

# Combining Reinforcement Learning and Tensor Networks, with an Application to Dynamical Large Deviations

Edward Gillman,<sup>1,2</sup> Dominic C. Rose,<sup>3</sup> and Juan P. Garrahan<sup>1,2</sup>

<sup>1</sup>*School of Physics and Astronomy, University of Nottingham, Nottingham, NG7 2RD, UK*

<sup>2</sup>*Centre for the Mathematics and Theoretical Physics of Quantum Non-Equilibrium Systems, University of Nottingham, Nottingham, NG7 2RD, UK*

<sup>3</sup>*Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, United Kingdom*

We present a framework to integrate tensor network (TN) methods with reinforcement learning (RL) for solving dynamical optimisation tasks. We consider the RL actor-critic method, a model-free approach for solving RL problems, and introduce TNs as the approximators for its policy and value functions. Our “actor-critic with tensor networks” (ACTeN) method is especially well suited to problems with large and factorisable state and action spaces. As an illustration of the applicability of ACTeN we solve the exponentially hard task of sampling rare trajectories in two paradigmatic stochastic models, the East model of glasses and the asymmetric simple exclusion process (ASEP), the latter being particularly challenging to other methods due to the absence of detailed balance. With substantial potential for further integration with the vast array of existing RL methods, the approach introduced here is promising both for applications in physics and to multi-agent RL problems more generally.

**Introduction.** Tensor networks (TNs), routinely used in the study of quantum many-body systems [1–7], are increasing being applied in machine learning (ML), see e.g. Refs. [8–20]. Both domains often deal with systems with state-spaces which are exponentially large in the number of degrees of freedom, say the number of qubits in a quantum system, or of pixels in images to be classified. In such situations TNs provide a powerful way to represent functions, vectors and distributions, while allowing for efficient sampling and computation of quantities such as inner products and norms.

To date, the intersection of TNs and ML has been mostly in supervised and unsupervised learning, see e.g. Refs. [8–20]. In contrast, the combination of TNs and reinforcement learning (RL) [21] has been more limited, despite recent major advances in RL [22–26]. While some promising related directions have been explored, such as the approximation of Q-functions in the context of large state-spaces [27], and/or large action-spaces [28], the flexible integration of TNs with RL remains an open problem, along with demonstrating useful applications.

Here, we introduce the *actor-critic with tensor networks* (ACTeN) method, a general framework for integrating TNs into RL via actor-critic (AC) techniques. By combining decision-making “actors” with “critics” that judge an actor’s quality, AC methods are used in many state-of-the-art RL applications. Using TNs as the basis for modelling actors and critics within AC and RL represents a powerful combination to tackle problems with both large state and action spaces.

To demonstrate the effectiveness of our approach, we consider the problem of computing the large deviation (LD) statistics of dynamical observables in classical stochastic systems [29–35], and of optimally sampling the associated rare long-time trajectories [36–48]. Such problems are of wide interest in statistical mechan-

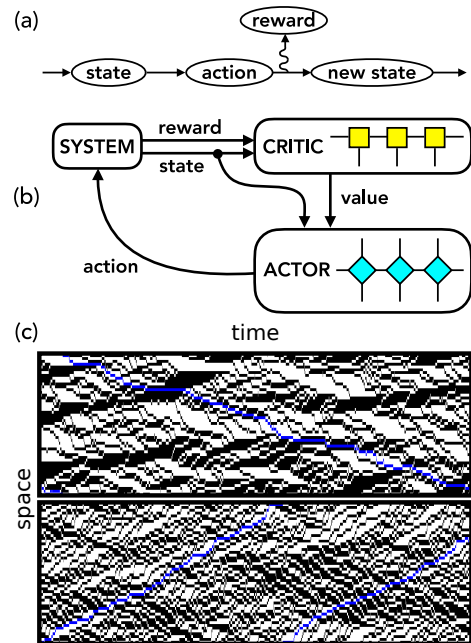


FIG. 1. **Actor-Critic with tensor networks (ACTeN)** (a) Sketch of a Markov decision process. (b) In actor-critic RL, the state is passed to an “actor”, which chooses the action, and to a “critic”, which values the state given the reward. This value is used to improve the actor’s policy. In ACTeN, the function approximators for actor and critic are tensor networks. (c) Top: typical trajectory of the ASEP at half-filling and  $L = 50$  sites with one particle highlighted (blue), shown for 3000 steps. Bottom: trajectory with a current large deviation, sampled from the ACTeN solution for biasing (counting) field  $\lambda = -3$ . See the text for details.

ics and can be phrased straightforwardly as an optimization problems that may be solved with RL [49–52] and similar techniques [53–61]. For concreteness we consider two models: (i) the East model, a kinetically constrained

model used to study slow glassy dynamics, and (ii) the asymmetric exclusion process (ASEP), a paradigmatic model of non-equilibrium, in which particles hop around a lattice while blocking each others movement. In particular, and in contrast to the East model, the ASEP (with periodic boundaries) does not obey detailed balance [62], and thus evades straightforward use of TNs to compute spectral properties of the relevant dynamical generators. We demonstrate that ACTeN can be applied to both problems irrespective of the equilibrium/non-equilibrium distinction, by computing their dynamical LDs for sizes well beyond those achievable with exact methods. Given the vast array of options for improving the RL algorithm that we use, our results indicate that the overall framework outlined here is highly promising for applications more generally.

**Background: Reinforcement Learning and Actor-Critic.** A discrete-time Markov decision process (MDP) [21] consists at each time  $t \in [0, T]$  of stochastic variables  $X_t = (S_t, a_t, R_t)$ , named *state*, *action* and *reward*. We assume these are drawn from  $t$ -independent finite sets,  $\mathcal{S}, \mathcal{A}$  and  $\mathcal{R}$ , where the action set may depend on the current state,  $\mathcal{A}(S)$  for  $S \in \mathcal{S}$ . The action and state variables are associated with the *policy*,  $\pi(a|S)$ , and *environment*,  $P(S'|S, a)$ , distributions. These are sampled in a sequence of steps to generate a trajectory of the MDP,  $\omega = (X_0, X_1, \dots, X_T)$ , see Figure 1(a). We assume that the reward is a deterministic function of the state and action variables.

In the typical scenario of *policy optimisation*, the policy is controllable and known, while the environment is fixed and potentially unknown. We focus on MDPs that are “continuing”, and admit a steady state distribution independent of  $X_0$ . We can then define the average reward per time-step when following a given policy as  $r(\pi) = \lim_{t \rightarrow \infty} \mathbb{E}_\pi[R_t]$ , where  $\mathbb{E}_\pi[\cdot]$  is the stationary state expectation over states and over transitions from those states according to policy  $\pi$ . The task of policy optimisation is to find the policy  $\pi^*$  that maximises  $r(\pi)$ .

Reinforcement learning (RL) refers to the group of methods that aim to discover optimal policies by using the experience gained from sampling trajectories of an MDP. In *policy gradient* methods, the policy is approximated by a function  $\pi_w(a|S)$  with parameters  $w$ , and optimized using the gradient of  $r(\pi_w)$  with respect to  $w$ . Building on this, actor-critic (AC) methods then assess  $\pi_w(a|S)$  by computing the value,  $v_\pi(S)$ , of the states that result when following the policy. This value is defined as the difference between rewards in the future of a given state and the average reward,  $v_\pi(S) = \mathbb{E}_\pi[\sum_{\tau=t}^{\infty} [R_\tau - r(\pi)] | S]$ . The gradient of  $r(\pi_w)$  can be written exactly in terms of these values as

$$\nabla_w r(\pi) = \mathbb{E}_\pi [\delta_t^\pi \nabla_w \ln \pi_w(a_t|S_t)] , \quad (1)$$

where we have introduced the temporal difference (TD) error,  $\delta_t^\pi = v_\pi(S_{t+1}) + R_{t+1} - r(\pi_w) - v_\pi(S_t)$  [21, 49], which quantifies if a resultant state is better than the

current one. AC methods use this information to alter the probability of taking that action in the future.

In reality, calculating the true value of every state under the current policy is impractical, and thus an auxiliary approximation for the value-function,  $v_\psi(S)$  with parameters  $\psi$ , is introduced, the so-called “critic”. To optimize the critic, we note that the value of a state is related to the value of states reachable from it, as encoded in the differential Bellman equation [21],

$$v_\pi(S_t) = \mathbb{E}_\pi [v_\pi(S_{t+1}) + R_{t+1} - r(\pi_w) | S_t] . \quad (2)$$

Minimizing the error in the Bellman equation when substituting the critic for the true values can be done by updating the weight as  $\psi' = \psi + \alpha \mathbb{E}_\pi [\delta_t \nabla_\psi v_\psi(S)(S_t)]$ , in terms of the approximate TD error,  $\delta_t = v_\psi(S_{t+1}) + R_{t+1} - r(\pi_w) - v_\psi(S_t)$ , and a learning rate  $\alpha$ . Intuitively, the estimated expected reward before a transition occurs,  $v_\psi(S_t)$ , is compared to the expected reward afterwards plus the true reward for that time-step,  $v_\psi(S_{t+1}) + R_{t+1} - r(\pi_w)$ , and  $v_\psi(S)$  adjusted to make these closer. The policy is then updated by following the gradient in Eq. (1) with the exact TD error  $\delta^\pi$  replaced by the critic’s approximate TD error  $\delta$ .

**Analytic Example: Two-Site East Model.** To illustrate these ideas, consider two spins,  $s_{1,2} = 0, 1$ , evolving with a constrained set of transitions as in the East model studied below, such that a spin can flip only when its left-neighbour (in this case simply the other spin due to periodic boundaries) is 1. The states are  $\mathcal{S} = \{00, 01, 10, 11\}$ . The dynamics of this model can be implemented as an MDP using a policy that (stochastically) selects which spins to flip. Denoting no-flip/flip by 0/1 and requiring at most one spin flip per time-step, the action-sets for this model are then:  $\mathcal{A}(10) = \{00, 01\}$ ,  $\mathcal{A}(01) = \{00, 10\}$ , and  $\mathcal{A}(11) = \{10, 01\}$ , with the state 00 being disconnected from the rest. An example policy that selects from all possible actions equally would assign a probability of 1/2 to each of these, e.g.  $\pi(a = 00|S = 10) = \pi(a = 01|S = 10) = 1/2$  and similarly for the other states. The transition to a new state is then enacted by the environment. We can choose this to be deterministic, and simply apply the spin-flip operations selected by the actor to the current state,  $s_i \rightarrow (1 - a_i)s_i + a_i(1 - s_i)$ . Finally, an example reward can be defined via the function  $R(S, a, S') = -\lambda(1 - \delta_{S', S})$ , which awards  $-\lambda$  every time a spin flip occurs ( $a \neq 0$ ), encouraging activity/inactivity for negative/positive  $\lambda$ . For this reward, the optimal policy for negative  $\lambda$  will maximize activity, flipping a spin at every step, requiring going from 10 and 01 to 11 with probability 1, i.e.  $\pi^*(01|10) = 1$ , etc.

**Method: Actor-Critic with Tensor Networks (ACTeN).** We now focus on applying AC to solve problems with large state and action spaces. For example, we may wish to find the optimal dynamics of a system of  $L$  binary components, resulting in  $2^L$  states, with individual agents and their actions associated to each component. To ensure optimal choices for a given task, actions

may need to be correlated not only with other agents states, but also the actions other agents are about to take. In such problems, simple approaches such as tabular RL fail due to exponentially large memory requirements and sampling costs, and a common alternative is to use neural networks (NNs) when defining  $\pi_w(a|S)$  and  $v_\psi(S)$ . TNs offer another approach, with polynomial memory and computational costs, and showing state-of-the-art performance in many settings; see e.g. the review [6]. Motivated by this, we define a general framework (which we call ACTeN) that exploits TNs to efficiently represent  $\pi_w(a|S)$  and  $v_\psi(S)$ .

A TN is a set of tensors,  $\mathcal{T} = \{T_{i_1 j_1 k_1 \dots}^{[1]}, T_{i_2 j_2 k_2 \dots}^{[2]}, \dots\}$ , contracted in some pattern, cf. Fig. 1(b). This results in a single tensor that can be viewed as defining a multivariate function,  $\varphi(x) = T_x$ ,  $T_x = \mathcal{C}[\mathcal{T}]$ , where  $\mathcal{C}$  indicates the chosen contractions and  $x$  all remaining uncontracted indices. For a given problem, the selection of an appropriate TN depends on factors such as dimensionality and geometry. Here we consider applying ACTeN to study one-dimensional (1D) systems with periodic boundaries (PBs) and  $L$  components, such that a state is  $S = (s_1, \dots, s_L)$  with  $s_i$  taking  $d$  values. To represent the value function  $v_\psi(S)$  we use a translation invariant matrix product state (MPS) which mirrors the chain geometry of the system. This TN is built from a single real-valued tensor  $\mathcal{A}_{s_{ij}}$  of shape  $(d, \chi, \chi)$  (or equivalently  $d$  square  $\chi$ -dimensional matrices  $[A_s]_{ij} = A_{s_{ij}}$ ) whose elements encode the parameters of  $v_\psi(S)$ ,  $\psi = \{A_{s_{ij}}\}_{s=1}^d$ . For a given  $S = (s_1, \dots, s_L)$  this TN is defined as,

$$\varphi(S) = \text{Tr} [A_{s_1} A_{s_2} \dots A_{s_L}] , \quad (3)$$

i.e., for each site we select the corresponding matrix  $A_{s_i}$ , multiplying the  $L$  matrices together with a trace to produce a real scalar. To take advantage of translation invariance and apply approximations from smaller  $L$  to larger  $L$  systems, we then define the value function in terms of  $\varphi(S)$  after the additional application of a square and log, which prevents the exponential growth or decay of values as  $L$  is changed for fixed  $A_{s_{ij}}$ . Hence,

$$v_\psi(S) = \log [\varphi(S)^2] . \quad (4)$$

To define  $\pi_w(a|S)$  we use a matrix product operator (MPO). This TN is built from a single real-valued  $(d_S, d_A, \chi, \chi)$ -shaped tensor,  $\mathcal{A}_{a_{s_{ij}}}$ , equivalent to  $d_s \times d_A$   $\chi$ -dimensional square matrices  $[A_{as}]_{ij} = A_{a_{s_{ij}}}$ , i.e., one matrix per combination of local state and action. Given a state  $S = (s_1, \dots, s_L)$  and actions  $a = (a_1, \dots, a_L)$ , the contraction is given by the traced matrix product,

$$\varphi(a, s) = \text{Tr} [A_{a_1 s_1} A_{a_2 s_2} \dots A_{a_L s_L}] . \quad (5)$$

To use this to define a policy, we need to ensure positivity and normalization, as well as preventing the policy from producing invalid actions. To achieve this we define

$$\pi_w(a|S) = \mathcal{C}(a, S) [\mathcal{N}(S)]^{-1} \varphi(a, S)^2 , \quad (6)$$

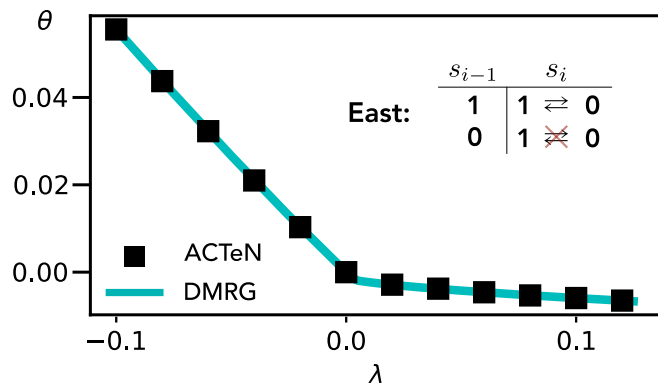


FIG. 2. **Dynamical large deviations in the East model using ACTeN.** Scaled-cumulant generating function for the dynamical activity of the East model as a function of biasing field  $\lambda$  from ACTeN (symbols), for  $L = 50$  and PBC. Our RL results coincide with those obtained from the current state-of-the-art method using DMRG, cf. Ref. [64] (which is possible since the East model obeys detailed balance). Inset: Kinetic constraint of the East model; a spin,  $s_i$ , can flip,  $s_i \rightarrow 1 - s_i$ , only if the spin to the left is up,  $s_{i-1} = 1$ .

where  $\mathcal{N}(S) = \sum_{a|\mathcal{C}(a,S)=1} \varphi(a, S)^2$  is the (state-dependent) normalisation factor and  $\mathcal{C}(a, S)$  returns one if an action  $a$  is possible in state  $S$  or zero otherwise [63].

**Application: Dynamical Large Deviations.** To test ACTeN we consider the problem of computing the large deviations (LDs) of trajectory observables [29, 31, 33] in the East model and the ASEP in 1D with PBs. Both models are many-body binary spin systems with large state-spaces for large  $L$ ,  $S = \{s_i\}_{i=1}^L$ , with  $s_i = 0, 1$  (see the Appendix for details on the models). The dynamics of these systems are subject to local constraints that lead to rich behaviours in their trajectories,  $\omega_0^T = \{S_t\}_0^T$ . This can be observed in the time-integrals of time-local quantities,  $O(\omega_0^T) = \sum_{t=1}^T o(S_t, S_{t-1})$ , the moments of which are contained in derivatives of the moment generating function (MGF),  $Z_T(\lambda) = \sum_{\omega_0^T} e^{-\lambda O(\omega_0^T)} P(\omega_0^T)$ , where  $P(\omega_0^T) = \prod_{t=1}^T P(S_t|S_{t-1})P(S_0)$  is the trajectory probability under the dynamics.

In the long-time limit, the MGF obeys a large deviation principle [29, 31, 33] with the scaled cumulant generating function (SCGF),  $\theta(\lambda) = \lim_{T \rightarrow \infty} \frac{1}{T} \ln Z_T(\lambda)$ , playing the role of a free-energy for trajectories. In principle the SCGF can be obtained by sampling methods. However, this is exponentially hard (in time and space) using the original dynamics. An alternative is to find a more efficient sampling dynamics which may then be combined with importance sampling to obtain unbiased statistics. This can be formulated as a RL problem as follows: can we find a parameterized dynamics  $P_w(S_t|S_{t-1})$  such that  $P_w(\omega_0^T) = e^{-\lambda O(\omega_0^T)} P(\omega_0^T) / Z_T(\lambda)$ , i.e. it reproduces a trajectory ensemble biased towards rare trajectories of the original dynamics. This dynamics is connected to an underlying policy  $\pi_w(a|S)$  by a deterministic environment

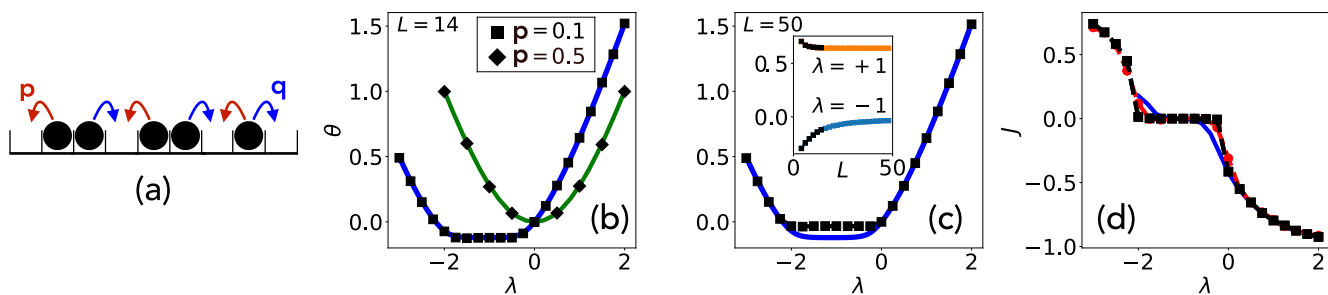


FIG. 3. **Dynamical large deviations in the ASEP using ACTeN.** (a) In the ASEP particles can only move to an unoccupied neighbouring site, with probability  $p$  to the left and  $q = 1 - p$  to the right. (b) SCGF for the time-integrated particle current as a function of biasing field. We show results from ACTeN for  $p = 0.1$  (squares) and  $p = 1/2$  (diamonds). The lack of detailed balance for PBC and  $p \neq 1/2$  prevents straightforward application of DRMGM, but for small sizes (here  $L = 14$ ) we can compare to exact diagonalisation (blue curve for  $p = 0.1$ , green for  $p = 1/2$ ). (c) SCGF for  $p = 0.1$  from ACTeN for size  $L = 50$  which is beyond the scope of ED. Compared to  $L = 14$  (blue curve from ED), we see that ACTeN captures the flattening of the SCGF for larger sizes indicative of a LD phase transition, cf. Ref. [32]. The inset shows the smooth convergence of our ACTeN numerics with  $L$  for two values of  $\lambda$ . (d) Since ACTeN provides direct access to the optimal dynamics, observables such as the time-integrated current can be evaluated directly (black squares for  $L = 50$ ). We show for comparison the numerical differentiation of the ACTeN SCGF (red circles) and of the ED SCGF at  $L = 14$  (blue line).

which returns states after receiving an associated action, i.e.  $P_w[S' = f(a, S)|S] = \pi_w(a|S)$ , where for each  $S$ ,  $f(a, S)$  returns a unique  $S'$  for each  $a$ . For example, in the East model if we take the action  $a = \{a_i\}_{i=1}^L$  then sites with  $a_i = 1$  are flipped and those with  $a_i = 0$  are not flipped; the new state is then  $S' = f(a, S) = \{(1 - a_i)s_i + a_i(1 - s_i)\}_{i=1}^L$ .

Optimizing the KL divergence between the two trajectory ensembles gives a regularized form of RL with a reward depending on the policy [49]

$$R_t = -\lambda o(S_t, S_{t-1}) - \ln \left( \frac{P_w(S_t|S_{t-1})}{P_{\text{orig}}(S_t|S_{t-1})} \right), \quad (7)$$

with its expected value becoming the SCGF at optimality [49]. Intuitively, choosing actions (e.g. flips) to maximize the first term increases the likelihood of rare events with extreme values of the observable, while maximizing the second term minimizes the difference between the parameterized and original dynamics, thus making the event more probable. Maximizing this reward is a balancing act between these two aims, resulting in dynamics biased towards rare events in a way representative of their occurrence in the original dynamics. In the appendices, we illustrate ACTeN by solving explicitly the 2-site East model and showing how this can be exactly represented by the TN ansatz.

(i) *East model and dynamical activity*: Figure 2 shows the SCGF of the dynamical activity [total number of spin flips in a trajectory, defined by  $o(S_t, S_{t-1}) = 1 - \delta_{S_t, S_{t-1}}$ ], calculated using ACTeN (symbols). Since the East model obeys detailed balance, the SCGF is the log of the largest eigenvalue of a Hermitian operator and can be estimated via density matrix renormalisation group (DMRG) methods, cf. Ref. [64–70] (here we ITensors.jl [71]). Figure 2 shows that the DMRG results (blue curve) coincide with ACTeN (black squares)

for size  $L = 50$ , which is well beyond what is accessible to exact diagonalisation (ED). Note that DMRG with PBs tends to be much less numerically stable than for open boundaries. Nonetheless, ACTeN can reach  $L \gtrsim 50$  without the need for any special stabilisation techniques.

(ii) *ASEP and particle current*: Figure 3 presents the LDs of the time-integrated particle current, defined by  $o(s_t, s_{t-1}) = \frac{1}{2} \sum_{i=1}^L s_{t-1}^i s_t^{i+1} - s_t^i s_{t-1}^{i+1}$ . Figure 3(b) shows the SCGF obtained via ACTeN (black squares/diamonds). Unlike the East model, for asymmetric hops ( $p \neq 1/2$ ) Hermitian DMRG cannot be applied directly to the ASEP, so for comparison we show results from exact diagonalisation for both  $p = 0.1$  (blue line) and  $p = 1/2$  for  $L = 14$ . Beyond  $L = 14$  ED becomes prohibitive, while ACTeN remains feasible. Figures 3(c,d), show the expected phase transition behaviour [32] and convergence with  $L$  up to  $L = 50$ . The optimal dynamics itself, i.e. the learnt policy, can be used to generate trajectories representative of  $\lambda \neq 0$ , see Fig. 1(c), and directly sample rare values of the integrated current, see Fig. 3(d).

**Outlook.** ACTeN compares very favourably with state-of-art methods for computing rare events without some of the limitations, such as boundary conditions or detailed balance. From the corpus of research in both TNs and RL, our approach has considerable potential for further improvement and exploration. These include: numerical improvements to precision via hyper-parameter searches; stabilisation strategies for large systems; integration with trajectory methods such as transition path sampling or cloning; integration with advanced RL methods such as those offered by the DeepMind ecosystem [72]; generalisation to continuous-time dynamics; and applications to other multi-agent RL problems, such as PistonBall [73], via integration with additional pro-



cessing layers particularly those for image recognition.

**Acknowledgement.** We would like to thank Christopher J. Turner for useful discussions. We acknowledge funding from The Leverhulme Trust grant no. RPG-2018-181, EPSRC Grant No. EP/V031201/1, and University of Nottingham grant no. FiF1/3. We are grateful for access to the University of Nottingham’s Augusta HPC service. DCR was supported by funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No. 853368). We thank the creators and community of the Python programming language [74], and acknowledge use of the packages JAX [75], NumPy [76], pandas [77], Matplotlib [78] and h5py [79]. We also thank the creators and community of the Julia programming language [80], and acknowledge use of the packages ITensors.jl [71] and HDF5.jl [81].

### Appendix on kinetically constrained models:

(i) *East Model.* Flipping of a spin is constrained on the spin to its left being in state 1 [31, 82]. Dynamics then amounts to two steps: first, select a random site  $i$  with probability  $1/L$ ; second, if spin  $s_{i-1} = 1$  then flip  $s_i$ . Given  $N$  spins in state 1 and periodic boundary conditions, the transition probability is  $P(s'|s) = \frac{1}{L}$  for each possible new state  $s' \neq s$ , and probability  $P(s|s) = 1 - \frac{N}{L}$  for no flip occurring.

(ii) *Asymmetric simple exclusion process.* The constraint is particle exclusion: a particle at site  $i$  ( $s_i = 1$ ) can move left or right only if the destination is unoccupied [83]. The movement of a particle to, say, the right thus corresponds to  $10 \rightarrow 01$ , i.e. a flip of both spin variables. The dynamics again amounts to two steps: first, select a particle with probability  $1/N$ , with  $N = \sum_i s_i$  the particle number; second, choose whether this particle hops right or left with probabilities  $p$  or  $1-p$ , respectively, with the hop occurring if the new site is unoccupied. The transition probabilities are then:  $P(s'|s) = p/N$  for a right hop;  $P(s'|s) = (1-p)/N$  for a left hop; and, given the number of neighbouring particles  $N_{\text{nn}} = \sum_{i=1}^L s_i s_{i+1}$ , with  $s_{L+1} := s_1$ , the probability of no change is  $P(s|s) = 1 - N_{\text{nn}}/N$ . In the main text, we consider the case of half-filling,  $N = L/2$ .

### Appendix on analytical solution of the two-site East model with tensor network ansatz:

An exact solution for the optimal policy in the two-site East model can be found by analytically constructing the so-called “Doob dynamics” [34, 84] as follows. First we define a “tilted” evolution operator [29, 31, 33]  $P_\lambda(S_t|S_{t-1}) = e^{-\lambda o(S_t, S_{t-1})} P(S_t|S_{t-1})$ , such that  $Z_T(\lambda) = \langle -| P_\lambda^T |P_{\text{ss}} \rangle$  where  $\langle -| = \sum_S \langle S|$  and  $|P_{\text{ss}} \rangle$  is the stationary state vector for  $P$ . In the large  $T$  limit the SCGF is the log of the largest eigenvalue of  $P_\lambda$  [29, 31, 33]. The corresponding left eigenvector  $l_\lambda$  is related to the dynamics that maximises the expected value of Eq. (7), given by the so-called Doob (or optimal) dynamics  $P_\lambda^D(S_t|S_{t-1}) = \frac{l_\lambda(S_{t-1})P_\lambda(S_t|S_{t-1})}{e^{\theta(\lambda)l_\lambda(S_{t-1})}}$ . Ap-

plied to the two-site East model, defining the function  $a(\lambda) = \frac{4}{1+\sqrt{1+8e^{-2\lambda}}}$ , we find  $\theta(\lambda) = -\ln(a(\lambda))$ , with optimal dynamics

$$\begin{aligned} P_\lambda^D(10|10) &= P_\lambda^D(01|01) = \frac{a(\lambda)}{2}, \\ P_\lambda^D(11|10) &= P_\lambda^D(11|01) = 1 - \frac{a(\lambda)}{2}, \\ P_\lambda^D(10|11) &= P_\lambda^D(01|11) = \frac{1}{2}. \end{aligned}$$

We may then find the corresponding value function by solving the differential Bellman equation (2). To do this, note first that by symmetry the values of states 10 and 01 are identical, and second that Eq. (2) is invariant under an overall shift of the value function by a constant. Therefore, we may choose the value of states 10 and 01 to be 0, and thus find  $V_\lambda(11) = -\lambda - \theta(\lambda)$ .

These results are what is expected intuitively. Trajectories biased towards enhanced activity ( $\lambda < 0$ ) have  $a(\lambda) < 1$ , making  $P_\lambda^D(11|10) = P_\lambda^D(11|01) > 1/2$ , i.e. the system is more likely to transition to the state which is guaranteed to flip at the next step rather than remain in 01 or 10. Furthermore,  $V_\lambda(11) > 0$ , i.e. the state guaranteed to flip is more valuable. In contrast, trajectories biased towards reduced activity ( $\lambda > 0$ ) show the opposite behaviour.

To connect  $P_\lambda^D$  to our policy ansatz Eq. (6), we first rewrite it as an operator. Using projection operators  $P_1 = |1\rangle\langle 1|$ ,  $P_0 = |0\rangle\langle 0|$ , and flip operators  $\sigma_+ = |1\rangle\langle 0|$ ,  $\sigma_- = |0\rangle\langle 1|$ , we define  $A(\lambda) = \sigma_- + a(\lambda)\sigma_+ + [2-a(\lambda)]P_0$ . We may thus write  $P_\lambda^D = 0.5(P_1 \otimes A(\lambda) + A(\lambda) \otimes P_1)$ .

The policy ansatz Eq. (6) involves the element-wise square of an operator  $\varphi(S', S)$ : we thus seek the element-wise square root of  $P_\lambda^D$ . We define  $[\hat{A}(\lambda)]_{ij} = \sqrt{[A(\lambda)]_{ij}}$ . Due to their sparsity structures, we have  $\hat{A}(\lambda) \odot P_1 = 0$ , where  $\odot$  is element-wise (Hadamard) multiplication. Since  $(A \otimes B) \odot (C \otimes D) = (A \odot C) \otimes (B \odot D)$ , we find  $\varphi = \sqrt{0.5(P_1 \otimes \hat{A}(\lambda) + \hat{A}(\lambda) \otimes P_1)}$  is such that  $\varphi \odot \varphi = P_\lambda^D$ . It remains to factor  $\varphi$  into an MPO of the form of Eq. (5). We may rewrite this as  $\varphi = \sum_{d_1, d_2=1}^X T_{d_1 d_2} \otimes T_{d_2 d_1}$  where each  $T_{d' d}$  is a  $2 \times 2$  matrix acting on the single site state space. This  $T$  can be constructed by taking  $\chi = 2$  with  $T_{11} = T_{22} = 0$ ,  $T_{12} = 0.5^{1/4} P_1$ , and  $T_{21} = 0.5^{1/4} \hat{A}(\lambda)$ . We thus find this order-4 tensor  $T$  reproduces the exact Doob dynamics from our translation invariant MPO-based ansatz.

For the value function, taking an exponential and square-root element-wise of the value function to invert Eq. (4) leads to the vector  $\tilde{V}^\lambda = (e^{V_\lambda(11)/2} - 1) |1\rangle \otimes |1\rangle + |-\rangle \otimes |-\rangle$ . Since this is already a sum of symmetric products, it is easy to rewrite it as a translation invariant MPS with  $\chi = 2$ , i.e.  $\tilde{V}_{s_1 s_2}^\lambda = \sum_{d_1, d_2=1}^2 v_{s_1}^{d_1 d_2} v_{s_2}^{d_2 d_1}$ , with order-3 tensor  $v$  such that  $v_1^{11} = \sqrt{e^{V_\lambda(11)/2} - 1}$ ,  $v_1^{22} = v_0^{22} = 1$ , and  $v_s^{d' d} = 0$  otherwise.

### Appendix on training procedure:

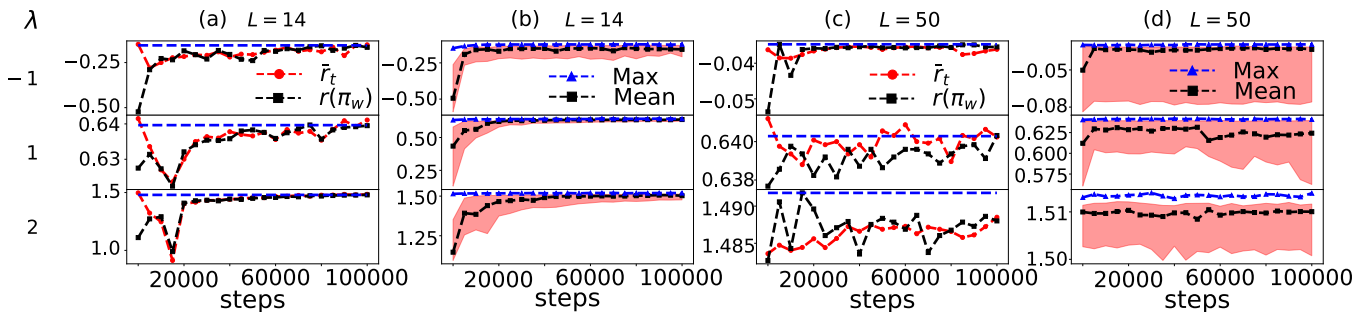


FIG. 4. **Training Procedure and Learning Curves (ASEP)** (a) For each bias [we show  $\lambda = -1$  (top row),  $\lambda = 1$  (middle row),  $\lambda = 2$  (bottom row)] TN-based policies and value-functions are produced via actor-critic optimization. These are initiated at random for  $L = 4$  with  $\chi = 16$  and trained for  $10^6$  steps. Every 5000 training steps the average reward of the policy is evaluated over  $10^4$  steps (black squares) and the weights of the policy (which we call a “snapshot” for that time) are stored. The evaluated values can be compared to the training estimate of  $r(\pi)$  (red circles), which tends to overestimate  $r(\pi)$  initially. The policy snapshot with the highest evaluated  $r$  (blue dashed line) is used to initiate the policy for higher values of  $L$ . This is repeated every  $\Delta L = 2$  up to  $L = 50$ , with  $L = 14$  shown here. (b) For each bias, several policies (here six) are independently trained via the same procedure from different random initial conditions. This produces a distribution of evaluated average rewards, here represented by the median (black squares) and inter-quartile range (red-shaded region). The policy with the maximum average reward at each  $L$  is selected as the optimal dynamics (blue triangles). (c) Same as (a) for  $L = 50$ . The learning curves appear noisier than in (a) but note that the vertical scale is much smaller. The learning rate is kept fixed throughout. (d) The distribution of  $r$  across parallel agents for  $L = 50$  is again much tighter than for  $L = 14$ .

We now provide more details on the training procedure used to obtain the policies of the main text. First, we outline the update step used to improve the policy and value function approximations. Second, we outline size-annealing, where we apply transfer learning by using systems of increasing size. Finally, we discuss policy evaluation and selection, whereby the best policy is chosen from a set of candidates.

(i) **Basic Outline of Training.** We start by initializing the parameters  $w_0$ ,  $\psi_0$  and  $\bar{r}_0$ , where  $\bar{r}_t$  is an estimate of the average reward per-time step,  $r(\pi_w)$ , after  $t$  training steps, along with the environment and initial state  $s_0$ . Choosing the three learning rates  $\alpha_\pi$ ,  $\alpha_v$  and  $\alpha_r$ , for each step  $t \in [0, T]$  we:

1. Sample an action  $a_t \sim \pi_w(\cdot|s_t)$  [where  $x \sim Y(\cdot)$  stands for  $x$  sampled from  $Y$ ], and from it get its log probability and eligibility,  $\ln \pi_w(a_t|s_t)$ ,  $\nabla_w \ln \pi_w(a_t|s_t)$ .
2. Get the next state and reward given the current state and action,  $(s_{t+1}, r_{t+1}) \sim P(\cdot, \cdot|a_t, s_t)$ .
3. Get the temporal difference error with the current value function,  $\delta_{t+1} = v_\psi(s_{t+1}) + r_{t+1} - \bar{r}_t - v_\psi(s_t)$ .
4. Update the parameters of the value function,  $\psi_{t+1} = \psi_t + \alpha_v \delta_{t+1} \nabla_\psi v_\psi(s_t)$ .
5. Update the parameters of the policy,  $w_{t+1} = w_t + \alpha_\pi \delta_{t+1} \nabla_w \ln \pi_w(a_t|s_t)$ .
6. Update the estimate of the average reward per time step,  $\bar{r}_{t+1} = \bar{r}_t + \alpha_r \delta_{t+1}$ .

(ii) **Annealing and Transfer Learning.** In the context of machine learning, “annealing” (sequentially

solving an optimisation problem reusing solutions to improve an initial guess) can be considered a form of *transfer learning*. In our case, we anneal the size of the system: the optimal policies for two system sizes will be similar as long as  $L' \gtrsim L$ , and in the settings considered the optimal dynamics should converge as  $L \rightarrow \infty$ .

We first approximate the optimal policy for a small system,  $L = 4$ , starting from random initial weights. The weights after optimisation at this size are then used as the initial weights for  $L = 6$ . This is repeated in steps of  $\Delta L = 2$ , up to the maximum desired  $L$ . This process ensures that effectively much longer training times are used for larger systems, and produces smooth convergence curves in  $L$ , which can be used both as diagnostic tools and for extrapolation [c.f. Fig. 3(c) inset].

(iii) **Policy Evaluation and Selection.** To determine the quality of a policy, we use it to generate trajectories without any change to the policy weights [c.f. Fig 1.(c) of the main text]. The set of rewards along these trajectories can then be averaged to estimate  $r(\pi_w)$  for the policy, allowing for different policies to be compared.

To ensure that we obtain the best policies possible, we then employ policy selection in two ways. Firstly, throughout training a given policy we store its weights periodically. After some number of periods, these weight snapshots are then evaluated and the best one is selected, ensuring that the policy can only improve with more training. Secondly, we run parallel policy optimisations and evaluations, starting from different random initial weights, with the best one selected.

The specific processes of policy evaluation and selection used to produce the results in the main text are illustrated for the ASEP in Fig. 4 (details in caption).

- 
- [1] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).
- [2] S. R. White, Phys. Rev. B **48**, 10345 (1993).
- [3] U. Schollwöck, Ann. Phys. **326**, 96 (2011).
- [4] J. Eisert, M. Cramer, and M. B. Plenio, Rev. Mod. Phys. **82**, 277 (2010).
- [5] J. Eisert, Modeling and Simulation **3**, 520 (2013), 1308.3318.
- [6] S. Montangero, *Introduction to Tensor Network Methods* (Springer International Publishing, 2018).
- [7] R. Orús, Nature Rev. Phys. **1**, 538 (2019), ISSN 2522-5820.
- [8] E. Stoudenmire and D. J. Schwab, in *Advances in Neural Information Processing Systems 29*, edited by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016), pp. 4799–4807, 1605.05775.
- [9] Y. Liu, X. Zhang, M. Lewenstein, and S.-J. Ran, arXiv:1803.09111 (2018).
- [10] Z.-Z. Sun, C. Peng, D. Liu, S.-J. Ran, and G. Su, arXiv:1903.10742 (2019).
- [11] C. Roberts, A. Milsted, M. Ganahl, A. Zalcman, B. Fontaine, Y. Zou, J. Hidary, G. Vidal, and S. Leichenauer, arXiv:1905.01330 (2019).
- [12] Z. Ghahramani and M. I. Jordan, Mach. Learn. **29**, 245 (1997), ISSN 1573-0565.
- [13] S. Efthymiou, J. Hidary, and S. Leichenauer, arXiv:1906.06329 (2019).
- [14] E. M. Stoudenmire, Quantum Sci. Technol. **3**, 034003 (2018).
- [15] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, Phys. Rev. X **8**, 031012 (2018).
- [16] Z.-F. Gao, S. Cheng, R.-Q. He, Z. Y. Xie, H.-H. Zhao, Z.-Y. Lu, and T. Xiang, arXiv:1904.06194 (2019).
- [17] Y. Levine, O. Sharir, N. Cohen, and A. Shashua, Phys. Rev. Lett. **122**, 065301 (2019).
- [18] C. Guo, Z. Jie, W. Lu, and D. Poletti, Phys. Rev. E **98**, 042114 (2018).
- [19] I. Glasser, N. Pancotti, and J. I. Cirac, arXiv:1806.05964 (2018).
- [20] I. Glasser, R. Sweke, N. Pancotti, J. Eisert, and J. I. Cirac, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. Alche-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), pp. 1496–1508.
- [21] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, 2018).
- [22] 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), ISSN 00280836.
- [23] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al., arXiv:1812.05905 (2018).
- [24] 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).
- [25] J. Schulman, S. Levine, P. Abbeel, M. I. Jordan, and P. Moritz, in *International Conference on Machine Learning* (2015).
- [26] T. Wang, X. Bao, I. Clavera, J. Hoang, Y. Wen, E. Langlois, S. Zhang, G. Zhang, P. Abbeel, and J. Ba, arXiv:1907.02057 (2019).
- [27] F. Metz and M. Bukov, arXiv:2201.11790 (2022).
- [28] A. Mahajan, M. Samvelyan, L. Mao, V. Makoviy-chuk, A. Garg, J. Kossaifi, S. Whiteson, Y. Zhu, and A. Anandkumar, arXiv:2106.00136 (2021).
- [29] H. Touchette, Physics Reports **478**, 1 (2009), ISSN 0370-1573, URL <https://www.sciencedirect.com/science/article/pii/S0370157309001410>.
- [30] J. P. Garrahan, J. Stat. Mech. **2016**, 073208 (2016).
- [31] J. P. Garrahan, Physica A **504**, 130 (2018).
- [32] R. L. Jack and P. Sollich, Eur. Phys. J. Spec. Top. **224**, 2351 (2015), ISSN 1951-6401.
- [33] R. L. Jack (2019), 1910.09883.
- [34] R. Chetrite and H. Touchette, Ann Henri Poincaré (2015).
- [35] C. Casert, T. Viejira, S. Whitelam, and I. Tamblyn, Phys. Rev. Lett. **127**, 120602 (2021), URL <https://link.aps.org/doi/10.1103/PhysRevLett.127.120602>.
- [36] S. N. Majumdar and H. Orland, J. Stat. Mech. **2015**, P06039 (2015).
- [37] R. Chetrite and H. Touchette, J. Stat. Mech. **2015**, P12001 (2015), ISSN 1742-5468.
- [38] P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, Annu. Rev. Phys. Chem. **53**, 291 (2002).
- [39] C. Giardinà, J. Kurchan, and L. Peliti, Phys. Rev. Lett. **96**, 120603 (2006).
- [40] C. Giardinà, J. Kurchan, V. Lecomte, and J. Tailleur, J. Stat. Phys. **145**, 787 (2011).
- [41] F. Cérou and A. Guyader, Stoch. Anal. Appl. **25**, 417 (2007).
- [42] V. Lecomte and J. Tailleur, J. Stat. Mech. **2007**, P03004 (2007), URL <http://stacks.iop.org/1742-5468/2007/i=03/a=P03004>.
- [43] S. Whitelam, D. Jacobson, and I. Tamblyn, J. Chem. Phys. **153**, 044113 (2020), URL <https://doi.org/10.1063/5.0015301>.
- [44] B. De Bruyne, S. N. Majumdar, and G. Schehr, Phys. Rev. E **104**, 024117 (2021), URL <https://link.aps.org/doi/10.1103/PhysRevE.104.024117>.
- [45] B. D. Bruyne, S. N. Majumdar, H. Orland, and G. Schehr, J. Stat. Mech. **2021**, 123204 (2021), URL <https://doi.org/10.1088/1742-5468/ac3e70>.
- [46] B. D. Bruyne, S. N. Majumdar, and G. Schehr, J. Phys. A **54**, 385004 (2021), URL <https://doi.org/10.1088/1751-8121/ac1d8e>.
- [47] B. De Bruyne, S. N. Majumdar, and G. Schehr, Phys. Rev. Lett. **128**, 200603 (2022), URL <https://link.aps.org/doi/10.1103/PhysRevLett.128.200603>.
- [48] G. Pozzoli and B. D. Bruyne, arXiv:2207.06865 (2022).
- [49] D. C. Rose, J. F. Mair, and J. P. Garrahan, New J. Phys. **23**, 013013 (2021), URL <https://doi.org/10.1088/1367-2630/abd7bd>.
- [50] A. Das and D. T. Limmer, J. Chem. Phys. **151**, 244123 (2019), URL <https://doi.org/10.1063/1.5128956>.
- [51] A. Das, D. C. Rose, J. P. Garrahan, and D. T. Limmer, J. Chem. Phys. **155**, 134105 (2021), URL <https://doi.org/10.1063/5.0057323>.
- [52] A. Das, B. Kuznets-Speck, and D. T. Limmer, Phys. Rev. Lett. **128**, 028005 (2022), URL <https://link.aps.org/doi/10.1103/PhysRevLett.128.028005>.
- [53] V. S. Borkar, S. Juneja, and A. A. Kherani, Commun. Inf. Syst. **3**, 259 (2003).
- [54] G. Ferré and H. Touchette, J. Stat. Phys. **172**, 1525 (2018).
- [55] J. Yan, H. Touchette, and G. M. Rotskoff, Phys. Rev. E

- 105**, 024115 (2022), URL <https://link.aps.org/doi/10.1103/PhysRevE.105.024115>.
- [56] J. Yan and G. M. Rotskoff, *J. Chem. Phys.* **157**, 074101 (2022), ISSN 1089-7690, URL <http://dx.doi.org/10.1063/5.0095593>.
- [57] T. Nemoto and S.-i. Sasa, *Phys. Rev. Lett.* **112**, 090602 (2014).
- [58] T. Nemoto, R. L. Jack, and V. Lecomte, *Phys. Rev. Lett.* **118**, 115702 (2017).
- [59] H. J. Kappen and H. C. Ruiz, *J. Stat. Phys.* **162**, 1244 (2016), ISSN 1572-9613.
- [60] P. G. Bolhuis, Z. F. Brotzakis, and B. G. Keller, arXiv:2207.04558 (2022).
- [61] L. Holdijk, Y. Du, F. Hooft, P. Jaini, B. Ensing, and M. Welling, arXiv:2207.02149 (2022).
- [62] Note1, by obeying detailed balance we specifically mean that a system has a dynamical generator that can be made Hermitian via a similarity transformation defined through the stationary state (i.e. equilibrium) distribution.
- [63] Note2, there are two situations where one might simply set  $\mathcal{C}(a, S) = 1$  for all state-action pairs: (i) when all actions are possible from any state; (ii) the form of the constraint is unknown and must be learnt, for example by penalising the policy when disallowed actions are selected.
- [64] M. C. Bañuls and J. P. Garrahan, *Phys. Rev. Lett.* **123**, 200601 (2019).
- [65] P. Helms, U. Ray, and G. K.-L. Chan, *Phys. Rev. E* **100**, 022101 (2019).
- [66] P. Helms and G. K.-L. Chan, *Phys. Rev. Lett.* **125**, 140601 (2020), URL <https://link.aps.org/doi/10.1103/PhysRevLett.125.140601>.
- [67] L. Causer, I. Lesanovsky, M. C. Bañuls, and J. P. Garrahan, *Phys. Rev. E* **102**, 052132 (2020), URL <https://link.aps.org/doi/10.1103/PhysRevE.102.052132>.
- [68] L. Causer, M. C. Bañuls, and J. P. Garrahan, *Phys. Rev. E* **103**, 062144 (2021), URL <https://link.aps.org/doi/10.1103/PhysRevE.103.062144>.
- [69] L. Causer, M. C. Bañuls, and J. P. Garrahan, *Phys. Rev. Lett.* **128**, 090605 (2022), URL <https://link.aps.org/doi/10.1103/PhysRevLett.128.090605>.
- [70] L. Causer, J. P. Garrahan, and A. Lamacraft, *Phys. Rev. E* **106**, 014128 (2022), URL <https://link.aps.org/doi/10.1103/PhysRevE.106.014128>.
- [71] M. Fishman, S. White, and E. Stoudenmire, *SciPost Physics Codebases* (2022), URL <http://dx.doi.org/10.21468/scipostphyscodeb.4>.
- [72] DeepMind, *Using jax to accelerate our research* (2022), URL <https://www.deepmind.com/blog/using-jax-to-accelerate-our-research>.
- [73] J. K. Terry, B. Black, N. Grammel, M. Jayakumar, A. Hari, R. Sullivan, L. Santos, R. Perez, C. Horsch, C. Dieffendahl, et al., arXiv:2009.14471 (2020).
- [74] G. Van Rossum and F. L. Drake, *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009), ISBN 1441412697.
- [75] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, et al., *JAX: composable transformations of Python+NumPy programs* (2018), URL <http://github.com/google/jax>.
- [76] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, et al., *Nature* **585**, 357 (2020), URL <https://doi.org/10.1038/s41586-020-2649-2>.
- [77] The pandas development team (2022), URL <https://doi.org/10.5281/zenodo.7093122>.
- [78] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [79] A. Collette, *Python and HDF5* (O'Reilly, 2013).
- [80] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, *SIAM Rev. Soc. Ind. Appl. Math.* **59**, 65 (2017), URL <https://doi.org/10.1137/141000671>.
- [81] *HDF5.jl* (2022), URL <https://github.com/JuliaIO/HDF5.jl>.
- [82] J. Jäckle and S. Eisinger, *Zeitschrift für Physik B Condensed Matter* **84**, 115 (1991), ISSN 1431-584X, URL <https://doi.org/10.1007/BF01453764>.
- [83] F. Spitzer, *Advances in Mathematics* **5**, 246 (1970), ISSN 0001-8708, URL <https://www.sciencedirect.com/science/article/pii/0001870870900344>.
- [84] R. L. Jack and P. Sollich, *Prog. Theor. Phys. Supp.* **184**, 304 (2010), ISSN 0375-9687, URL <https://doi.org/10.1143/PTPS.184.304>.
- [85] E. Gillman and D. C. Rose, *Acten, actor critic with tensor networks*, [https://github.com/RL-with-TNs/acten\\_code](https://github.com/RL-with-TNs/acten_code) (2022).

## Supplemental Material

### FORWARD PASSES FOR THE EAST MODEL AND SIMPLE EXCLUSION PROCESS

In this section we discuss the so-called forward passes required for the discovery of optimal dynamics using the function approximations described in the main text. Along with the information here, example scripts that implement the forward passes can be found on the associated GitHub repository, which should be referred to for more details [85]. Note that, while we focus here on the case of kinetically constrained spin systems, specifically the east model and asymmetric simple exclusion process (ASEP), much of the following is generic and can easily be adapted to a variety of other applications.

Generally, to solve policy optimisation problems with actor-critic (AC), for a given state  $s$  we must be able to:

1. Evaluate the function approximation for the state-value function,  $v_\psi(S)$ .
2. Sample an action,  $a$ , from the function approximation for the policy,  $\pi_w(a|S)$ .
3. Calculate  $\log[\pi_w(a|S)]$  for that action.



The computations that implement these are known as the forward passes, while the so-called backward passes implement the gradients of these quantities with respect to the parameters (weights). Since gradients can be calculated automatically in programming frameworks such as JAX [75], we are only required to implement the necessary forward passes explicitly for a working AC method.

### Evaluation of $v_\psi(S)$

Recall that in the main text we defined,

$$v_\psi(S) = \log [\varphi(s)^2] , \quad (\text{S1})$$

where

$$\varphi(s) = \text{Tr} [A_{s_1} A_{s_2} \dots A_{s_L}] . \quad (\text{S2})$$

Since the form of  $v_\psi(S)$  is independent of the dynamics in question, we start by describing its implementation here, before moving onto the specific implementation of the policy forward passes for each model in turn below. To ensure the implementation presented is readily applicable to automatic differentiation in frameworks such as JAX [75], we describe it here in terms of standard functions (specifically “scan” functions) that are amenable to automatic differentiation.

For a set of steps,  $k = 1, 2, \dots, T$ , a scan-function is defined by the repeated application of a given function,  $f$ , to a set of inputs,  $x_k$ , and a single “carry” object,  $C$ , that is updated at each step. That is, for each  $k$  the scan-function computes,

$$C, y_k = f(C, x_k) , \quad (\text{S3})$$

where  $y_k$  is an output at each step which is not carried forward to the remaining steps, but instead is returned from the overall scan function as an array of the  $y_k$  values from each step. We will not make use of this output, instead using the final carry as the result of the scan, and as such will ignore it.

To implement the forward-pass for  $\varphi(s)$  in Eq. (S2), the inputs  $x_k$  are vectors defined as  $x_k = (1 - s_k, s_k)$  with  $k = 1, 2, \dots, L$  i.e.  $L = T$ . The carry is initiated as the identity matrix,  $C = \mathcal{I}$ , with shape  $(\chi, \chi)$ . The carry output of  $f$  is then defined in terms of the tensor components  $A_{s,\alpha,\beta} = [A_s]_{\alpha,\beta}$  as,

$$C_{\alpha,\gamma} = \sum_{\beta=0}^{\chi-1} \sum_{n=0}^1 C_{\alpha,\beta} A_{n,\beta,\gamma} [x_k]_n . \quad (\text{S4})$$

The additional output  $y_k$  is not required and can be discarded. Collecting the inputs and outputs as vectors,  $x$  and  $y$ , then the value of  $\varphi(s)$  is then given by applying the scan-function for this  $f$ , followed by a trace over the carry,

$$\begin{aligned} C, y &= \text{scan}_f(C, x) \\ \varphi(s) &= \text{Tr} [C] . \end{aligned} \quad (\text{S5})$$

Implemented in this manner, e.g. as found in [85], the gradient of the value-function can be easily evaluated using auto-differentiation.

### Policy Function Approximations For Local Kinetic Constraints

Before turning to details of implementing the forward passes for  $\pi_w(a|S)$  in the east model and ASEP, we briefly discuss here how the constraints on the dynamics can be captured explicitly in the structure of  $\pi_w(a|S)$  via the constraint function  $\mathcal{C}(a, s)$ . To this end, recall first that in the main text we defined the policy function approximation as,

$$\pi_w(a|S) = \mathcal{C}(a, s) \frac{\varphi(a, s)^2}{\mathcal{N}(s)} , \quad (\text{S6})$$

where

$$\varphi(a, s) = \text{Tr} [A_{a_1 s_1} A_{a_2 s_2} \dots A_{a_L s_L}] . \quad (\text{S7})$$

Here,  $\mathcal{N}(s) = \sum_a \mathcal{C}(a,s)=1 \varphi(a,s)^2 = \sum_a \mathcal{C}(a,s) \varphi(a,s)^2$  and  $\mathcal{C}(a,s)$  is the constraint function, which returns one if an action  $a$  is possible given a state  $s$ , or zero otherwise.

The constraint function,  $\mathcal{C}(a,s)$ , allows us to include explicit constraints on the actions selected by our function approximations. This is particularly powerful when modelling the dynamics of spin systems whose constraints are both known and such that only a few states can be reached from any other. In that case we can construct  $\mathcal{C}(a,s)$  so that our policy reflects these constraints exactly, whereas in other scenarios this must be approximated or learnt.

**Single Spin-Flip Constraint:** In the dynamics studied in the main-text, we consider two varieties of constraint. The first, which pertains to both models, requires that at most a single spin is allowed to flip at a given time. For convenience, we will define the set of actions  $\tilde{a}_k$  for  $k \in [1, L+1]$  that represent a flip at site  $k$  when  $k \leq L$  and no flip at any site when  $k = L+1$  (in which case the state is unchanged by the action). Taking  $L = 4$  for illustration, in terms of the variables in the main text where  $a = (a_1, a_2, a_3, a_4)$  with each  $a_k$  indicating a flip at site  $k$ , we have  $\tilde{a}_1 = (1, 0, 0, 0)$ ,  $\tilde{a}_2 = (0, 1, 0, 0)$ ,  $\tilde{a}_3 = (0, 0, 1, 0)$ ,  $\tilde{a}_4 = (0, 0, 0, 1)$ , and  $\tilde{a}_5 = (0, 0, 0, 0)$ . This constraint can be included in  $\pi_w(a|S)$  straightforwardly via the choice,

$$\mathcal{C}(a,s) = \sum_{k=1}^{L+1} \delta_{\tilde{a}_k, a}. \quad (\text{S8})$$

With the choice of constraint function (S8), sampling  $\pi_w(a|S)$  will select actions only from the  $L+1$  possibilities  $\tilde{a}_k$ , with other actions having strictly zero probability of occurring. As such, sampling can equivalently be performed by selecting an action from the probabilities  $\{\pi_w(\tilde{a}_k|s)\}_{k=1}^{L+1}$  alone. This simplifies the problem of sampling  $\pi_w(a|S)$ , because the normalisation factor  $\mathcal{N}(s)$  [c.f. (S6)] –which in general is hard to calculate for conditional probabilities– can be calculated explicitly by enumerating  $\mathcal{C}(\tilde{a}_k, s) \varphi(\tilde{a}_k, s)^2$  for all  $k = 1, 2, \dots, L+1$ . Thus, at most  $L+1$  computations are required to sample an action from a policy with this constraint. While, in principle, these could be computed in parallel, for the problems here we present an alternative method (as applied in the main text) where the policy is instead sampled via a “sweep”, similar to that performed for more standard tensor network algorithms.

**Local Kinetic Constraint:** The second aspect of the constraints is the local kinetic constraint. Here, whether a spin at site  $k = 1, 2, \dots, L$  can flip depends only on the states of the neighbouring sites at  $k-1, k$  and  $k+1$ . For example, in the case where the possibility of  $\tilde{a}_k$  depends on a three-site neighbourhood we can further write that,

$$\mathcal{C}(a,s) = \sum_{k=1}^L \delta_{\tilde{a}_k, a} \mathcal{C}(\tilde{a}_k, s_{k-1}, s_k, s_{k+1}) + \delta_{\tilde{a}_{L+1}, a} \mathcal{C}(\tilde{a}_{L+1}, s). \quad (\text{S9})$$

Note that here we have separated out the “no-flip” action,  $\tilde{a}_{L+1}$ , as this must typically be treated separately in a given problem.

While both the east model and ASEP are subject to local kinetic constraints, the specific form of the local constraint function,  $\mathcal{C}(\tilde{a}_k, s_{k-1}, s_k, s_{k+1})$ , will depend on the model in hand. As such, the function approximations for the policies will differ slightly and, therefore, so will the implementation of the forward passes.

### Forward Pass for $\pi_w(a|S)$ in the East Model

We now describe the implementation of the forward pass for  $\pi_w(a|S)$  in the east model, see [85] for an explicit example of implementation. In the east model, a spin can only flip if the spin to its left is up. As such, this local kinetic constraint can be captured by the local constraint function,

$$\mathcal{C}(\tilde{a}_k, s) = \mathcal{C}(\tilde{a}_k, s_{k-1}, s_k, s_{k+1}) = s_{k-1} \text{ for } k \in [1, L]. \quad (\text{S10})$$

For the case of no-flips,  $\tilde{a}_{L+1}$ , we take this to be always possible unless every spin is up, i.e.,

$$\mathcal{C}(\tilde{a}_{L+1}, s) = 1 - \delta_{s, (1, 1, 1, \dots, 1)}. \quad (\text{S11})$$

Due to the constraint Eq. (S8), we need consider only the possible actions  $\tilde{a}_k$ . For these actions, the matrix product operator used in the function approximation for the policy (S6) takes the form,

$$\varphi(\tilde{a}_k, s) = \text{Tr} \left[ \left( \prod_{l=1}^{k-1} A_{0s_l} \right) A_{1s_k} \left( \prod_{l=k+1}^L A_{0s_l} \right) \right]. \quad (\text{S12})$$

We then define the product of matrices  $A_{0s_m}$  to the left and right of some site  $m$  as,

$$M_m^{\text{left}} = \prod_{l=1}^{m-1} A_{0s_l} , \quad (\text{S13})$$

and

$$M_m^{\text{right}} = \prod_{l=m+1}^L A_{0s_l} . \quad (\text{S14})$$

These matrices can be constructed iteratively as,

$$M_{m+1}^{\text{left}} = M_m^{\text{left}} A_{0s_m} , \quad (\text{S15})$$

and

$$M_{m-1}^{\text{right}} = A_{0s_m} M_m^{\text{right}} . \quad (\text{S16})$$

To relate these to the probabilities of taking a given action, we then define the left-environment,

$$\mathcal{E}_m^{\text{left}} = \mathcal{C}(a, s) M_m^{\text{left}} A_{1s_m} , \quad (\text{S17})$$

which includes the site  $m$  and the constraint function,  $\mathcal{C}(a, s)$ . With this,

$$\varphi(\tilde{a}_k, s) = \text{Tr} \left[ \mathcal{E}_k^{\text{left}} M_k^{\text{right}} \right] . \quad (\text{S18})$$

The iterative form of Eq. (S15) and the expression Eq. (S18) shows that we can construct all required  $\varphi(\tilde{a}_k, s)$  iteratively by “sweeping” from left to right (obtaining  $\mathcal{E}_k^{\text{left}}$  for  $k = 1, 2, 3, \dots, L$ ) and then from right to left (obtaining  $M_k^{\text{right}}$  for  $k = L, L-1, \dots, 1$ ). These sweeps can be implemented efficiently with the use of a scan function, just as with  $v_\psi(S)$  [c.f. Eq. (S5)]. However, while for  $v_\psi(S)$  just one sweep to the right was required (i.e. a single scan function), here an additional sweep to the left is required.

Specifically, the rightward sweep is implemented by a scan of the “right-step” function  $f_R$ , such that  $C, y_k = f_R(C, x_k)$ . The carry in this case consists of two objects,  $C = \{\mathcal{E}_L, M\}$ . The first is the set of left environments,  $\mathcal{E}_L$ , whose components are the left-environments for a specific site  $m$ , i.e.  $[\mathcal{E}^{\text{left}}]_m = \mathcal{E}_m^{\text{left}}$ . The second object is the matrix  $M$ , which at the start of any step  $k$  is simply equal to  $M_k^{\text{left}}$ . At step  $k$ , these are updated as,

$$[\mathcal{E}_L]_k \rightarrow \mathcal{C}(\tilde{a}_k, s) M A_{1s_k} , \quad (\text{S19})$$

$$M \rightarrow M A_{0s_k} . \quad (\text{S20})$$

As with  $v_\psi(S)$ , the output  $y_k$  of  $f_R$  is not needed and can be discarded. Scanning  $f_R$  then gives,

$$\{\mathcal{E}_L, M\}, y = \text{scan}_{f_R}(\{\mathcal{E}_L, M\}, x) , \quad (\text{S21})$$

where once again the input is a set of vectors,  $x = \{(1 - s_l, s_l)\}_{l=1}^L$ .

According to Eq. (S18), to complete the computations of  $\varphi(\tilde{a}_k, s)$  we must then compute  $M_k^{\text{right}} \forall k \in [1, L]$ . From the iterative expression Eq. (S16) it is clear how to achieve this, see [85] for an example implementation.

### Forward Pass for $\pi_w(a|S)$ in the ASEP

In the ASEP we consider, the local kinetic constraint is such that particles (i.e. up spins) can move to the left/right only if there is an unoccupied space (i.e. a down spin) in that position. Moreover, analogous to the east model where at most only a single spin could flip per time-step, for the ASEP at most a single particle can move left or right. These constraints can be realised in  $\pi_w(a|S)$  by setting the constraint function  $\mathcal{C}(a, s)$  such that the only possible actions are the  $L$  actions that flip pairs of spins,  $\tilde{a}_k : (s_k, s_{k+1}) \rightarrow (1 - s_k, 1 - s_{k+1})$ , and the no-flip action  $\tilde{a}_{L+1}$ . Taking

$L = 4$  for illustration, in terms of the fundamental single-site spin flip actions:  $\tilde{a}_1 = (1, 1, 0, 0)$ ,  $\tilde{a}_2 = (0, 1, 1, 0)$ ,  $\tilde{a}_3 = (0, 0, 1, 1)$ ,  $\tilde{a}_4 = (1, 0, 0, 1)$  and  $\tilde{a}_5 = (0, 0, 0, 0)$ . The constraint function is then,

$$\mathcal{C}(a, s) = \sum_{k=1}^L \delta_{a, \tilde{a}_k} \mathcal{C}(\tilde{a}_k, s) + \delta_{a, \tilde{a}_{L+1}} \mathcal{C}(\tilde{a}_{L+1}, s), \quad (\text{S22})$$

$$= \sum_{k=1}^L \delta_{a, \tilde{a}_k} \mathcal{C}(\tilde{a}_k, s_k, s_{k+1}) + \delta_{a, \tilde{a}_{L+1}} \mathcal{C}(\tilde{a}_{L+1}, s), \quad (\text{S23})$$

$$= \sum_{k=1}^L \delta_{a, \tilde{a}_k} [(1 - s_k)s_{k+1} + s_k(1 - s_{k+1})] + \delta_{a, \tilde{a}_{L+1}} \mathcal{C}(\tilde{a}_{L+1}, s). \quad (\text{S24})$$

Here, in the third line, the choice of local constraint function,  $\mathcal{C}(\tilde{a}_k, s_k, s_{k+1})$ , ensures action  $\tilde{a}_k$  only occurs when there is a single particle in the pair of sites to be flipped.

The special case of no-flip, realised in  $\mathcal{C}(\tilde{a}_{L+1}, s)$ , can always occur except when no particles have a neighbour (e.g.  $(0, 0, 1, 0)$  for  $L = 4$ ). This can be expressed as,

$$\mathcal{C}(\tilde{a}_{L+1}, s) = 1 - \mathbb{I} \left( \sum_{k=1}^L s_k s_{k+1} = 0 \right), \quad (\text{S25})$$

where we have introduced the so-called indicator function,  $\mathbb{I}$ , which returns 1 or 0 if the condition in its argument is true or false respectively.

The expression for  $\varphi(\tilde{a}_k | s_t)$  in the case of the ASEP is,

$$\varphi(\tilde{a}_k, s) = \text{Tr} \left[ \left( \prod_{l=1}^{k-1} A_{0s_l} \right) A_{1s_k} A_{1s_{k+1}} \left( \prod_{l=k+2}^L A_{0s_l} \right) \right]. \quad (\text{S26})$$

Using the previous definitions for the left-environments,  $\mathcal{E}_k^{\text{left}}$  and the matrices  $M_k^{\text{right}}$ , this is equivalent to

$$\varphi(\tilde{a}_k, s) = \text{Tr} \left[ \mathcal{E}_k^{\text{left}} A_{1s_{k+1}} M_{k+1}^{\text{right}} \right]. \quad (\text{S27})$$

Expressed in this manner, it is clear that to implement the forward pass for  $\varphi(\tilde{a}_k, s)$  we can again apply a scan function, just as with the east model. Indeed, the details of this procedure are very similar to those of the east model outlined previously, although slightly more involved due to the fact the  $L + 1$  potential actions  $\tilde{a}_k$  change two sites rather than one. For further details, a full example implementation is given in [85].