

NRC Publications Archive Archives des publications du CNRC

Finding the ground state of spin Hamiltonians with reinforcement learning

Mills, Kyle; Ronagh, Pooya; Tamblyn, Isaac

This publication could be one of several versions: author's original, accepted manuscript or the publisher's version. / La version de cette publication peut être l'une des suivantes : la version prépublication de l'auteur, la version acceptée du manuscrit ou la version de l'éditeur.

For the publisher's version, please access the DOI link below. / Pour consulter la version de l'éditeur, utilisez le lien DOI ci-dessous.

Publisher's version / Version de l'éditeur:

<https://doi.org/10.1038/s42256-020-0226-x>

Nature Machine Intelligence, 2, 9, pp. 509-517, 2020-09-07

NRC Publications Archive Record / Notice des Archives des publications du CNRC :

<https://nrc-publications.canada.ca/eng/view/object/?id=885266fe-78e3-4d42-ae68-723d2015f0ca>

<https://publications-cnrc.canada.ca/fra/voir/objet/?id=885266fe-78e3-4d42-ae68-723d2015f0ca>

Access and use of this website and the material on it are subject to the Terms and Conditions set forth at

<https://nrc-publications.canada.ca/eng/copyright>

READ THESE TERMS AND CONDITIONS CAREFULLY BEFORE USING THIS WEBSITE.

L'accès à ce site Web et l'utilisation de son contenu sont assujettis aux conditions présentées dans le site

<https://publications-cnrc.canada.ca/fra/droits>

LISEZ CES CONDITIONS ATTENTIVEMENT AVANT D'UTILISER CE SITE WEB.

Questions? Contact the NRC Publications Archive team at

PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca. If you wish to email the authors directly, please see the first page of the publication for their contact information.

Vous avez des questions? Nous pouvons vous aider. Pour communiquer directement avec un auteur, consultez la première page de la revue dans laquelle son article a été publié afin de trouver ses coordonnées. Si vous n'arrivez pas à les repérer, communiquez avec nous à PublicationsArchive-ArchivesPublications@nrc-cnrc.gc.ca.

Controlled Online Optimization Learning (COOL): Finding the ground state of spin Hamiltonians with reinforcement learning

Kyle Mills*

1QB Information Technologies (1QBit), Vancouver, British Columbia, Canada
University of Ontario Institute of Technology, Oshawa, Ontario, Canada
Vector Institute for Artificial Intelligence, Toronto, Ontario, Canada

Pooya Ronagh†

1QB Information Technologies (1QBit), Vancouver, British Columbia, Canada
Institute for Quantum Computing (IQC), Waterloo, Ontario, Canada
Department of Physics and Astronomy, University of Waterloo, Ontario, Canada

Isaac Tamblyn‡

National Research Council Canada, Ottawa, Ontario, Canada
University of Ottawa, Ottawa, Ontario, Canada
Vector Institute for Artificial Intelligence,
Toronto, Ontario, Canada

(Dated: March 3, 2020)

Reinforcement learning (RL) has become a proven method for optimizing a procedure for which success has been defined, but the specific actions needed to achieve it have not. We apply the so-called “black box” method of RL to what has been referred as the “black art” of simulated annealing (SA), demonstrating that an RL agent based on proximal policy optimization can, through experience alone, arrive at a temperature schedule that surpasses the performance of standard heuristic temperature schedules for two classes of Hamiltonians. When the system is initialized at a cool temperature, the RL agent learns to heat the system to “melt” it, and then slowly cool it in an effort to anneal to the ground state; if the system is initialized at a high temperature, the algorithm immediately cools the system. We investigate the performance of our RL-driven SA agent in generalizing to all Hamiltonians of a specific class; when trained on random Hamiltonians of nearest-neighbour spin glasses, the RL agent is able to control the SA process for other Hamiltonians, reaching the ground state with a higher probability than a simple linear annealing schedule. Furthermore, the scaling performance (with respect to system size) of the RL approach is far more favourable, achieving a performance improvement of one order of magnitude on $L = 14^2$ systems. We demonstrate the robustness of the RL approach when the system operates in a “destructive observation” mode, an allusion to a quantum system where measurements destroy the state of the system. The success of the RL agent could have far-reaching impact, from classical optimization, to quantum annealing, to the simulation of physical systems.

I. INTRODUCTION

In metallurgy and materials science, the process of annealing is used to equilibrate the positions of atoms to obtain perfect low-energy crystals. Heat provides the energy necessary to break atomic bonds, and high-stress interfaces are eliminated by the migration of defects. By slowly cooling the metal to room temperature, the metal atoms become energetically locked in a lattice structure more favourable than the original structure. Metallurgists can tune the temperature schedule to arrive at final products that have desired characteristics, such as ductility and hardness. Annealing is a biased stochastic search for the ground state.

An analogous *in silico* technique, simulated annealing (SA) [1], can be used to find the ground state of spin-

glass models, an NP-hard problem [2]. A spin glass is a graphical model consisting of binary spins S_i . The connections between spins are defined by the coupling constants J_{ij} , and a linear term with coefficients h_i can apply a bias to individual spins. The Hamiltonian

$$\mathcal{H} = - \sum_{i \neq j} J_{ij} S_i S_j - \sum_i h_i S_i, \quad S_i = \pm 1$$

defines the energy of the microstates [3]. The choices of the quadratic coupling coefficients J_{ij} and the linear bias coefficients h_i effect the interesting dynamics of the model: J_{ij} can be randomly distributed according to a Gaussian distribution [3], encompass all i, j combinations for a fully connected Hamiltonian, or be limited to short-range (e.g., nearest-neighbour, $\langle i, j \rangle$) interactions, to name a few. For example, when the positive, unit-magnitude coupling is limited to nearest-neighbour pairs, the ubiquitous ferromagnetic Ising model [4] is recovered. Examples of the Hamiltonians we investigate in this work are presented in Figure 1 and discussed in further detail in Section VI.

* kyle.mills@1qbit.com

† pooya.ronagh@1qbit.com

‡ isaac.tamblyn@nrc.ca

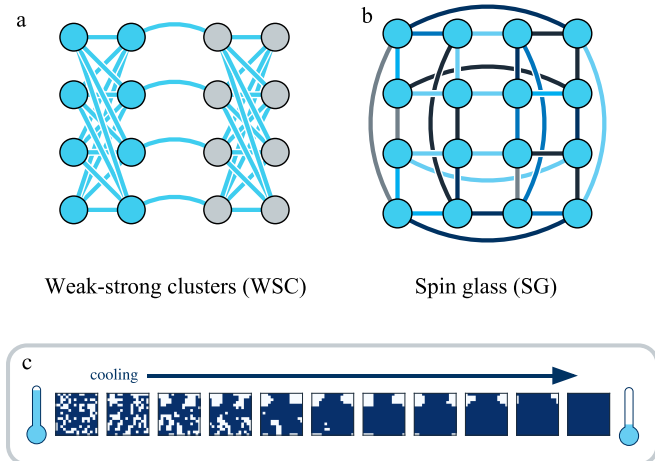


FIG. 1. Two classes of Hamiltonian problems are depicted. (a) The weak-strong clusters (WSC) model comprises two bipartite clusters. The left cluster is biased upward; the right cluster is biased downward. All couplings are equal and of unit magnitude. The two clusters are coupled via the eight central nodes. This model exhibits a deep local minimum very close in energy to the model’s global minimum. When initialized in the local minimum, the RL agent is able to learn schemes to escape the local minimum and arrive at the global minimum, without any explicit knowledge of the Hamiltonian. (b) Here we present an example spin-glass model. The nodes are coupled to nearest neighbours with random Gaussian-distributed coupling coefficients. The nodes are unbiased, and the couplings are changed at each instantiation of the model. The RL algorithm is able to learn a dynamic temperature schedule by observing the system throughout the annealing process, without explicit knowledge of the form of the Hamiltonian, and the learned policy can be applied to all instances of randomly generated couplings. We demonstrate this on variably sized spin glasses and investigate the scaling with respect to a classic linear SA schedule. In (c), we show snapshots of a sample progression of a configuration undergoing SA under the ferromagnetic Ising model Hamiltonian and a constant cooling schedule. The terminal state, all spins-up, is the ground state; this anneal would be considered successful.

Finding the ground state of (i.e., “solving”) such systems is interesting from the perspective of thermodynamics, as one can observe phenomena such as phase transitions [5, 6], but also practically useful as discrete optimization problems can be mapped to spin-glass models (e.g., the travelling salesperson problem or the knapsack problem) [7]. The Metropolis–Hastings algorithm [8, 9] can be used to simulate the spin glass at arbitrary temperature; thus, it is used ubiquitously for SA. By beginning the simulation at a high temperature, one can slowly cool the system over time, providing sufficient thermal energy to escape local minima, and arrive at the ground state “solution” to the problem. The challenge is to find a temperature schedule that minimizes computational effort while still arriving at a satisfactory solution; if the temperature is reduced too rapidly, the sys-

tem will become trapped in a local minimum, and reducing the temperature too slowly results in an unnecessary computational expense. Kirkpatrick et al. [1, 10] proposed starting at a temperature that results in an 80% acceptance ratio (i.e., 80% of Metropolis spin flips are accepted) and reducing the temperature geometrically. They also recommended monitoring the objective function and reducing the cooling rate if the objective value (e.g., the energy) drops too quickly. More-sophisticated adaptive temperature schedules have been investigated [11]. Nevertheless, in his 1987 paper, Bounds [12] said that “choosing an annealing schedule for practical purposes is still something of a black art”.

When framed in the advent of quantum computation and quantum control, establishing robust and dynamic scheduling of control parameters becomes even more relevant. For example, the same optimization problems that can be cast as classical spin glasses are also amenable to quantum annealing [13–17], exploiting, in lieu of thermal fluctuations, the phenomenon of quantum tunnelling [18–20] to escape local minima. Quantum annealing (QA) was proposed by Finnila et al. [21] and Kadowaki and Nishimori [22], and, in recent years, physical realizations of devices capable of performing QA (quantum annealers), have been developed [23–26], and are being rapidly commercialized. As these technologies progress and become more commercially viable, practical applications [17, 27] will continue to be identified and resource scarcity will spur the already extant discussion of the efficient use of annealing hardware [28, 29].

Nonetheless, there are still instances where the classical (SA) outperforms the quantum (QA) [30], and improving the former should not be undervalued. *In silico* and hardware annealing solutions such as Fujitsu’s FPGA-based Digital Annealer [31], NTT’s laser-pumped coherent Ising machine (CIM) [32], and the quantum circuit model algorithm known as QAOA [33, 34] all demand the scheduling of control parameters, whether it is the temperature in the case of the Digital Annealer, or the power of the laser pump in the case of CIM. Heuristic methods based on trial-and-error experiments are commonly used to schedule these control parameters, and an automatic approach could expedite development, and improve the stability of such techniques.

In this work, we demonstrate the use of a reinforcement learning (RL) method to learn the “black art” of classic SA temperature scheduling, and show that an RL agent is able to learn dynamic control parameter schedules for various problem Hamiltonians. The schedules that the RL agent produces are dynamic and reactive, adjusting to the current observations of the system to reach the ground state quickly and consistently without *a priori* knowledge of a given Hamiltonian. Our technique, aside from being directly useful for *in silico* simulation, is an important milestone for future work in quantum information processing, including for hardware- and software-based control problems.

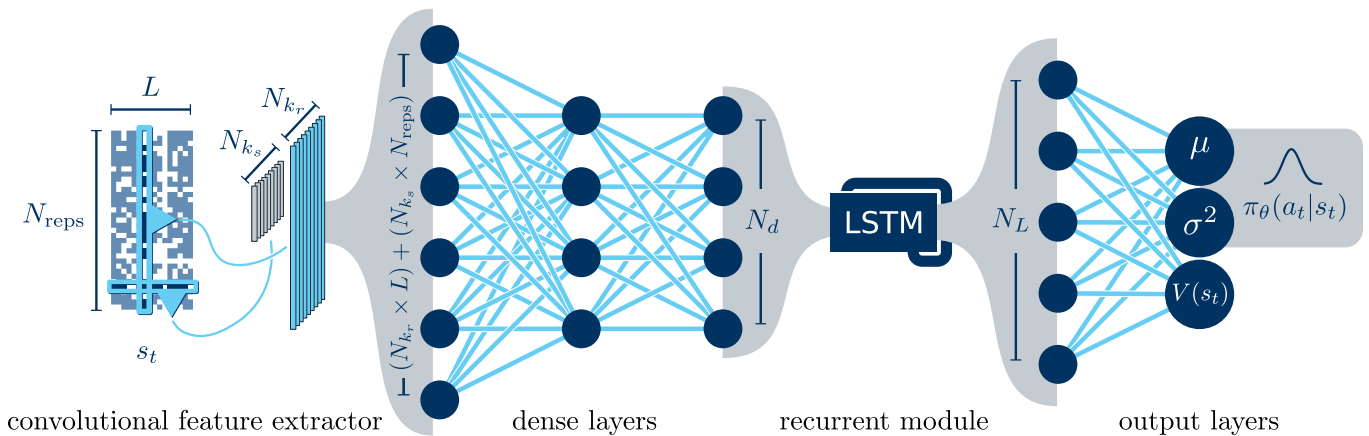


FIG. 2. A neural network is used to learn the control parameters for several SA experiments. By observing a lattice of spins, the neural network can learn to control the temperature of the system in a dynamic fashion, annealing the system to the ground state. The spins at time t form the state s_t fed into the network. Two concurrent convolutional layers extract features from the state. These features are combined with a dense layer and fed into a recurrent module (an LSTM module) capable of capturing temporal characteristics. The LSTM module output is reduced to two parameters used to form the policy distribution $\pi_{\theta}(a_t | s_t)$ as well as to approximate the value function $V(s_t)$ used for the generalized advantage estimate.

II. REINFORCEMENT LEARNING

Reinforcement learning is a branch of dynamic programming whereby an agent, residing in state s_t at time t , learns to take an action a_t that maximizes a cumulative reward signal R by dynamically interacting with an environment [35]. Through the training process, the agent arrives at a policy π that depends on some observation (or “state”) of the system, s . In recent years, neural networks have taken over as the *de facto* function approximator for the policy. Deep reinforcement learning has seen unprecedented success, achieving superhuman performance in a variety of video games [36–39], board games [40–42], and other puzzles [43, 44]. While many reinforcement learning algorithms exist, we have chosen to use proximal policy optimization (PPO) [45], implemented within Stable Baselines [46] for its competitive performance on problems with continuous action spaces.

III. THE ENVIRONMENT

We developed an OpenAI gym [47] environment which serves as the interface to the “game” of simulated annealing. Let us now define some terminology and parameters important to simulated annealing. For a given Hamiltonian, defining the interactions of L spins, we create N_{reps} randomly initialized replicas (unless otherwise specified). The initial spins of each replica are drawn from a Bernoulli distribution with probability of a spin-up being randomly drawn from a uniform distribution. These independent replicas are annealed in parallel. The replicas follow an identical temperature schedule with their uncoupled nature providing a mechanism for statistics of the system to be represented through an ensemble

of measurements. In the context of the Metropolis–Hastings framework, we define one “sweep” to be L random spin flips (per replica), and one “step” to be N_{sweeps} . After every step, the environment returns an observation of the current state s_t of the system, an $N_{\text{reps}} \times L$ array consisting of the binary spin values present. This observation can be used to make an informed decision of the action a_t that should be taken. The action, a single scalar value, corresponds to the total inverse temperature change $\Delta\beta$ that should be carried out over the subsequent step. The choice of action is provided to the environment, and the process repeats until N_{steps} steps have been taken, comprising one full anneal, or “episode” in the language of RL. If the chosen action would result in the temperature becoming negative, no change is made to the temperature and the system continues to evolve under the previous temperature.

A. Observations

For the classical version of the problem, an observation consists of the explicit spins of an ensemble of replicas. In the case of an unknown Hamiltonian, the ensemble measurement is important as the instantaneous state of a single replica does not provide sufficient information about the current temperature of the system. Providing the agent with multiple replicas allows it to compute statistics and have the possibility of inferring the temperature. For example, if there is considerable variation among replicas, then the system is likely hot, whereas if most replicas look the same, the system is probably cool.

When discussing a quantum system, where the spins represent qubits, direct mid-anneal measurement of the system is not possible as measurement causes a collapse of the wavefunction. To address this, we discuss experi-

ments conducted in a “destructive observation” environment, where measurement of the spins is treated as a “one-time” opportunity for inclusion in RL training data. The subsequent observation is then based on a different set of replicas that have evolved through the same schedule, but from different initializations.

IV. REINFORCEMENT LEARNING ARCHITECTURE

Through the framework of reinforcement learning, we wish to produce a policy function $\pi_\theta(a_t | s_t)$ that takes the observed binary spin state $s_t \in \{-1, 1\}^{N_{\text{reps}} \times L}$ and produces an action a_t corresponding to the optimal change in the inverse temperature. Here π is a distribution represented by a neural network and the subscript θ denotes parameterization by learnable weights θ . We define the function

$$\phi_k(s_t) \in \{-1, 1\}^{1 \times L}$$

as an indexing function that returns the binary spin values for the k -th rep of state s_t .

The neural network is composed of two parts: a convolutional feature extractor, and a recurrent network to capture the temporal characteristics of the problem. The feature extractor comprises two parallel two-dimensional convolutional layers. The first convolutional layer has N_{k_r} kernels of size $1 \times L$, and aggregates along the replicas dimension, enabling the collection of spin-wise statistics across the replicas. The second convolutional layer has N_{k_s} kernels of size $N_{\text{reps}} \times 1$ and slides along the spin dimension, enabling the aggregation of replica-wise statistics across the spins. The outputs of these layers are flattened, concatenated, and fed into a dense layer of size N_d hidden nodes. This operates as a latent space encoding for input to a recurrent neural network (a long short-term memory, or LSTM, module [48]), used to capture the sequential nature of our application. The latent output of the LSTM module is of size N_L . For simplicity, we set $N_{k_r} = N_{k_s} = N_d = N_L = 64$. All activation functions are hyperbolic tangent (tanh) activations. Since a_t can assume a continuum of real values, this task is referred to as having a continuous action space, and thus standard practice is for the network to output two values corresponding to the first and second moments of a normal distribution. During training, when exploration is desired, an entropic regularization in the PPO cost function can be used to encourage a high variance (i.e., encouraging σ^2 to remain sizable). Additionally, PPO requires an estimate of the generalized advantage function [49], the difference between the reward received by taking action a_t while in state s_t , and the expected value of the cumulative future reward prior to an action being taken. The latter, termed the “value function”, or $V(s_t)$, cannot possibly be computed because we know only the reward from the action that was chosen, and nothing

about the actions that were not chosen, but we can estimate it using a third output from our neural network. Thus, as seen in Figure 2, the neural network in this work takes the state s_t and maps it to three scalar quantities, μ , σ^2 , and $V(s_t)$, defining the two moments of a normal distribution and an estimate of the value function, respectively. At the core of RL is the concept of reward engineering, that is, developing a reward scheme to inject a notion of success into the system. As we only care about reaching the ground state by the end of a given episode, we use a sparse reward scheme, with a reward of zero for every time step before the terminal step, and a reward equal to the negative of the minimum energy as the reward for the terminal step, that is,

$$R_t = \begin{cases} 0, & t < N_{\text{steps}} \\ -\min_k \mathcal{H}(\phi_k(s_t)), & t = N_{\text{steps}} \end{cases}, \quad (1)$$

where $k \in [1, N_{\text{reps}}]$. With this reward scheme, we encourage the agent to arrive at the lowest possible energy by the time the episode terminates, without regard to what it does in the interim. In searching for the ground state, the end justifies the means.

When optimizing the neural network, we use a PPO discount factor of $\gamma = 0.99$, eight episodes between weight updates, a value function coefficient of $c_1 = 0.5$, an entropy coefficient of $c_2 = 0.001$, a clip range of $\epsilon = 0.05$, a learning rate of $\alpha = 1 \times 10^{-6}$, and a single minibatch per update. Each agent is trained over the course of 25,000 episodes (anneals), with $N_{\text{steps}} = 40$ steps per episode, and with $N_{\text{sweeps}} = 100$ sweeps separating each observation. We used $N_{\text{reps}} = 64$ replicas for each observation.

V. EVALUATION

Whereas the RL policy can be made deterministic, meaning a given state always produces the same action, the underlying Metropolis algorithm is stochastic; thus, we must statistically define the metric for success. We borrow this evaluation scheme from Aramon et al. [50]. Each RL episode will either result in “success” or “failure”. Let us define the “time to solution” as

$$T_s = \tau n_{99}, \quad (2)$$

that is, the number of episodes that must be run to be 99% sure the ground state has been observed at least one time (n_{99}), multiplied by the time τ taken for one episode.

Let us also define X_i as the binary outcome of the i -th episode, with $X_i = 1$ (0) if at least one (none) of the N_{reps} replicas are observed to be in the ground state at episode termination. The quantity $Y \equiv \sum_{i=1}^n X_i$ is the number of successful episodes after a total of n episodes, and $p \equiv P(X_i = 1)$ denotes the probability that an anneal i will be successful. Thus the probability of exactly k out

of n episodes succeeding is given by the probability mass function of the binomial distribution

$$P(Y = k | n, p) = \binom{n}{k} p^k (1-p)^{n-k}. \quad (3)$$

To compute the time to solution, our quantity of interest is the number of episodes n_{99} where $P = 0.99$, that is,

$$P(Y \geq 1 | n_{99}, p) = 0.99.$$

From this and (3), it can be shown that

$$n_{99} = \frac{\log(1 - 0.99)}{\log(1 - p)}.$$

In the work of Aramon et al. [50], p is estimated using Bayesian inference due to their large system sizes sometimes resulting in zero successes, precluding the direct calculation of p . In our case, to evaluate a policy, we perform 100 runs for each of 100 instances and compute p directly from the ratio of successful to total episodes, that is, $p = \bar{X}$.

VI. HAMILTONIANS

We present an analysis of two classes of Hamiltonians. The first, which we call the weak-strong clusters model (WSC; see Figure 1a), is an $L = 16$ bipartite graph with two fully connected clusters, inspired by the ‘‘Chimera’’ structure used in D-Wave Systems’ quantum annealing hardware [51]. In our case, one cluster is negatively biased with $h_i = -0.44$ and the other positively biased with $h_i = 1.0$. All couplings are ferromagnetic and have unit magnitude. This results in an energy landscape with a deep local minimum where both clusters are aligned to their respective biases, but a slightly lower global minimum when the two clusters are aligned together, sacrificing the benefit of bias-alignment for the satisfaction of the intercluster couplings. For all WSC runs, the spins of the lattice are initialized in the local minimum.

The second class of Hamiltonians are nearest-neighbour square spin glasses (SG; see Figure 1b). Couplings are periodic (i.e., the model is defined on a torus), and drawn from a normal distribution with standard deviation 1.0. All biases are zero. Hamiltonian instances are generated as needed during training. To evaluate our method and compare against the simulated annealing standard, we must have a testing set of instances for which we know the true ground state. For each lattice size investigated ($\sqrt{L} = [4, 6, 8, 10, 12, 14, 16]$) we generate $N_{\text{test}} = 100$ unique instances and obtain the true ground state energy for each instance using the branch-and-cut method [52] through the Spin Glass Server [53].

VII. RESULTS

A. Weak-strong clusters model

We demonstrate the use of RL on the WSC model shown in Figure 1a. RL is able to learn a simple temperature schedule that anneals the WSC model to the ground state in 100% of episodes, regardless of the temperature in which the system is initialized. In Figure 3b, we compare the RL policy to standard simulated annealing (i.e., Metropolis–Hastings) with several linear inverse temperature schedules (i.e., constant cooling rates).

When the system is initially hot (small β), both RL and SA are capable of reaching the ground state with 100% success as there exists sufficient thermal energy to escape the local minimum. In Figure 3c, we plot an example schedule. The RL policy (red) increases the temperature slightly at first, but then begins to cool the system almost immediately. An abrupt decrease in the Metropolis acceptance rate is observed (Figure 3e). The blue dashed line in Figure 3c represents the average schedule of the RL policy over 100 independent anneals. The standard deviation is shaded. It is apparent that the schedule is quite consistent between runs at a given starting temperature, with some slight variation in the rate of cooling.

When the system is initially cold (large β), there exists insufficient thermal energy to overcome the energy barrier between the local and global minima, and SA fails with a constant cooling rate. The RL policy, however, is able to identify, through observation, that the temperature is too low and can rapidly decrease β initially, heating the system to provide sufficient thermal energy to avoid the local minimum. In Figure 3f, we see an increase in the Metropolis acceptance ratio, followed by a decrease, qualitatively consistent with the human-devised heuristic schedules that have been traditionally suggested [1, 10, 11]. In Figure 3d, we plot an example schedule. Initially, the RL algorithm increases the temperature to provide thermal energy to escape the minimum, then begins the process of cooling. Similar to Figure 3c, the broadness of the variance of the policies is greatest in the cooling phase, with some instances being cooled more rapidly than others. The RL agent does not have access to the current temperature directly, and bases its policy solely on the spins. The orthogonal unit-width convolutions provide a mechanism for statistics over spins and replicas, and the LSTM module provides a mechanism to capture the time-dependent dynamics of the system.

B. Spin-glass model

We now investigate the performance of the RL algorithm in learning a general policy for an entire class of Hamiltonians, investigating whether the RL algorithm can learn to generalize its learning to accommodate a theoretically infinite set of Hamiltonians of a specific class.

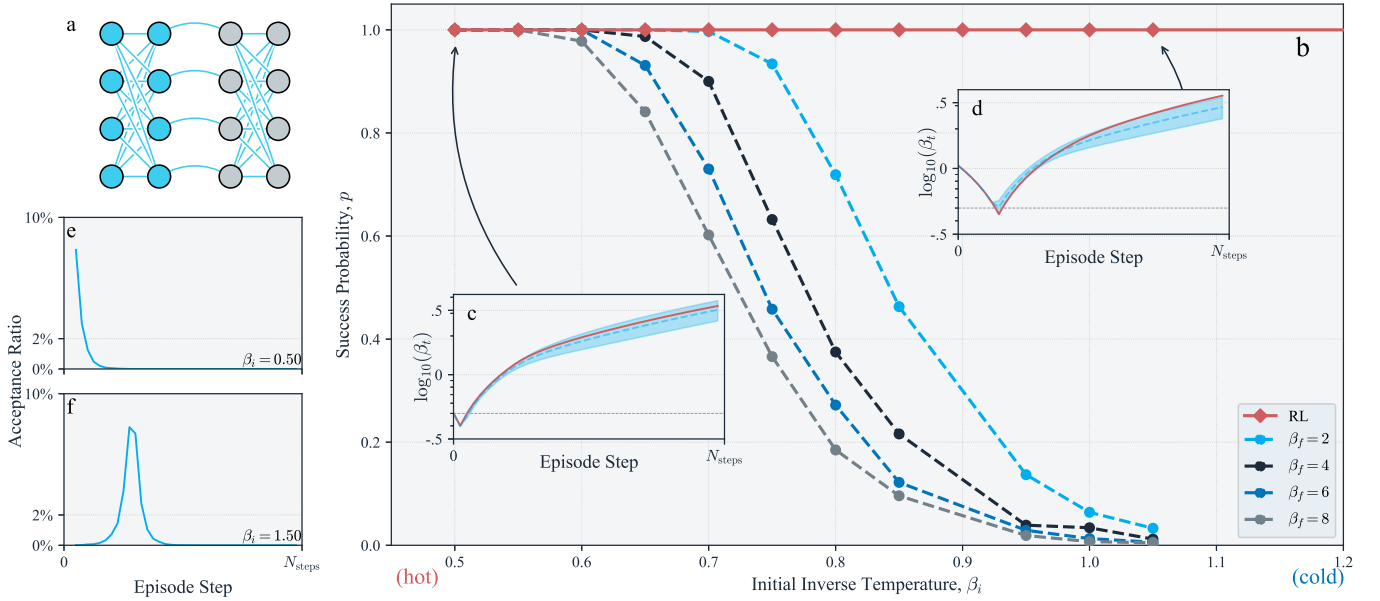


FIG. 3. An RL policy learns to anneal the WSC model (shown in (a)). (b) We plot the performance of various linear (with respect to β) simulated annealing schedules, cooling from β_i to β_f , as well as the performance of the RL policy for a variety of starting temperatures. When the initial inverse temperature is sufficiently small, both the RL and SA algorithms achieve 100% success (i.e., every episode reaches the ground state). When the system is initialized with a large β_i , there is insufficient thermal energy for SA to overcome the energy barrier and reach the ground state, and consequently a very low success probability. A single RL policy achieves almost perfect success across all initial temperatures. In (c) and (d), we plot the RL inverse temperature schedule in red for episodes initialized with respective low and high inverse temperatures. In blue, we show the average RL policy for the specific starting temperature. The RL algorithm can identify a cold initialization from observation, and increases the temperature before then decreasing it (as shown in (d)). In (e) and (f), we plot the Metropolis acceptance ratio for two episodes, initialized at two extreme temperatures (e) low β_i , and (f) high β_i .

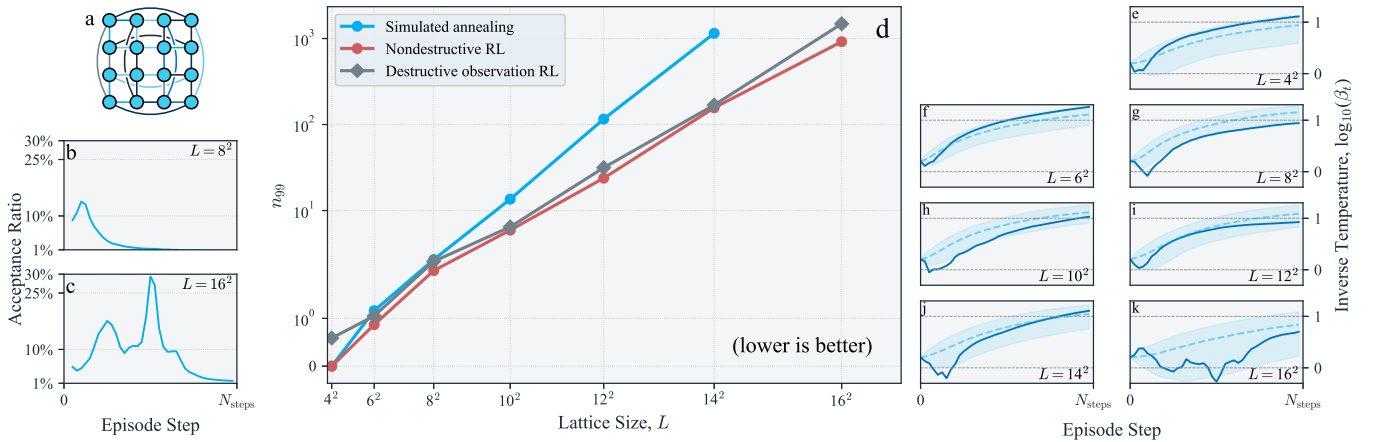


FIG. 4. An RL policy learns to anneal spin-glass models. An example ($L = 4^2$) lattice is shown in (a). In (b) and (c), we plot the acceptance ratios over time for three episodes for each of the $L = 8^2$ and $L = 16^2$ lattices. In (d), we compare the scaling of the RL policy with respect to system size and compare it to a simple linear simulated annealing schedule. We plot the n_{99} value (the number of anneals required to be 99% certain of observing the ground state) as a function of system size for both the RL and the best linear simulated annealing schedule we observed. For all system sizes investigated, the learned RL policy is able to reach the ground state in significantly fewer runs. Additionally, we plot the destructive observation results, which in most cases, still rival or beat the linear schedules. We note that the destructive observation requires far more Monte Carlo steps per episode to simulate the destructive measurements; this plot should not be interpreted as a comparison of run time with regard to the destructive observation result. In (e) through (k), we plot an example inverse temperature schedule as a solid line, as well as the average inverse temperature schedule (for all testing episodes) as a dashed line, with the shaded region denoting the standard deviation.

Furthermore, we investigate how RL performs with various lattice sizes, and compare the trained RL model to a linear (with respect to β) classic simulated annealing schedule. The results of this investigation are shown in Figure 4.

In all cases, the RL schedule obtains a better (lower) n_{99} value, meaning far fewer episodes are required for us to be confident that the ground state has been observed. Furthermore, the n_{99} value exhibits much better scaling with respect to the system size (i.e., the number of optimization variables). In Figure 4e-k, we plot some of the schedules that the RL algorithm produces. In many cases, we see initial heating, followed by cooling, although in the case of the $L = 16^2$ model (Figure 4k) we see much more complex, but still successful, behaviour. In all cases, the variance of the policies with respect to time (shown as the shaded regions in Figure 4e-k), indicate the agent is using information from the provided state to make decisions, and not just basing its decisions on the elapsed time using the internal state of the LSTM module. If schedules were based purely on some internal representation of time, there would be no variance between episodes.

C. Destructive observation

A key element of the nature of quantum systems is the collapse of the wavefunction when a measurement of the quantum state is made. When dealing with quantum systems, one must make control decisions based on quantum states that have evolved through an identical policy but have never before been measured. We model this restriction on quantum measurements by allowing any replica observed in the anneal to be consumed as training data for the RL algorithm only once. We simulate this behaviour by keeping track of the policy decisions (the changes in inverse temperature) in an action buffer as we play through each episode. When a set of N_{reps} replicas are measured, they are consumed and the system is reinitialized. The actions held in the buffer are replayed on the new replicas.

In this situation, the agent cannot base its decision on any replica-specific temporal correlations between given measurements; this should not be a problem early in each episode, as the correlation time scale of a hot system is very short, and the system, even under nondestructive observation, would have evolved sufficiently, in the time window between steps, to be uncorrelated. However, as the system cools, the correlation time scale increases exponentially, and destructive observation prevents the agent from relying on temporal correlations of any given replica.

We evaluate an agent trained in this “quantum-inspired” way and plot its performance alongside the non-destructive (i.e., classical) case in Figure 4d. In the case of destructive observation, the agent performs marginally less well than the nondestructive case, but still performs

better than SA in most cases. As it is a more complicated task to make observations when the system is temporally uncorrelated, it is understandable that the performance would be inferior to the nondestructive case. Nonetheless, RL is capable of outperforming SA in both the destructive and nondestructive cases.

The relative performance in terms of computational demand between destructive observation and SA alludes to an important future direction in the field of RL, especially when applied to physical systems where observation is destructive, costly, and altogether difficult. With destructive observations, N_{steps} systems must be initialized and then evolved together under the same policy. Each copy is consumed one by one, as observations are required for decision making, thus incurring an unavoidable $N_{\text{steps}}^2/2$ penalty in the destructive case. In this sense, it is difficult to consider RL to be superior; prescheduled SA simply does not require observation. However, if the choice to observe were to be incorporated into the action set of the RL algorithm, the agent would choose when observation would be necessary.

For example, in the systems presented in this work, the correlation time of the observations is initially small; the temperatures are high, and frequent observations are required to guide the system through phase space. As the system cools, however, the correlation time grows exponentially, and the observations become much more similar to each previous observation; in this case, it would be beneficial to forgo some expensive observations, as the system would not be evolving substantially. With such a scheme, RL stands a better chance at achieving greater performance.

D. Policy analysis

To glean some understanding into what the RL agent is learning, we train an additional model on a well-understood Hamiltonian, the ferromagnetic Ising model of size 16×16 . In this case, the temperatures are initialized randomly (as in the WSC model). This model is the extreme case of a spin glass, with all $J_{ij} = 1$. In Figure 5a, we display the density of states $g(M, E)$ of the Ising model, plotted in phase space, with axes of magnetization per spin (M/L) and energy per spin (E/L). The density of states is greatest in the high-entropy $M = E = 0$ region, and lowest in the low-entropy “corners”. We show the spin configurations at the three corners (“checkerboard”, spin-up, and spin-down) for clarity. The density of states is obtained numerically using Wang–Landau sampling [54]. Magnetization and energy combinations outside of the dashed “triangle” are impossible.

In Figure 5b, we plot a histogram of the average value function $V(s_t)$ on the phase plane, as well as three trajectories. Note that since each observation s_t is composed of N_{reps} replicas, we count each observation as N_{reps} separate points on the phase plot when comput-

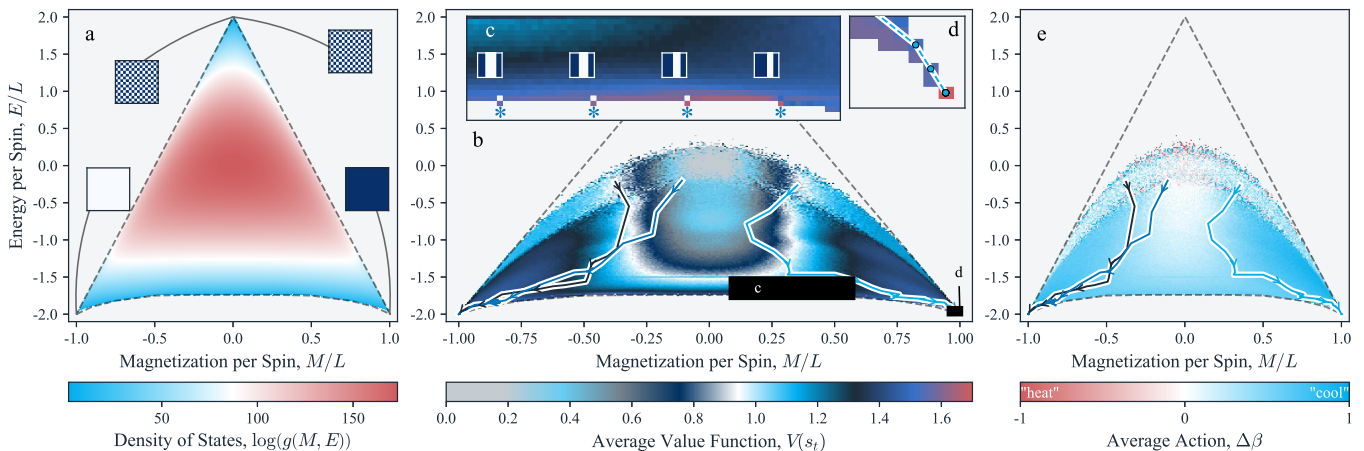


FIG. 5. We train an agent on a special case of the spin-glass Hamiltonians: the 16×16 ferromagnetic Ising model where all couplings $J_{ij} = 1$. (a) We plot the density of states $\log(g(M, E))$ for the 16×16 Ising model in the phase space of energy and magnetization, sampled numerically using the Wang–Landau algorithm [54], and indicate four of the novel high- and low-energy spin configurations on a grid. (b) For the trained model, we plot the average of the learned value function $V(s_t)$ for each possible energy–magnetization pair. Additionally, we plot the trajectories of the first replica for three episodes of annealing to demonstrate the path through phase space the algorithm learns to take. In (c) and (d), we enlarge two high-value regions of interest. In (e), we plot the average action taken at each point in phase space, as well as the same two trajectories plotted in (b).

ing the histogram, each with an identical contribution of $V(s_t)$ to the average. As expected, the learned value function trends higher toward the two global energy minima. The lowest values are present in the initialization region (the high-energy band along the top). We expand two regions of interest in Figure 5c–d. In Figure 5d, we can see that the global minimum is assigned the highest value; this is justifiable in that if the agent reaches this point, it is likely to remain here and reap a high reward so long as the agent keeps the temperature low for the remainder of the episode.

In Figure 5c, we identify four noteworthy energy–magnetization combinations, using asterisks. These four energy–magnetization combinations have identical energies, with increasing magnetization, and correspond to banded spin structures of decreasing width (four example spin configurations are shown). The agent learns to assign a higher value to the higher-magnetization structures, even though the energy, which is the true measure of “success”, is identical. This is because the higher-magnetization bands are closer to the right-most global minimum in action space, that is, the agent can traverse from the small-band configuration to the ground state in fewer spin flips than if traversing from the wide-band configurations.

In Figure 5e, we plot a histogram of the average action taken at each point in phase space. The upper high-energy band exhibits more randomness in the actions chosen, as this is the region in which the system lands upon initialization. When initialized, the temperature is at a randomly drawn value, and sometimes the agent must first heat the system to escape a local minimum before then cooling, and thus the first action is, on average,

of very low magnitude. As the agent progresses toward the minimum, the agent becomes more aggressive in cooling the system, thereby thermally trapping itself in lower energy states.

VIII. CONCLUSION

In this work, we show that reinforcement learning is a viable method for learning dynamic control schemes for the task of simulated annealing (SA). We show that, on a simple spin model, a reinforcement learning (RL) agent is capable of devising a temperature control scheme that can consistently escape a local minimum, and then anneal to the ground state. It arrives at a policy that generalizes to a range of initialization temperatures; in all cases, it learns to cool the system. However, if the initial temperature is too low, the RL agent learns to first increase the temperature to provide sufficient thermal energy to escape the local minimum.

We then demonstrate that the RL agent is capable of learning a policy that can generalize to an entire class of Hamiltonians, and that the problem need not be restricted to a single set of couplings. By training multiple RL agents on increasing numbers of variables (increasing lattice sizes), we investigate the scaling of the RL algorithm and find that it outperforms a linear SA schedule both in absolute terms and in terms of its scaling.

We analyze the value function that the agent learns and see that it attributes an intuitive representation of value to specific regions of phase space.

We discuss the nature of RL in the physical sciences, specifically in situations where observing systems is de-

structive (“destructive observation”) or costly (e.g., performing quantum computations where observations collapse the wavefunction, or conducting chemical analysis techniques that destroy a sample material). We demonstrate that our implementation of RL is capable of performing well in a destructive observation situation, albeit inefficiently. We propose that the future of physical RL (i.e., RL in the physical sciences) will be one of “controlled observation”, where the algorithm can choose when an observation is necessary, minimizing the inher-

ent costs incurred when observations are expensive, slow, or difficult.

IX. ACKNOWLEDGEMENTS

I.T. acknowledges support from NSERC. K.M. acknowledges support from Mitacs. The authors thank Marko Bucyk for reviewing and editing the manuscript.

-
- [1] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, *Science* **220**, 671 (1983).
- [2] F. Barahona, *Journal of Physics A: Mathematical and General* **15**, 3241 (1982).
- [3] D. Sherrington and S. Kirkpatrick, *Physical Review Letters* **35**, 1792 (1975).
- [4] E. Ising, *Zeitschrift für Physik* **31**, 253 (1925).
- [5] L. Onsager, *Physical Review* **65**, 117 (1944).
- [6] A. E. Ferdinand and M. E. Fisher, *Physical Review* **185**, 832 (1969).
- [7] A. Lucas, *Frontiers in Physics* **2**, 1 (2014).
- [8] B. Y. W. K. Hastings, , 97 (1970).
- [9] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, *The Journal of Chemical Physics* **21**, 1087 (1953).
- [10] S. Kirkpatrick, *Journal of Statistical Physics* **34**, 975 (1984).
- [11] P. J. M. van Laarhoven and E. H. L. Aarts, *Simulated Annealing: Theory and Applications* (Springer Netherlands, Dordrecht, 1987).
- [12] D. G. Bounds, *Nature* **329**, 215 (1987).
- [13] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, *Science* **292**, 472 (2001), arXiv:0104129 [quant-ph].
- [14] I. Hen and A. P. Young, *Physical Review A - Atomic, Molecular, and Optical Physics* **86**, 1 (2012), arXiv:1207.1712.
- [15] S. Boixo, T. F. Rønnow, S. V. Isakov, Z. Wang, D. Wecker, D. A. Lidar, J. M. Martinis, and M. Troyer, *Nature Physics* **10**, 218 (2014), arXiv:1304.4595.
- [16] Z. Bian, F. Chudak, R. Israel, B. Lackey, W. G. Mcready, and A. Roy, *Frontiers in Physics* **2**, 1 (2014).
- [17] D. Venturelli, D. J. J. Marchand, and G. Rojo, , 1 (2015), arXiv:1506.08479.
- [18] P. Ray, B. K. Chakrabarti, and A. Chakrabarti, *Physical Review B* **39**, 11828 (1989).
- [19] R. Martoák, G. E. Santoro, and E. Tosatti, *Physical Review B - Condensed Matter and Materials Physics* **66**, 1 (2002).
- [20] G. E. Santoro, *Science* **295**, 2427 (2002), arXiv:0205280 [cond-mat].
- [21] A. Finnila, M. Gomez, C. Sebenik, C. Stenson, and J. Doll, *Chemical Physics Letters* **219**, 343 (1994).
- [22] T. Kadowaki and H. Nishimori, *Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics* **58**, 5355 (1998), arXiv:9804280 [cond-mat].
- [23] R. Harris, J. Johansson, A. J. Berkley, M. W. Johnson, T. Lanting, S. Han, P. Bunyk, E. Ladizinsky, T. Oh, I. Perminov, E. Tolkacheva, S. Uchaikin, E. M. Chapple, C. Enderud, C. Rich, M. Thom, J. Wang, B. Wilson, and G. Rose, *Physical Review B - Condensed Matter and Materials Physics* **81**, 1 (2010), arXiv:0909.4321.
- [24] R. Harris, M. W. Johnson, T. Lanting, A. J. Berkley, J. Johansson, P. Bunyk, E. Tolkacheva, E. Ladizinsky, N. Ladizinsky, T. Oh, F. Cioata, I. Perminov, P. Spear, C. Enderud, C. Rich, S. Uchaikin, M. C. Thom, E. M. Chapple, J. Wang, B. Wilson, M. H. Amin, N. Dickson, K. Karimi, B. MacReady, C. J. Truncik, and G. Rose, *Physical Review B - Condensed Matter and Materials Physics* **82**, 1 (2010), arXiv:1004.1628.
- [25] M. W. Johnson, M. H. Amin, S. Gildert, T. Lanting, F. Hamze, N. Dickson, R. Harris, A. J. Berkley, J. Johansson, P. Bunyk, E. M. Chapple, C. Enderud, J. P. Hilton, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Oh, I. Perminov, C. Rich, M. C. Thom, E. Tolkacheva, C. J. Truncik, S. Uchaikin, J. Wang, B. Wilson, and G. Rose, *Nature* **473**, 194 (2011).
- [26] C. C. McGeoch and C. Wang, *Proceedings of the ACM International Conference on Computing Frontiers*, CF 2013 10.1145/2482767.2482797 (2013).
- [27] K. Ikeda, Y. Nakamura, and T. S. Humble, *Scientific Reports* **9**, 1 (2019), arXiv:1904.12139.
- [28] N. G. Dickson, M. W. Johnson, M. H. Amin, R. Harris, F. Altomare, A. J. Berkley, P. Bunyk, J. Cai, E. M. Chapple, P. Chavez, F. Cioata, T. Cirip, P. Debuen, M. Drew-Brook, C. Enderud, S. Gildert, F. Hamze, J. P. Hilton, E. Hoskinson, K. Karimi, E. Ladizinsky, N. Ladizinsky, T. Lanting, T. Mahon, R. Neufeld, T. Oh, I. Perminov, C. Petroff, A. Przybysz, C. Rich, P. Spear, A. Tcaciuc, M. C. Thom, E. Tolkacheva, S. Uchaikin, J. Wang, A. B. Wilson, Z. Merali, and G. Rose, *Nature Communications* **4**, 1 (2013).
- [29] S. Okada, M. Ohzeki, and K. Tanaka, (2019), arXiv:1904.01522.
- [30] D. A. Battaglia, G. E. Santoro, and E. Tosatti, *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* **71**, 1 (2005), arXiv:0502468 [cond-mat].
- [31] S. Tsukamoto, M. Takatsu, S. Matsubara, and H. Tamura, *Fujitsu Scientific and Technical Journal* **53**, 8 (2017).
- [32] T. Inagaki, Y. Haribara, K. Igarashi, T. Sonobe, S. Tamate, T. Honjo, A. Marandi, P. L. McMahon, T. Umeki, K. Enbutsu, O. Tadanaga, H. Takenouchi, K. Aihara, K. I. Kawarabayashi, K. Inoue, S. Utsunomiya, and H. Takesue, *Science* **354**, 603 (2016).
- [33] E. Farhi, J. Goldstone, and S. Gutmann, , 1 (2014), arXiv:1411.4028.

- [34] E. Farhi, J. Goldstone, S. Gutmann, and L. Zhou, , 1 (2019), arXiv:1910.08187.
- [35] R. Sutton and A. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. (The MIT Press, 2018).
- [36] C. Berner, G. Brockman, B. Chan, V. Cheung, P. P. Dębiak, C. Dennison, D. Farhi, Q. Fischer, S. Hashme, C. Hesse, R. Józefowicz, S. Gray, C. Olsson, J. Pachocki, M. Petrov, H. Pondé, D. O. Pinto, J. Raiman, T. Salimans, J. Schlatter, J. Schneider, S. Sidor, I. Sutskever, J. Tang, F. Wolski, and S. Zhang, (2019).
- [37] Z. Zhang, H. Li, L. Zhang, T. Zheng, T. Zhang, X. Hao, X. Chen, M. Chen, F. Xiao, and W. Zhou, (2019), arXiv:1901.08004.
- [38] O. Vinyals, I. Babuschkin, W. M. Czarnecki, M. Mathieu, A. Dudzik, J. Chung, D. H. Choi, R. Powell, T. Ewalds, P. Georgiev, J. Oh, D. Horgan, M. Kroiss, I. Danihelka, A. Huang, L. Sifre, T. Cai, J. P. Agapiou, M. Jaderberg, A. S. Vezhnevets, R. Leblond, T. Pohlen, V. Dalibard, D. Budden, Y. Sulsky, J. Molloy, T. L. Paine, C. Gulcehre, Z. Wang, T. Pfaff, Y. Wu, R. Ring, D. Yogatama, D. Wünsch, K. McKinney, O. Smith, T. Schaul, T. Lillicrap, K. Kavukcuoglu, D. Hassabis, C. Apps, and D. Silver, *Nature* 10.1038/s41586-019-1724-z (2019).
- [39] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, arXiv preprint arXiv: {...} , 1 (2013), arXiv:1312.5602.
- [40] D. Silver, A. Huang, C. J. Maddison, A. Guez, L. Sifre, G. van den Driessche, J. Schrittwieser, I. Antonoglou, V. Panneershelvam, M. Lanctot, S. Dieleman, D. Grewe, J. Nham, N. Kalchbrenner, I. Sutskever, T. Lillicrap, M. Leach, K. Kavukcuoglu, T. Graepel, and D. Hassabis, *Nature* **529**, 484 (2016).
- [41] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, Y. Chen, T. Lillicrap, F. Hui, L. Sifre, G. van den Driessche, T. Graepel, and D. Hassabis, *Nature* **550**, 354 (2017).
- [42] D. Silver, T. Hubert, J. Schrittwieser, I. Antonoglou, M. Lai, A. Guez, M. Lanctot, L. Sifre, D. Kumaran, T. Graepel, T. Lillicrap, K. Simonyan, and D. Hassabis, *Science* **362**, 1140 (2018).
- [43] F. Agostinelli, S. McAleer, A. Shmakov, and P. Baldi, *Nature Machine Intelligence* **1**, 356 (2019).
- [44] OpenAI, I. Akkaya, M. Andrychowicz, M. Chociej, M. Litwin, B. McGrew, A. Petron, A. Paino, M. Plappert, G. Powell, R. Ribas, J. Schneider, N. Tezak, J. Tworek, P. Welinder, L. Weng, Q. Yuan, W. Zaremba, and L. Zhang, , 1 (2019), arXiv:1910.07113.
- [45] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, , 1 (2017), arXiv:1707.06347.
- [46] A. Hill, A. Raffin, M. Ernestus, A. Gleave, A. Kanervisto, R. Traore, P. Dhariwal, C. Hesse, O. Klimov, A. Nichol, M. Plappert, A. Radford, J. Schulman, S. Sidor, and Y. Wu, *Stable baselines*, <https://github.com/hill-a/stable-baselines> (2018).
- [47] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, , 1 (2016), arXiv:1606.01540.
- [48] S. Hochreiter and J. Schmidhuber, *Neural Computation* **9**, 1735 (1997).
- [49] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, , 1 (2015), arXiv:1506.02438.
- [50] M. Aramon, G. Rosenberg, E. Valiante, T. Miyazawa, H. Tamura, and H. G. Katzgraber, *Frontiers in Physics* **7**, 10.3389/fphy.2019.00048 (2019), arXiv:1806.08815.
- [51] P. I. Bunyk, E. M. Hoskinson, M. W. Johnson, E. Tolkacheva, F. Altomare, A. J. Berkley, R. Harris, J. P. Hilton, T. Lanting, A. J. Przybysz, and J. Whittaker, *IEEE Transactions on Applied Superconductivity* **24**, 1 (2014), arXiv:1401.5504.
- [52] F. Liers, M. Jünger, G. Reinelt, and G. Rinaldi, *New Optimization Algorithms in Physics* , 47 (2005).
- [53] M. Jünger, *Spin Glass Server* (<https://informatik.uni-koeln.de/spinglass/>).
- [54] F. Wang and D. P. Landau, *Physical Review Letters* **86**, 2050 (2001), arXiv:0011174 [cond-mat].