is selected when the system is not at the correct temperature, then isochoric heating (or cooling) is performed first to reach the necessary temperature for the isothermal process to occur. Note that the isochoric processes are not used in the Carnot cycle, but are required to make approximations of it when using fixed changes in volume because the system must reach specific volumes in order to be adiabatically heated from T_c to T_h and cooled from T_h to T_c .

Upon undertaking any action $s_t \rightarrow s_{t+1}$, we record the resulting changes of work, $\Delta W_{s_t s_{t+1}}$, and heat input from the hot reservoir, $\Delta Q_{s_t s_{t+1}}^{\text{in}} = \Delta Q_{s_t s_{t+1}} H(\Delta Q_{s_t s_{t+1}} \delta_{T_f, T_h})$; these are listed in Tab. I. Here $H(\cdot)$ is the Heaviside function, equal to 1 for positive values of $\Delta Q_{s_t s_{t+1}}$ and 0 otherwise, and T_f is the temperature of the system following the move. $\delta_{\alpha, \beta}$ is the Kronecker delta (1 if $\alpha = \beta$ and 0 otherwise). We define the thermodynamic efficiency of a K -step trajectory

as

$$\eta_K \equiv \frac{\sum_{t=0}^{K-1} \Delta W_{s_t s_{t+1}}}{\sum_{t=0}^{K-1} \Delta Q_{s_t s_{t+1}}^{\text{in}}}. \quad (1)$$

The thermal efficiency, a path-extensive quantity, is used as a means of ranking trajectories, and the networks that generate them, during our evolutionary learning procedure. The neural network selects processes deterministically based on observations, therefore once a trajectory produces a single cycle, that cycle will repeat until the maximum number of trajectory steps have occurred. For this reason, the maximum value $\eta = \max_K \eta_K$ for all K points along a long trajectory is sufficient to identify efficient thermodynamic cycles. We could terminate a trajectory after it produces a single cycle, however we chose this way for consistency with the gradient-based reinforcement learning method used where we do require multiple cycles.

III. NEURAL NETWORK-BASED POLICY

The neural network, which contains two layers of tunable weights, performs computations as follows. Two input neurons receive the current observation s_t , and the output is comprised of $M \leq 8$ neurons, each corresponding to one of the actions shown in Tab. I (in some simulations we prohibit certain actions). The network possesses one hidden layer of 1024 neurons, each connected to every input and output neuron. This architecture was chosen after several tests as it can produce optimal results while still keeping a low computational cost relative to larger (more parameters) networks that produce similar results. Let the indices $i \in \{1, 2\}$, $j \in \{1, \dots, 1024\}$, and $k \in \{1, \dots, M\}$ label the neurons of the input, hidden, and output layers, respectively. The input I_i of the two nodes $i = 0, 1$ of the input layer are, respectively, scaled versions of the current temperature $(T - T_c)/(T_h - T_c) \in [0, 1]$ and volume $(V - V_{\min})/(V_{\max} - V_{\min}) \in [0, 1]$ of the system. We set the output signal S_i of each input-layer node as $S_i = I_i$.

The input I_j to neuron j in the hidden layer is

$$I_j = \sum_{i=1}^2 S_i w_{ij}, \quad (2)$$

where the sum runs over the two neurons in the input layer, and w_{ij} is the weight of the connection between nodes i and j . We set the output signal S_j of neuron j to be

$$S_j = \frac{1}{2} [\tanh(I_j + b_j)], \quad (3)$$

where b_j is a bias associated with neuron j .

The input I_k to neuron k in the output layer is

$$I_k = \sum_{j=1}^{1024} S_j w_{jk}, \quad (4)$$

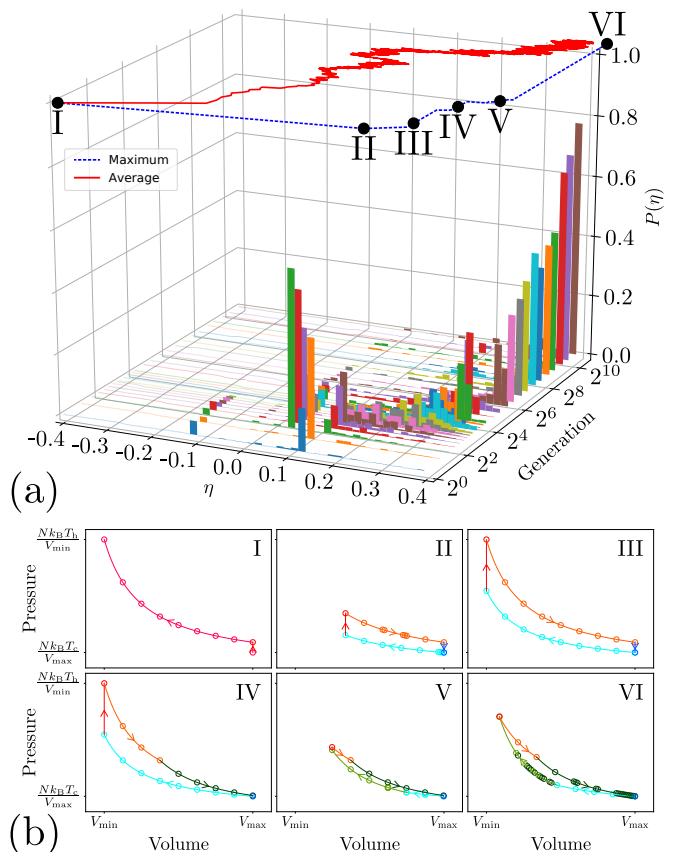


FIG. 2. (a) The evolution, as a function of generation number, of the probability distribution $P(\eta)$ of trajectories of the model heat engine. The maximum and average efficiency of the population are shown above. The Carnot efficiency is $\eta_{\max} = 0.4$. (b) Trajectories in P - V space produced by the best-performing networks in generations $2^0, 2^1, 2^2, 2^4, 2^5,$ and 2^{12} , in the boxes labeled I-VI, respectively. The colors (shades), curvature, and direction of the branches correspond to the processes shown in Fig. 1. Highly-evolved networks enact the Carnot cycle.

where the sum runs over all 1024 neurons of the hidden layer. Finally, we take the output signal S_k from each output-layer neuron to be equal to I_k . To choose an action we pick the output neuron, k^* , with the largest value of S_k . Given a current observation s_t , this action a_t^* defines a new observation s_{t+1}^* via Tab. I. The probability $\pi_\theta(a_t^*|s_t)$ is then unity, and all other $\pi_\theta(a_t|s_t)$ are zero. We denote by $\theta = \{\{w\}, \{b\}\}$ the set of all weights and biases of the network. Initially each weight and bias is chosen from a Gaussian distribution with zero mean and unit variance.

IV. EVOLUTIONARY LEARNING DYNAMICS (GRADIENT-FREE)

With the thermodynamic system and means of evolving it defined, we introduce an evolutionary learning dy-

TABLE I. All possible actions that can be taken on our model heat engine and their corresponding ΔW and ΔQ equations.

Action	ΔW	ΔQ
Adiabatic Compression	$-\frac{3}{2}Nk_B T_i \left(\left(\frac{V_i}{V_f} \right)^{\frac{2}{3}} - 1 \right)$	0
Adiabatic Expansion	$-\frac{3}{2}Nk_B T_i \left(\left(\frac{V_i}{V_f} \right)^{\frac{2}{3}} - 1 \right)$	0
Isothermal Compression at T_h ($T = T_h$)	$Nk_B T_h \log \left(\frac{V_f}{V_i} \right)$	$Nk_B T_h \log \left(\frac{V_f}{V_i} \right)$
Isothermal Expansion at T_h ($T = T_h$)	$Nk_B T_h \log \left(\frac{V_f}{V_i} \right)$	$Nk_B T_h \log \left(\frac{V_f}{V_i} \right)$
Isothermal Compression at T_h ($T \neq T_h$)	$Nk_B T_h \log \left(\frac{V_f}{V_i} \right)$	$Nk_B T_h \log \left(\frac{V_f}{V_i} \right) + \frac{3}{2}Nk_B (T_h - T_i)$
Isothermal Expansion at T_h ($T \neq T_h$)	$Nk_B T_h \log \left(\frac{V_f}{V_i} \right)$	$Nk_B T_h \log \left(\frac{V_f}{V_i} \right) + \frac{3}{2}Nk_B (T_h - T_i)$
Isothermal Compression at T_c ($T = T_c$)	$Nk_B T_c \log \left(\frac{V_f}{V_i} \right)$	$Nk_B T_c \log \left(\frac{V_f}{V_i} \right)$
Isothermal Expansion at T_c ($T = T_c$)	$Nk_B T_c \log \left(\frac{V_f}{V_i} \right)$	$Nk_B T_c \log \left(\frac{V_f}{V_i} \right)$
Isothermal Compression at T_c ($T \neq T_c$)	$Nk_B T_c \log \left(\frac{V_f}{V_i} \right)$	$Nk_B T_c \log \left(\frac{V_f}{V_i} \right) + \frac{3}{2}Nk_B (T_c - T_i)$
Isothermal Expansion at T_c ($T \neq T_c$)	$Nk_B T_c \log \left(\frac{V_f}{V_i} \right)$	$Nk_B T_c \log \left(\frac{V_f}{V_i} \right) + \frac{3}{2}Nk_B (T_c - T_i)$
Isochoric Heating	0	$\frac{3}{2}Nk_B (T_h - T_i)$
Isochoric Cooling	0	$\frac{3}{2}Nk_B (T_c - T_i)$

namics designed to produce networks able to propagate efficient thermodynamic trajectories. We start with a population of 100 networks, with the internal parameters θ of each initialized in the random fashion described above. We name this population generation 1. This population produces thermodynamic trajectories ω of K steps with the distribution $P(\eta)$ of efficiencies η shown in Fig. 2(a). Some of the networks in earlier generations do not appear in these distributions due to their undefined or very negative efficiencies. Even the best-performing members of this population produce efficiencies much lower than the Carnot efficiency, which is the most efficient trajectory possible given the set of allowed thermodynamic processes [42]. We have $\eta_{\max} = 0.4$ with our choice of parameters.

We next perform the first step of evolutionary learning dynamics. We keep the 25 generation-1 networks whose trajectories have the largest η , and we discard the rest. We create 75 new networks by drawing uniformly from the set of 25, each time “mutating” all weights w and biases b : for each weight or bias we draw a random number δ from a Gaussian distribution with zero mean and unit variance, and update the weight or bias as $w \rightarrow w + \epsilon\delta$ or $b \rightarrow b + \epsilon\delta$, where $\epsilon = 0.05$ is an evolutionary learning rate.

The new population of the 25 best generation-1 networks and their 75 mutant offspring constitute generation 2. We simulate those 100 networks for K steps, producing the distribution of efficiencies shown in Fig. 2(a). Continuing this alternation of evolutionary dynamics (retaining and mutating the best networks of the current generation) and physical dynamics (using the new generation of networks to generate a set of trajectories) gives rise to networks able to propagate increasingly efficient trajectories [Fig. 2(a)]. After about 100 generations, we obtain networks whose efficiencies are equal to that of

the Carnot cycle (to within four decimal places). Inspection of the trajectories corresponding to these values of η show that they indeed form Carnot cycles; see Fig. 2(b).

Several features of this learning process are notable. In learning to maximize the efficiency of a thermodynamic trajectory, networks have learned to propagate cycles, as opposed to non-closed trajectories in P - V space, because cycles lead in general to larger efficiencies. As the thermal efficiency of the Carnot cycle, defined as

$$\eta_{\max} = 1 - \frac{T_c}{T_h}, \quad (5)$$

is independent of volume, the Carnot cycle has no absolute scale associated with it, meaning the Carnot cycle on two heat engines with two distinct total volumes are equivalent. However, this is not true for the Stirling, Otto, or the discretized approximations of the Carnot cycles, whose thermal efficiencies can be simply derived using the Eq. (1) with the processes in Tab. I [41] and are defined as

$$\eta_S = \frac{T_h - T_c}{T_h + \frac{3(T_h - T_c)}{2 \log(V_r)}}, \quad (6)$$

for the Stirling cycle, where $V_r = V_{\max}/V_{\min}$,

$$\eta_O = 1 - V_r^{-\frac{2}{3}}, \quad (7)$$

for the Otto cycle, and

$$\eta'_C = \frac{T_h \left(\log \left(\frac{V_2}{V_{\min}} \right) - \frac{3}{2} \left(\left(\frac{V_2}{V_{\max}} \right)^{\frac{2}{3}} - 1 \right) \right)}{\frac{3}{2} \left(T_h - T_c \left(\frac{V_1}{V_{\min}} \right)^{\frac{2}{3}} \right) + T_h \log \left(\frac{V_2}{V_{\min}} \right)} + \frac{T_c \left(\log \left(\frac{V_1}{V_{\max}} \right) - \frac{3}{2} \left(\left(\frac{V_1}{V_{\min}} \right)^{\frac{2}{3}} - 1 \right) \right)}{\frac{3}{2} \left(T_h - T_c \left(\frac{V_1}{V_{\min}} \right)^{\frac{2}{3}} \right) + T_h \log \left(\frac{V_2}{V_{\min}} \right)} \quad (8)$$

for the discretized Carnot cycle, where V_1 is the volume when adiabatic compression begins, and V_2 is the volume when adiabatic expansion begins. If $V_1 = (T_h/T_c)^{3/2} V_{\min}$ and $V_2 = (T_c/T_h)^{3/2} V_{\max}$ then the volume terms cancel and we get $\eta'_C = \eta_{\max}$. Otherwise, all of these equations contain volume terms and therefore the Stirling, Otto, and discretized Carnot cycles are volume dependent. Interestingly, η_O is independent of the thermal baths, however it can only be performed if $T_h > T_c V^{2/3}$. When this is the case we have

$$\eta_{\max} = 1 - \frac{T_c}{T_h} > 1 - \frac{T_c}{T_c V_r^{\frac{2}{3}}} = 1 - V_r^{-\frac{2}{3}} = \eta_O, \quad (9)$$

however the Carnot cycle is not technically possible in this scenario as it is impossible to heat the system the required amount with adiabatic processes, so we only briefly consider it. For fixed T_h and T_c , η_S and η_O are maximized by maximizing the volume ratio V_r , however η'_C is maximized by minimizing the difference between V_1 and V_2 with the volume values used in the true Carnot cycle. Given the fixed step sizes permitted for the processes in Tab. I, networks have learned to enact the size of a cycle that allows the best approximation of the Carnot cycle.

Given only a set of processes and a path-extensive measure of efficiency, our neural network-based evolutionary learning framework is able to maximize path efficiency and so deduce a classic result of physics. This learning framework is also successful if it is presented with a different set of processes. When denied the adiabatic processes of Tab. I, our gradient-free evolutionary algorithm learns the Stirling cycle [45], which is maximally-efficient in this context (with our systems parameters, $\eta_S \approx 0.291$); when denied the isothermal processes it similarly learns the maximally-efficient Otto cycle [46]. In fact, these are the only types of cycles possible in those scenarios. For the Otto cycle, T_h must be increased so we set it to 1000 K ($\eta_O \approx 0.658$ and $\eta_{\max} = 0.8$).

Extensions to unknown thermodynamic processes are straightforward, and inspection of the resulting solutions provides physical insight in an unfamiliar setting. As an illustration, we replace the standard monatomic ideal gas adiabatic process, for which $TV^{2/3}$ is constant, with a fictitious irreversible process for which

$$TV^{2/3} \propto (1 - k)^{\Delta V / (V_0 - V_1)}; \quad (10)$$

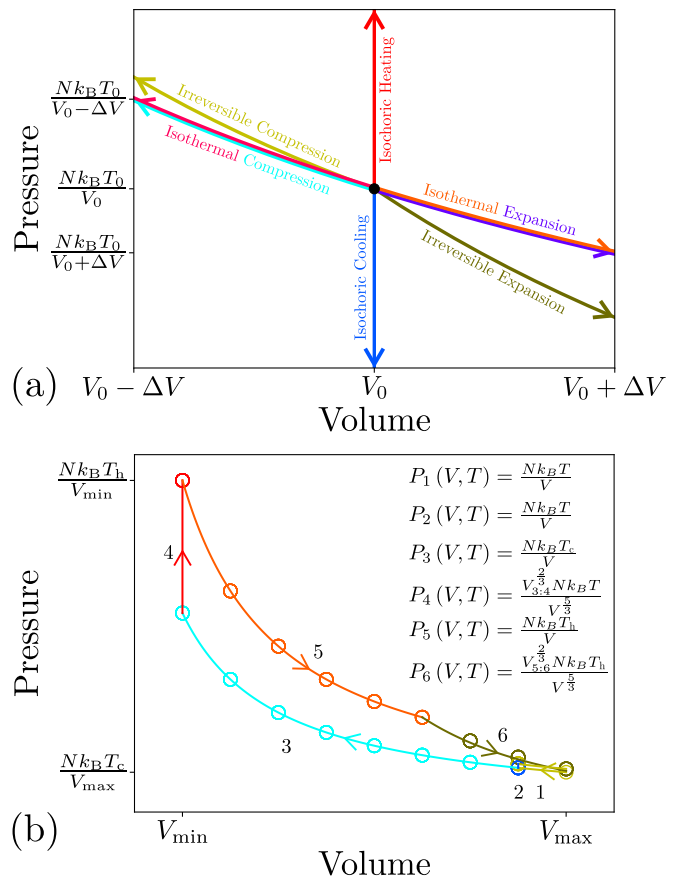


FIG. 3. We apply the evolutionary process described in Fig. 2 to a new setting in which the adiabatic processes of Tab. I are replaced with the irreversible process (10); panel (a) summarizes the new set of accessible moves. (b) Highly-evolved networks learn to enact a hybrid of the Stirling and Carnot cycles, and the resulting equations of state can be identified by curve fitting.

here $k = 2/5$ and ΔV are the fraction of thermal energy lost to the surrounding environment, and the change in volume upon making the move, respectively. We allow the network access to this process and the others of Tab. I (excluding adiabatic processes), summarized in Fig. 3(a). In this setting, we do not know in advance the most efficient trajectory. In Fig. 3(b) we show that the solution identified by our evolutionary learning scheme is a hybrid of the Stirling and Carnot cycles. Upon investigation of this solution, it is clear that the detriment to the thermal efficiency caused by the increased use of heat from the isochore (labelled process 4) is more beneficial than trying to approximate just the Carnot cycle given these irreversible processes. In this modified heat engine, we define an approximation of the Carnot cycle by replacing the adiabatic processes with the irreversible processes such that they heat and cool the system in a qualitatively similar manner. The irreversible compression process is thermally inefficient compared to adiabatic compression, therefore requiring more work to

heat the system without using the thermal reservoirs. If k is sufficiently low, the approximated Carnot cycle is most efficient and as k approaches 1.0, the hybrid cycle approaches the Stirling cycle. Note that for sufficiently high k the irreversible process becomes physically unreasonable. By fitting equations to each branch of the cycle, we identify the equations of state that result from the irreversible process (10). These results highlight the general applicability of the learning scheme and indicate the physical insight that can be obtained by interrogating solutions identified by machine learning.

V. GRADIENT-BASED REINFORCEMENT LEARNING

As an alternative to the gradient-free evolutionary learning, we consider a method which includes gradient information. We note however, that while the observation $s_t = (V, T)$ contains sufficient information to correctly determine the required thermodynamic process at any step t , as shown above, it is not sufficient for determining thermal efficiency. For example, suppose the Carnot cycle requires K steps with $s_0 = (V_{\max}, T_c)$. K steps later we arrive back at the same observation, $s_K = s_0$, achieving a thermal efficiency of η_{\max} . Now if we perform the reverse Carnot cycle, we would begin and end at the same observations as before, however this time achieving a thermal efficiency of $-\eta_{\max}$, meaning that the relation of the current observation representation to thermal efficiency is one-to-many. Unlike gradient-free algorithms, traditional gradient-based reinforcement learning algorithms operate using the rewards associated with specific steps of a trajectory, instead of assigning a score to an entire trajectory. For this reason, these algorithms cannot be used on this specific simulated heat engine because it is not Markovian. A common example of a Markovian problem would be chess, where one needs to only know the current state of the board, not how the individual pieces were moved to their current positions in order to make an optimal move.

A. Markovian model heat engine

We now introduce a Markovian simulated heat engine designed for gradient-based reinforcement learning (we will also consider this Markovian system with our gradient-free approach). In this new problem, which will be referred to as the Markovian heat engine problem, we provide the agent with non-negative work and heat budgets, defined as $W_K^* = W_0^* + \sum_{t=0}^{K-1} \Delta W_{s_t s_{t+1}}$ and $Q_K^* = Q_0^* - \sum_{t=0}^{K-1} \Delta Q_{s_t s_{t+1}}^{\text{in}}$ respectively, where W_0^* and Q_0^* are the initial budget values. For a thermodynamic process to be performed at time t , $W_{t-1}^* \geq -\Delta W_{s_{t-1} s_t}$ and $Q_{t-1}^* \geq Q_{s_{t-1} s_t}^{\text{in}}$ are required to ensure W_t^* and Q_t^* are non-negative. These budgets represent the amount of work the agent can put into the system, where any work

produced during a trajectory is added to the budget, and the amount of heat that can be transferred from the hot thermal reservoir to the heat engine. We define the observation for this new heat engine as $s_t = (V, T, W_t^*, Q_t^*)$, however we still consider a cycle as a trajectory that returns the same values of V and T . With this new representation, the observations at the beginning and end of a given singular cycle are now unique and Markovian. This means that a given observation contains all the required information about the state of the simulated heat engine and therefore the previous trajectory is not required. Full Markovity is not necessarily required to apply gradient-based reinforcement learning, however it does ensure the mapping from observation-action pairs to rewards can be learned by the agent. As W_0^* and Q_0^* are fixed, the relation of observation to thermal efficiency is now many-to-one, solving the issue of mapping an observation to a single thermal efficiency.

As the observations for the start of each repeated cycle are unique, if a trajectory produces one cycle, it does not guarantee it will produce many cycles. This poses a greater challenge as each subsequent cycle must now be learned individually. This could be solved by re-initializing the system to the initial observation after a cycle is performed, however this would force a solution for single cycles where we would like to find a more general solution for the system. Thermal efficiency could be assigned to the final observation-action pair of a trajectory to update a given policy using gradients in order to produce the most thermally efficient cycle, similar to what was done in the previous case, however this does not encourage trajectories containing more than a single cycle as thermal efficiency does not change when a cycle is repeated exactly.

To avoid explicitly penalizing a trajectory for stopping after a single cycle, we instead use $\Delta W = W_K^* - W_0^*$ at the end of a trajectory. As $\Delta Q = Q_0^* - Q_K^*$ has a fixed finite maximum, using Eq. 1, maximizing ΔW maximizes thermal efficiency while also encouraging trajectories with many cycles, satisfying our desired requirements. To help clarify this, consider the following example. Suppose some cycle on this heat engine produces 0.2 units of work and consumes 0.5 units of heat. This would give a thermal efficiency of 0.4. Now if we performed this cycle 10 times we would produce 2.0 units of work and consume 5.0 units of heat. However this would still give a thermal efficiency of 0.4. If we use thermal efficiency as the reward, these two cases are equivalent, however if we use ΔW as the reward, the second case is preferred.

B. Proximal policy optimization

With our newly defined engine and reward scheme, we turn to proximal policy optimization (PPO) [13] as the gradient-based reinforcement learning algorithm to find a policy, π_θ , that produces the desired trajectories. PPO is a commonly used gradient-based reinforcement learn-

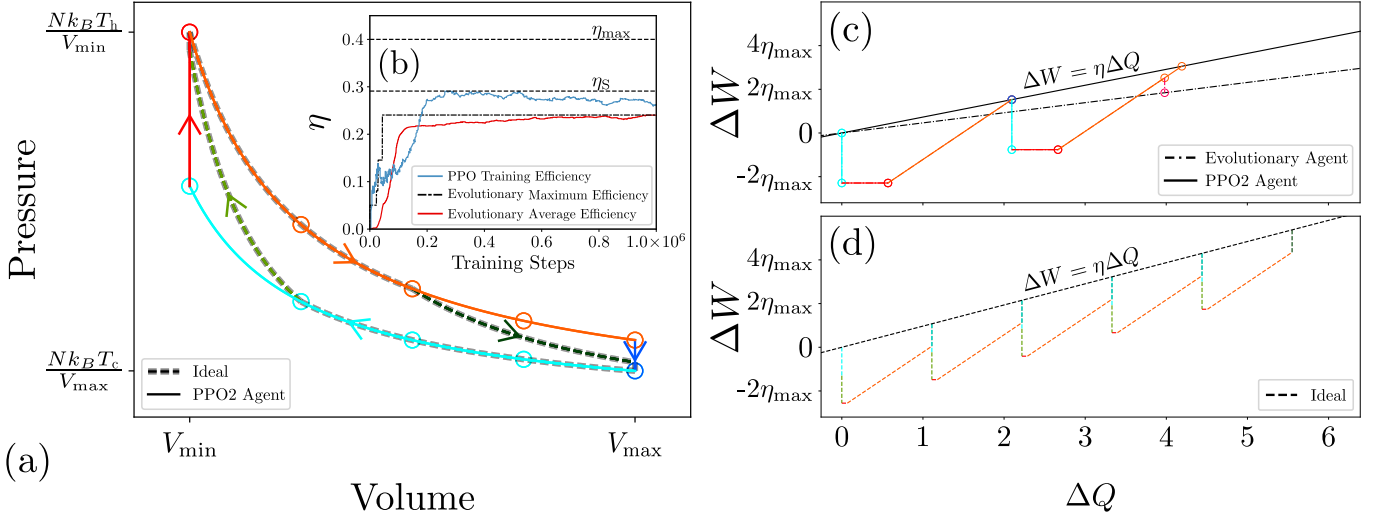


FIG. 4. We apply the PPO reinforcement learning process to our Markovian simulated heat engine; panel (a) shows the most thermally efficient trajectory produced in this process overlaid with the ideal solution, the discretized Carnot cycle. (b) The learning dynamics of thermal efficiency as a function of training steps for both the PPO and evolutionary agents as they learn to maximize ΔW . (c) ΔW as a function of ΔQ for the trajectory produced by the best performing PPO and evolutionary agents and (d) the ideal solution (discretized Carnot cycle), where $\Delta Q = 1$ maps to the amount of energy used during a single analytic Carnot cycle.

ing algorithm which has had good performance on many problems [47–51]. The true policy, π_θ^* , requires certain features from the system in order to produce these desired trajectories. While the observation may not explicitly contain these features, we assume we can extract them using a neural network approximation of π_θ^* . PPO aims to maximize the policy gradient objective function,

$$L_t^{\text{PG}}(\theta) = \hat{\mathbb{E}} \left[\log \pi_\theta(a_t | s_t) \hat{A}_t \right], \quad (11)$$

where the expectation is taken with respect to the actions a_t and $\hat{A}_t = \hat{A}_\theta(s_t | a_t)$ is the estimate of the advantage function defined as the difference between the Q -function $Q_{\pi_\theta}(s_t, a_t)$ and the value function $V_{\pi_\theta}(s_t)$. The Q -function is the expected discounted future reward if the process a_t is performed starting at s_t and π_θ is followed for the remainder of the trajectory starting at s_{t+1} , defined as

$$Q_{\pi_\theta}(s_t, a_t) = r_t + \gamma \max_a Q_{\pi_\theta}(s_{t+1}, a), \quad (12)$$

where $0 \leq \gamma \leq 1$ is the discount factor. The value function is the expected discounted future reward if π_θ is followed for the remainder of the trajectory starting at s_t , defined as

$$V_{\pi_\theta}(s_t) = \max_a Q_{\pi_\theta}(s_t, a). \quad (13)$$

Our estimate of future reward is rarely perfect, therefore it is discounted relative to the time scale of the reward. For example, when deciding on where to eat, we only care about the immediate reward we get, so we heavily discount our estimate of future reward in this case, however when investing in stocks, we care about the long-term reward, so our estimate of future reward should be

minimally discounted. Q_{π_θ} and V_{π_θ} are not known analytically but it is assumed they require the same features as π_θ and therefore we have our same neural network approximation of π_θ output the value at s_t , i.e. the output of the neural network is $(\pi_\theta(s_t), V_{\pi_\theta}(s_t))$. However, it is more efficient to instead maximize the trust region policy optimization (TRPO) [52] objective function,

$$L_t^{\text{CPI}}(\theta) = \hat{\mathbb{E}} \left[\frac{\pi_\theta(a_t | s_t)}{\pi_{\theta_{\text{old}}}(a_t | s_t)} \hat{A}_t \right] = \hat{\mathbb{E}} \left[r_{\theta, \theta_{\text{old}}}(a_t | s_t) \hat{A}_t \right], \quad (14)$$

where $\pi_{\theta_{\text{old}}}$ is the set of parameters of the previous policy. The trivial solution to this is to make extreme modifications to π_θ , therefore PPO clips the probability ratio $r_{\theta, \theta_{\text{old}}}$ and takes the minimum,

$$L_t^{\text{CLIP}}(\theta) = \hat{\mathbb{E}} \left[\min \left(r_\theta(a_t | s_t) \hat{A}_t, r_{\theta}^{\text{CLIP}}(a_t | s_t) \hat{A}_t \right) \right] r_{\theta, \theta_{\text{old}}}^{\text{CLIP}}(a_t | s_t) = \text{clip} \left(r_{\theta, \theta_{\text{old}}}(a_t | s_t), 1 - \epsilon, 1 + \epsilon \right), \quad (15)$$

where ϵ is a hyperparameter and the clip function is defined as

$$\text{clip}(x, a, b) = \min(\max(x, a), b) = \max(\min(x, b), a). \quad (16)$$

While this objective function is designed to optimize the policy, it is not designed to optimize the value function, therefore we use a mean squared error cost function,

$$L_t^{\text{VF}}(\theta) = (V_{\pi_\theta}(s_t) - V_t^{\text{targ}})^2, \quad (17)$$

where V_t^{targ} is the target value. Combing this cost function with our objective function and adding in a so called

“entropy” term S to encourage exploration, we have

$$L_t^{\text{PPO}}(\theta) = \hat{\mathbb{E}} [L_t^{\text{CLIP}}(\theta) - c_1 L_t^{\text{VF}}(\theta) + c_2 S[\pi_\theta](s_t)], \quad (18)$$

for our overall PPO objective function, where c_1 and c_2 are hyper-parameters and $S[\pi_\theta]$ is defined by

$$S[\pi_\theta](s_t) = \sum_a \pi_\theta(a_t|s_t) \log(\pi_\theta(a_t|s_t)). \quad (19)$$

Note that we have the negative of our cost term, so maximizing L_t^{PPO} simultaneously maximizes our objective function and minimizes our cost function. To update the parameters of our neural network, the gradient of L_t^{PPO} is estimated with respect to θ using backpropagation and we make changes to θ in the ascending direction of this gradient.

C. Results and Discussion

Using the identical network architecture as with the gradient-free method, we randomly initialize a policy and allow it to produce trajectories in the newly-defined thermodynamic system. As the policy interacts with this system, we collect tuples of the form (s_t, a_t, r_t, s_{t+1}) which are the experiences used to update the policy with PPO. We require all four of these because the reward r_t specifically arises from how the action a_t maps the observation s_t to the next observation s_{t+1} . r_t and s_{t+1} are also required to update our approximation of the value function and calculate \hat{A}_t . To perform these updates, we used the following hyper-parameters: discount factor of 0.99, $\epsilon = 0.2$, $c_1 = 0.5$, $c_2 = 0.01$, a batch size of 128, a total number of training steps of 2×10^6 , and a learning rate of 2.5×10^{-4} which determines the step size when updating the parameters θ .

Unlike the gradient-free approach previously discussed, PPO is unable to achieve an optimal result. However the agent was able to produce a trajectory that not only follows the Stirling cycle as shown in Fig. 4(a), and does so more than once. When looking at the learning dynamics shown in Fig. 4(b), we can see that the policy that produces the most thermally efficient trajectory occurs at $\sim 2.5 \times 10^5$ training steps, and then thermal efficiency steadily decreases in all subsequent policies. This is likely due to the Stirling cycle being a local maximum (in terms of ΔW). When a policy is at a local maximum for long enough, the value function is optimized on the trajectories found at this local maximum to the point that it becomes useless when evaluated on any trajectories found outside of the local maximum. Due to the deterministic nature of these problems, only a single trajectory is seen at this local maximum. Eventually, the entropy term used for stochastic exploration forces the policy to stray from this locally optimal trajectory, causing the policy to produce a trajectory unseen by the agent recently. As the value function is no longer accurate at this observation, the agent is unable to bring the policy back to this local

maximum causing performance to diminish. The Stirling cycle is a local maximum because it produces more work than the Carnot cycle, i.e. achieves a higher reward per cycle, however as seen when comparing Figs. 4(c) and (d), the Stirling cycle consumes a greater amount of heat than the Carnot cycle. This makes it more viable in the short term, however, due to the upper bound constraint placed on ΔQ , the Carnot cycle is still the ideal solution when tasked with maximizing ΔW .

With the maximum amount of heat allowed, the only way to increase ΔW beyond what is produced by two Stirling cycles, a trajectory would need to instead produce either three Carnot cycles and one Stirling cycle or five Carnot cycles. The differences between these trajectories and the one the agent produces are not trivial. Making any slight modifications to the found trajectory decreases ΔW . It is only when an entire Stirling cycle is replaced with multiple Carnot cycles in a trajectory that there would be any improvement. Additionally, the reward signal used to update a policy comes from the end of a trajectory, whereas due to the repetitive nature of the non-Markovian problem, we were able to select the maximal score along the entire trajectory. Now the policy is required to produce trajectories that not only follow optimal paths but also end in optimal positions, increasing the difficulty of the Markovian problem even further.

VI. EVOLUTIONARY LEARNING ON MARKOVIAN HEAT ENGINE

If ΔW is used as the scoring metric, the evolutionary learning method from before can also be applied to this variant of the simulated heat engine. Doing so yields trajectories with *lower* thermal efficiencies and less work produced compared with the gradient-based method as shown in Figs. 4(b) and (c). Similar to the gradient-based method, the gradient-free method starts its trajectory by producing a Stirling cycle, however it is unable to complete the second cycle. While the desired trajectory is identical to that of the non-Markovian problem, this task is a more challenging one because the cycle must be learned over and over again as the heat budget is consumed. This shows that the evolutionary learning method is not necessarily superior to the gradient-based method in all ways, but rather each approach is suited for different tasks. The gradient-based method was able to outperform the evolutionary learning method on the Markovian problem, however, it was the evolutionary learning method that was able to find the maximally efficient trajectory when applied to the non-Markovian problem. The evolutionary learning method was able to be applied to both tasks, however, because of the nature of gradient-based reinforcement learning, we could only apply it to the Markovian problem. The fact the gradient-based method finds this local maximum while the gradient-free does not supports the idea that




FIG. 5. T - S diagram for the (a) Carnot, (b) discretized Carnot, and (c) Stirling cycles. The red (dark gray) area (labelled Q_{out}) represents the heat removed from the system, the white area (labelled ΔW) represents the work produced by the system, and the red (dark gray) and white areas together represent the heat added to the system.

gradient-based reinforcement learning can propagate the reward assigned to specific observation–action pairs in order to produce at least a partially optimal solution. Evolutionary learning methods cannot straightforwardly make use of these specific assignments and this is likely the reason why it is outperformed by the gradient-based learning method on the Markovian problem. Although the evolutionary learning method seems to be the go-to method, being able to define a problem the way we did for the non-Markovian one is not always possible. It would be possible to make the Markovian problem easier for the gradient-based learning method, but not without providing such a high level of feedback that the problem becomes pointless to solve. For example, one could reward the gradient-based learning agent as it performs each of the correct steps in the most thermally efficient trajectory, however, this approach would only work in the case an optimal solution was known beforehand. Novelty search techniques might allow this gradient-based learning method to find the optimal trajectory, however, this would increase the computational cost of this method even further. Given that gradient-based methods are already more computationally expensive than gradient-free methods, and that the gradient-free method has already achieved the optimal result, we have chosen not to explore these options.

VII. CONCLUSIONS

Motivated by the correspondence between games and physical trajectories, we have shown that neural network-based evolutionary learning can optimize the efficiency of trajectories of classical thermodynamics. Given a set of physical processes and a path-extensive measure of efficiency, networks evolve to learn the maximally-efficient Carnot, Stirling, or Otto cycles, reproducing classic results of physics that were originally derived by application of physical insight. Given new processes, the evolutionary framework identifies solutions that when in-

terrogated provide physical insight into the problem at hand. We have also shown that with slight modifications, gradient-based reinforcement learning can optimize the efficiency of trajectories of classical thermodynamics, although not quite as effectively.

ACKNOWLEDGMENTS

CB, UY, RC, KM, and IT performed work at the National Research Council of Canada under the auspices of the AI4D Program. IT acknowledges NSERC. SW performed work at the Molecular Foundry, Lawrence Berkeley National Laboratory, supported by the Office of Science, Office of Basic Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

The simulated heat engines (or environments [15] as commonly referred to in the reinforcement learning community) used in this study can be found at <https://clean.energyscience.ca/gyms>.

Appendix A: Proof of Carnot’s theorem on infinite heat baths

Using T - S diagrams, shown in Fig. 5, the thermal efficiency of a cycle is determined by $\eta = \Delta W / (\Delta W + Q_{\text{out}})$ where ΔW is the work produced by the system and Q_{out} is the heat removed from the system. Using Fig. 5(a), the thermal efficiency of the Carnot cycle is given by Eq. (5), which is independent of the minimum and maximum entropies reached by the Carnot cycle, denoted by S_{min}^C and S_{max}^C respectively. For cycles that operate between S_{min}^C and S_{max}^C , since T_c is the minimum temperature the system can reach, η is maximized by maximizing ΔW and minimizing Q_{out} , which results in the Carnot cycle. Now consider cycles that operate outside of the entropy range S_{min}^C to S_{max}^C , such as the ones shown in Figs. 5(b) and (c). It is impossible for the system to simultaneously be

at T_h and an entropy less than S_{\min}^C , therefore the ratio of work produced and heat removed by the system between points II and III will be less than that of the Carnot cycle. Similarly, it is impossible for the system to simultaneously be at T_c and an entropy greater than S_{\max}^C , therefore the ratio of work produced and heat removed

by the system between points IV and I will be less than the Carnot cycle. Putting all these together, any cycle operating within S_{\min}^C and S_{\max}^C is at most efficient as the Carnot cycle, and any cycle operating outside of S_{\min}^C and S_{\max}^C is less efficient, therefore Carnot's theorem holds for infinitely many heat baths.

-
- [1] J. M. Roberts, M. J. Arth, and R. R. Bush, *American anthropologist* **61**, 597 (1959).
- [2] C. J. Watkins and P. Dayan, *Machine learning* **8**, 279 (1992).
- [3] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, arXiv preprint arXiv:1312.5602 (2013).
- [4] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al., *Nature* **518**, 529 (2015).
- [5] M. G. Bellemare, Y. Naddaf, J. Veness, and M. Bowling, *Journal of Artificial Intelligence Research* **47**, 253 (2013).
- [6] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, in *International conference on machine learning* (2016), pp. 1928–1937.
- [7] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. d. L. Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq, et al., arXiv preprint arXiv:1801.00690 (2018).
- [8] E. Todorov, T. Erez, and Y. Tassa, in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on* (IEEE, 2012), pp. 5026–5033.
- [9] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming* (John Wiley & Sons, 2014).
- [10] A. Asperti, D. Cortesi, and F. Sovrano, in *Machine Learning, Optimization, and Data Science*, edited by G. Nicosia, P. Pardalos, G. Giuffrida, R. Umeton, and V. Sciacca (Springer International Publishing, Cham, 2019), pp. 264–275.
- [11] M. Riedmiller, in *European Conference on Machine Learning* (Springer, 2005), pp. 317–328.
- [12] M. Riedmiller, T. Gabel, R. Hafner, and S. Lange, *Autonomous Robots* **27**, 55 (2009).
- [13] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, arXiv preprint arXiv:1707.06347 (2017).
- [14] F. P. Such, V. Madhavan, E. Conti, J. Lehman, K. O. Stanley, and J. Clune, arXiv preprint arXiv:1712.06567 (2017).
- [15] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, arXiv preprint arXiv:1606.01540 (2016).
- [16] M. Kempka, M. Wydmuch, G. Runc, J. Toczek, and W. Jaśkowski, in *Computational Intelligence and Games (CIG), 2016 IEEE Conference on* (IEEE, 2016), pp. 1–8.
- [17] M. Wydmuch, M. Kempka, and W. Jaśkowski, *IEEE Transactions on Games* **11**, 248 (2019).
- [18] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. Van Den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, et al., *nature* **529**, 484 (2016).
- [19] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al., *Nature* **550**, 354 (2017).
- [20] K. Mills, P. Ronagh, and I. Tamblyn, *Nature Machine Intelligence* **2**, 509 (2020).
- [21] P. Andreasson, J. Johansson, S. Liljestrand, and M. Granath, *Quantum* **3**, 183 (2019).
- [22] Z. Zhou, S. Kearnes, L. Li, R. N. Zare, and P. Riley, *Scientific reports* **9**, 1 (2019).
- [23] M. I. Radaideh, I. Wolverson, J. Joseph, J. J. Tusar, U. Otgonbaatar, N. Roy, B. Forget, and K. Shirvan, *Nuclear Engineering and Design* p. 110966 (2020).
- [24] T. Badloe, I. Kim, and J. Rho, *Physical Chemistry Chemical Physics* **22**, 2337 (2020).
- [25] M. Popova, O. Isayev, and A. Tropsha, *Science advances* **4**, eaap7885 (2018).
- [26] J. J. De Yoreo, P. U. Gilbert, N. A. Sommerdijk, R. L. Penn, S. Whitlam, D. Joester, H. Zhang, J. D. Rimer, A. Navrotsky, J. F. Banfield, et al., *Science* **349**, aaa6760 (2015).
- [27] M. F. Hagan and D. Chandler, *Biophysical Journal* **91**, 42 (2006).
- [28] A. W. Wilber, J. P. Doye, A. A. Louis, E. G. Noya, M. A. Miller, and P. Wong, *The Journal of Chemical Physics* **127**, 085106 (2007).
- [29] S. Whitlam and R. L. Jack, *Annual review of physical chemistry* **66**, 143 (2015).
- [30] D. T. Gillespie, *Annu. Rev. Phys. Chem.* **58**, 35 (2007).
- [31] T. McGrath, N. S. Jones, P. R. ten Wolde, and T. E. Ouldridge, *Physical Review Letters* **118**, 028101 (2017).
- [32] U. Seifert, *Reports on progress in Physics* **75**, 126001 (2012).
- [33] A. I. Brown and D. A. Sivak, *Proceedings of the National Academy of Sciences* **114**, 11057 (2017).
- [34] U. Seifert, *Physical Review Letters* **95**, 040602 (2005).
- [35] V. Lecomte, C. Appert-Rolland, and F. Van Wijland, *Journal of statistical physics* **127**, 51 (2007).
- [36] F. Ritort, *Advances in Chemical Physics* **137**, 31 (2008).
- [37] J. P. Garrahan, R. L. Jack, V. Lecomte, E. Pitard, K. van Duijvendijk, and F. van Wijland, *Journal of Physics A: Mathematical and Theoretical* **42**, 075007 (2009).
- [38] T. Speck, A. Engel, and U. Seifert, *Journal of Statistical Mechanics: Theory and Experiment* **2012**, P12001 (2012).
- [39] V. Lecomte, A. Imparato, and F. v. Wijland, *Progress of Theoretical Physics Supplement* **184**, 276 (2010).
- [40] R. J. Harris, *Journal of Statistical Mechanics: Theory and Experiment* **2015**, P07021 (2015).
- [41] H. B. Callen, *Thermodynamics and an introduction to thermostatistics* (John Wiley & Sons, New York, 1985), 2nd ed.
- [42] S. Carnot, *Reflections on the motive power of fire* by Sadi Carnot and other papers on the Second Law of Thermo-

dynamics by E. Clapeyron and R (1824).

- [43] M. S. Silberberg, *Principles of general chemistry* (McGraw-Hill Higher Education New York, 2007).
- [44] R. S. Sutton, A. G. Barto, F. Bach, et al., *Reinforcement learning: An introduction* (MIT press, 1998).
- [45] T. Finkelstein, Tech. Rep., AIAA-94-3951-CP (1829).
- [46] M. Mozurkewich and R. S. Berry, *Journal of Applied Physics* **53**, 34 (1982).
- [47] E. Bøhn, E. M. Coates, S. Moe, and T. A. Johansen, in *2019 International Conference on Unmanned Aircraft Systems (ICUAS)* (IEEE, 2019), pp. 523–533.
- [48] L. C. Melo and M. R. O. A. Máximo, in *2019 Latin american robotics symposium (LARS), 2019 Brazilian symposium on robotics (SBR) and 2019 workshop on robotics in education (WRE)* (IEEE, 2019), pp. 37–42.
- [49] H. Wei, X. Liu, L. Mashayekhy, and K. Decker, in *2019 IEEE Vehicular Networking Conference (VNC)* (IEEE, 2019), pp. 1–8.
- [50] L. Zhang, Y. Zhang, X. Zhao, and Z. Zou, *Image and Vision Computing* **108**, 104126 (2021).
- [51] L. De Vree and R. Carloni, *IEEE Transactions on Neural Systems and Rehabilitation Engineering* **29**, 607 (2021).
- [52] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, in *International conference on machine learning* (2015), pp. 1889–1897.