

OPEN

Prospects of reinforcement learning for the simultaneous damping of many mechanical modes

Christian Sommer^{1*}, Muhammad Asjad¹ & Claudiu Genes^{1,2}

We apply adaptive feedback for the partial refrigeration of a mechanical resonator, i.e. with the aim to simultaneously cool the classical thermal motion of more than one vibrational degree of freedom. The feedback is obtained from a neural network parametrized policy trained via a reinforcement learning strategy to choose the correct sequence of actions from a finite set in order to simultaneously reduce the energy of many modes of vibration. The actions are realized either as optical modulations of the spring constants in the so-called quadratic optomechanical coupling regime or as radiation pressure induced momentum kicks in the linear coupling regime. As a proof of principle we numerically illustrate efficient simultaneous cooling of four independent modes with an overall strong reduction of the total system temperature.

The radiation pressure effect of light onto the motion of mechanical resonators has been extensively employed to bring such macroscopic systems towards the quantum ground state^{1–11}. In a standard approach, the aim is to isolate a single vibrational mode and bring it to a state where the only relevant motion is given by the zero-point fluctuations. Cold-damping is one of the used techniques, where one detects motionally-induced phase changes in the cavity output and an electronic feedback loop is implemented to dynamically modify the cavity drive such as to produce an extra optical damping effect^{12–19}. Alternatively, in the good cavity limit where the photon loss rate is smaller than the mechanical frequency, the resolved sideband technique can be implemented by detuning the drive to the cooling sideband^{20–24}. As the effect stems from the inherent time delay between the action of the mechanical resonator onto the cavity field and the back-action of light, this can be seen as a sort of automatic cavity induced feedback. Both techniques are devised and have been successfully applied for single vibrational mode cooling. However, it is interesting to devise an alternative technique that can induce partial to full refrigeration of the mechanical resonator, i.e. to simultaneously cool a multitude of vibrational modes into which the thermal energy is distributed. An impediment is that the detected output signal only gives information on a generalized collective quadrature but not on all modes. This leads to efficient cooling of some collective mode (for example center of mass) while some collective modes become dark and remain in a high temperature state. It has been recently pointed out that some strategies such as multimode cold-damping could in principle lead to sympathetic cooling of many modes via disorder induced coupling between bright and dark modes²⁵.

Here, we propose a machine learning approach towards devising a strategy capable of providing refrigeration of the classical motion of a mechanical resonator based on the feedback obtained from the detection of a single optical mode. While the detected optical mode only gives information on a collective generalized quadrature obtained as a linear combination of individual mode displacements, the procedure is optimized such as at any instant in time a compromise is made between efficiently cooling a particular target mode while not affecting the others too much. We provide proof-of-principle multi-mode numerical simulations using a neural network parametrized policy trained by a reinforcement learning algorithm to generate the feedback signal capable of simultaneously extracting thermal energy from four distinct modes of a single mechanical resonator.

Machine learning techniques have been recently applied to various applications in quantum physics ranging from the identification of phases in many-body systems, predicting ground-state energies for electrostatic potentials, active learning approaches to propose and optimize experimental setup configurations and towards applications for quantum control and quantum-error correction^{26–34}. In particular, a few studies^{26,34,35} successfully applied the technique of reinforcement learning with neural networks³⁶. This approach originates from the idea,

¹Max Planck Institute for the Science of Light, Staudtstraße 2, D-91058, Erlangen, Germany. ²Department of Physics, University of Erlangen-Nuremberg, Staudtstraße 2, D-91058, Erlangen, Germany. *email: christian.sommer@mpl.mpg.de

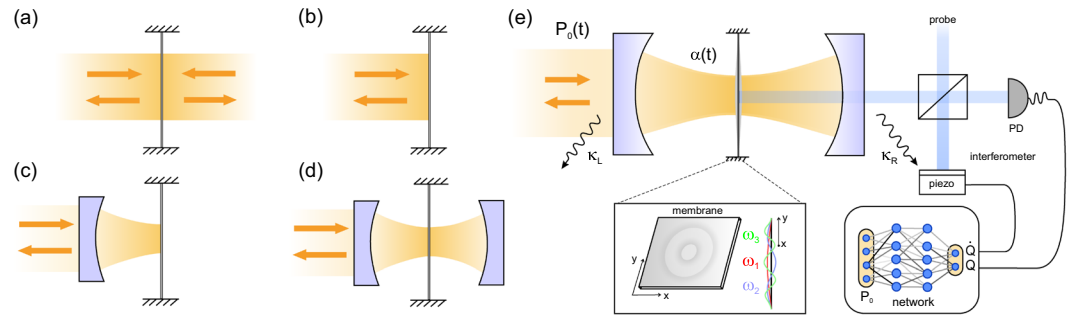


Figure 1. Cooling of thermal motion via two-sided kicking in (a) or via one-sided-kicking in (b). (c) Increased damping efficiency can be realized by an amplification of the photon-phonon coupling in the linear coupling regime via an optical cavity enhancement of the electric field amplitude. (d) Membrane-in-the-middle configuration leads to a quadratic coupling in the mechanical displacement allowing optical control of the mechanical mode's spring constant. (e) Simultaneous cooling of multiple oscillating modes (inset shows a few drum modes of a dielectric membrane) using feedback generated by a reinforcement trained policy, encoded in a neural network. While the illustration shows a quadratic membrane-in-the-middle setup, the validity extends to the end-mirror linear setup as well. The outgoing signal from a driven cavity carries information on the collective displacement of all membrane vibration modes. This signal is fed through a neural network and the network's suggested action is implemented as a modulation of the cavity input drive amplitude.

to let an intelligent agent that observes its environment choose an action, that is determined by a given policy trying to optimize a particular reward and/or minimize a punishment.

We employ such a reinforcement learning technique for optically assisted cooling of the classical thermal state of a multi-mode mechanical resonator system^{37–39}. The learning technique allows one to acquire a nonlinear function that chooses a feedback action that will be applied on the dynamical system upon taking the full or partial measured state of the system as an input. The training of this function that is given by a dense neural network is obtained by trial and error and quantified by an increased reward that is obtained by successfully reducing the energy of the resonators.

The physical systems considered are depicted in Fig. 1. The mechanical resonator is subject to environmental noise described by a standard Brownian motion stochastic force leading to thermalization at some equilibrium temperature T . The feedback action is implemented via the radiation pressure force, i.e. photon kicks either from one or two sides. The induced damping is straightforward in the two-sided kicking case [illustrated in Fig. 1a]: the read-out of motion is followed by appropriate kicking action from the side towards which the resonator is moving. However, one-sided kicking [illustrated in Fig. 1b] already suffices allowing setups such as the cavity optomechanical platform pictured in Fig. 1c. The typical weak free space photon-phonon interaction can also be drastically increased by the filtering of the action through the high-finesse optical cavity. Such a situation is characterized by a linear coupling of the photon number to the membrane's displacement and has been extensively studied in single mode cooling via cavity time delayed effects¹² or by implementation of cold damping techniques¹² especially in the bad cavity regime. The membrane-in-the-middle^{40–42} scenario in Fig. 1d,e corresponds to a quadratic coupling in displacement leading to the possibility of optically modulating the mechanical oscillation frequency⁴². We describe in Fig. 1e a possible approach for feedback cooling via cavity field detection and neural network assisted feedback.

We will consider the bad cavity case where losses are large compared to the mechanical resonator's vibration frequencies such that the cavity back-action is negligible. In such a case, the situations described in Fig. 1b,c are physically equivalent with the difference that in Fig. 1c the action of a single photon is multiplied by a large number roughly proportional to the finesse of the cavity. We also distinguish between a parametric regime with quadratic coupling implemented in the membrane-in-the-middle setup and the linear coupling regime realizable with a single-end mirror cavity or in free space. First we analyze the performance of a neural network suggested set of actions onto the cooling of a single mode via parametric modulation of the oscillation frequency: we describe the shape of the action and numerically show the efficient reduction of energy from the initial thermal distribution. We then apply the technique to the linear cooling of four distinct modes of the resonator and find a more complex set of actions required for efficient simultaneous cooling of all four modes (with limitations arising due to the numerical complexity of the simulations).

Model

We consider a membrane resonator with a few modes of oscillations of frequencies ω_j (where $j = 1, \dots, N$). We start with a quantum formulation of the system's dynamics aimed at future treatments of cooling in the presence of quantum noise. However the current formulation aims only at the reduction of classical thermal noise and is therefore obtained by inferring the equivalent classical stochastic equations of motion. The Hamiltonian for the collection of modes is written as $H_m = \sum_{j=1}^N \hbar \omega_j / 2 (p_j^2 + q_j^2)$, in terms of dimensionless position and momentum quadratures q_j and p_j for each independent membrane oscillation mode. The effect of the thermal reservoir can be easily included in a set of equations of motion supplemented with the proper input stochastic noise terms:

$$\dot{q}_j = \omega_j p_j, \quad (1a)$$

$$\dot{p}_j = -\omega_j q_j - \gamma_j p_j + \xi_j + F_j(t). \quad (1b)$$

The parameter γ_j describes the damping of the j 's resonator mode. Its associated zero-averaged Gaussian stochastic noise term leading to thermalization with the environment can be fully described by the two-time correlation function:

$$\langle \xi_j(t) \xi_j(t') \rangle = \frac{\gamma_j}{\omega_j} \int_0^\Omega \frac{d\omega}{2\pi} e^{-i\omega(t-t')} S_{th}(\omega) \delta_{jj}, \quad (2)$$

where Ω is the frequency cutoff of the reservoir and the thermal noise spectrum is given by $S_{th}(\omega) = \omega [\coth(\hbar\omega/2k_B T) + 1]$. For sufficiently high temperatures $k_B T \gg \hbar\omega_j$, the correlation function becomes a standard white noise input with delta correlations both in frequency and time. Specifically, one can approximate $\langle \xi_j(t) \xi_j(t') \rangle \approx (2\bar{n}_j + 1) \gamma_j \delta(t - t') \delta_{jj}$, where the occupancy of each vibrational mode is given by $\bar{n}_j = (\exp(\hbar\omega_j/k_B T) - 1)^{-1} \approx k_B T / \hbar\omega_j$. For numerical simulations we generate a stochastic input noise as a delta-correlated Wiener increment with variance proportional to the integration time-step (see Methods) and follow an approach described in ref.⁴³. For consistency we check (in the Methods) that the thermal bath indeed correctly describes the expected thermalization of an initially cold oscillator towards the equilibrium temperature T at a rate given by γ .

The momentum kicks selected by the network are encompassed in the action of the force terms $F_j(t)$. This can be realized for example by the radiation pressure effect of a laser beam, modulated by a device like an AOM (acousto-optic modulator). Here, forces acting on different resonators given by F_j and $F_{j'}$ for $j \neq j'$ differ only by a constant multiplication factor as they are all obtained from the same quantity (the output field).

To amplify the effect of the action force onto the mechanical resonator one can utilize optical cavities. A cavity also allows control over the coupling by placing the membrane either in a node (quadratic coupling) or anti-node (linear coupling) of the cavity mode. The Hamiltonian is now modified by the addition of the free cavity mode $\hbar\omega_c a^\dagger a$, laser driving resonant to the cavity mode $i\hbar\mathcal{E}(t)(a^\dagger - a)$ (in a frame rotating at ω_c) and optomechanical interaction of linear $\sum_j \hbar g_j^{(1)} a^\dagger a q_j$ or quadratic form $\sum_j \hbar g_j^{(2)} a^\dagger a q_j^2$. The amplification effect of the light field amplitude can be seen from the relation $\mathcal{E}(t) = \sqrt{2\mathcal{P}_0(t)\kappa_L/\hbar\omega_c}$ connecting the driving amplitude to the input laser power $\mathcal{P}_0(t)$ through the left mirror with losses at rate κ_L . For high-finesse cavities photons perform many round trips before leaking out through the mirrors resulting in a large momentum transfer onto the mirror: this can be seen by taking the limit of small κ_L resulting in a large value of $a(t)$ for a given $\mathcal{P}_0(t)$. Notice that we considered a double-sided cavity with left κ_L and right κ_R decay rates adding to the total loss rate $\kappa = \kappa_L + \kappa_R$. The coefficients $g_j^{(1)}$ and $g_j^{(2)}$ are the linear and quadratic per photon optomechanical coupling rates corresponding to the two situations depicted in Fig. 1c,d, respectively. While the cavity field amplitude inherently depends on the displacement of the mechanical mode, we will assume the unresolved sideband regime where this dependence is weak. Moreover, we are interested in the classical problem i.e. in simulating the proper set of actions that results in the shrinking of an initial large thermal distribution for the total energy of the oscillator. To this end we only consider the trivial dynamics of the cavity field classical amplitude $\alpha(t) = \langle a(t) \rangle$ which follows the driving field as $\dot{\alpha}(t) = -\kappa\alpha + \mathcal{E}(t)$. We can then reduce the dynamics of the system to

$$\dot{q}_j = \omega_j p_j, \quad (3a)$$

$$\dot{p}_j = -\omega_j q_j - \gamma_j p_j + \xi_j - g_j^{(1)} |\alpha(t)|^2, \quad (3b)$$

$$\dot{\alpha} = -\kappa\alpha + \mathcal{E}(t), \quad (3c)$$

which resemble Eqs. 1a,b, where we can identify the action forces $F_j(t) = -g_j^{(1)} |\alpha(t)|^2$ (the cavity field $\alpha(t)$ playing the role of the action delivering the cooling momentum kicks to the mechanical oscillators). As noted before, as the actions are obtained from the same cavity field intensity, they only differ by the multiplicative $g_j^{(1)}$ factor. Notice also that this configuration strongly resembles a cold damping approach¹²⁻¹⁴.

In contrast, for a quadratic coupling Hamiltonian, the changes in the momentum are of a very different nature

$$\dot{p}_j = -[\omega_j + 2g_j^{(2)} |\alpha(t)|^2] q_j - \gamma_j p_j + \xi_j, \quad (4)$$

as the cavity periodically modulates the oscillation frequencies of each mode.

To provide the neural network feedback onto the motional dynamics, we use the inferred q and \dot{q} at a given time $t - \Delta t$ as input values for the neural network [see Fig. 1]. The trained network then selects the appropriate action by choosing the value of $\mathcal{E}(t)$ (from a finite number of possible values) to be acted upon the system. The size of the time-step Δt is chosen such that $1/\omega_j \gg \Delta t$ for all j to minimize the error in the numerical integration. For a given drive amplitude, the set of actions on the different modes will be different according to the values of the optomechanical couplings (as they are proportional to $g_j^{(1)} |\alpha(t)|^2$ or $2g_j^{(2)} |\alpha(t)|^2$). We then use the Runge-Kutta

fourth-order method (RK4) for the numerical integration of the dynamical system where we iteratively sum for each time step. Additionally, at each time step we inject the measured parameters of the dynamical system as input data into the nonlinear function formed by the neural network to predict on the action and thereby the momentum kick or cavity field strength suitable for the next time step, which is acquired from the output nodes (neurons) of the network.

Reinforcement Learning

The neural network provides a nonlinear function, that for some given input data, which harbor information about the oscillator states at a given time step t , predicts the correct action for the next time step $t + \Delta t$ that helps to reduce the overall energy of the dynamical system at later times. This function forms the neural network parameterized policy π . To obtain an optimal (or nearly optimal) policy we employ the technique of reinforcement learning^{36,44} and in particular a policy gradient approach⁴⁵. Such a problem is in general referred to as a Markov decision process (MDP)⁴⁶ and described in detail in the Methods section. Here, the network acts as an agent that by observing parameters of the environment (resonator) improves its probabilistic policy that chooses the right actions a_t at a given time t to increase an overall reward $R = \sum_t R_t$ (full reward over a trajectory) that is connected to the reduction of the energy of the resonator modes. The actions are chosen from a finite set (of values of different amplitudes) and realized as momentum kicks or translated into frequency shifts. As an input to the network we feed information about the state of the environment given by $s_t = (q(t), \dot{q}(t))$. The network outputs the probabilities $\pi_\theta(a_t|s_t)$ for the actions a_t that could be applied to the dynamical system. Here, the parameter θ encompasses all the weights and biases of the network. We take the action with the highest probability and apply it in the next iteration of Eq. 1 up to Eq. 3b. The probabilities $\pi_\theta(a_t|s_t)$ can be optimized with respect to an increased reward return R_t by employing an update rule for the weights and biases of the neural network, following $\theta \leftarrow \theta + \Delta\theta$ and

$$\Delta\theta_j = \eta \partial_{\theta_j} \mathbb{E}[R] = \eta \sum_t \mathbb{E}[(R - b) \partial_{\theta_j} (\ln \pi_\theta(a_t|s_t))], \quad (5)$$

where \mathbb{E} is the expectation value over all state and action sequences (full trajectories), which here is approximated by averaging over a large enough set of oscillator trajectories (training batch) and their corresponding action sequences which we have obtained from the iterative summation of the dynamical equations (RK4) and from the predictions of the neural network at each time step for various randomly chosen initial conditions. The learning rate is given by the parameter η and b is a baseline to suppress fluctuations of the reward gradient^{45,47}. Here, the baseline is approximated by $b \approx b_n = (1/n - 1) \sum_{i=1}^{n-1} \bar{R}^{(i)}$, where $\bar{R}^{(i)}$ is the average total Reward from the i 's learning epoch. Here, the training epoch is defined as the number of updates $\theta \leftarrow \theta + \Delta\theta$.

The neural network which is represented by the array θ encompassing all weights and biases, consists of an input and output layer and two hidden layers whereby the number of input neurons depends on the number measured of variables of the system while the number of output neurons depends on the number of possible output actions, respectively (see Methods). The two hidden layers consists of up to 60 to 100 neurons each. The network is densely connected and we chose “relu” (rectified linear unit) as a nonlinear function acting on each neuron in the two hidden layers. The probabilities for each action given out by the output layer are obtained by using the “softmax” nonlinear function for the output neurons. From these probabilities the action is chosen by taking the neuron index with the highest probability value in the output.

Results

Single mode cooling. In a first step we numerically simulate the time dynamics of a single oscillating mode of frequency ω initially in a thermal distribution imposed by its coupling to an environment at some temperature T . This corresponds to the following distribution of energies

$$P(E) = Z^{-1} e^{-\beta E \hbar \omega}, \quad (6)$$

with a partition function $Z \approx (\beta \hbar \omega)^{-1}$ and the total occupation number normalized energy $E = (q^2 + p^2)/2$. We then randomly pick an initial energy value from the thermal distribution

$$E_0 = -(1/\beta \hbar \omega) \ln(1 - s), \quad (7)$$

by picking a random number s between zero and one. From the equipartition theorem we deduce $q(0)$ and $p(0)$ and train the neural network by recursively injecting sequences of $(q(t), \dot{q}(t))|_{[0, \dots, T]}$, obtained by applying the terms of Eqs. 3a, 3c and 4 recursively on the initial values, as training data into the network. A reward at each time step is only given when the action reduced the energy of the resonator at a given time step with respect to the previous time. The reward is defined by

$$R_t = (E_0 - E_{t+\Delta t}) \theta (E_t - E_{t+\Delta t}), \quad (8)$$

where E_t is the total energy at time t . The reward gets larger when the energy separation between the current and initial energy increases therefore optimizing the effective cooling rate (see Methods). The application of the reward to the single mode cooling is exemplified for the quadratic coupling configuration illustrated in Fig. 1d. The cooling dynamics is exemplified both as amplitude decreases Fig. 2a and in phase space Fig. 2b on two trajectories corresponding to two different initial states randomly picked from a thermal initial distribution with average occupancy $\bar{n} = 100$. The action sequences of the network in Fig. 2c show a periodic structure matching the frequency of the resonator mode, which is more visible in the zoom-in plot provided in Fig. 2e. There one can

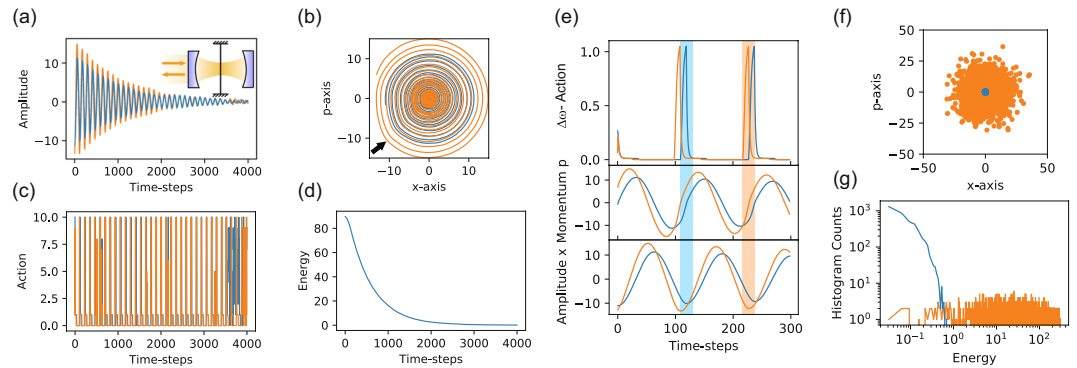


Figure 2. *Single mode parametric cooling.* (a) Time dynamics followed on two trajectories with initial conditions drawn from a Boltzmann distribution corresponding to an initial thermal state with average occupancy of $\bar{n} = 100$. The inset refers to the choice of cooling performed parametrically by modulation of the spring constant. (b) Corresponding phase space trajectories. The black arrow indicates where the force is applied. (c) Action sequences chosen by the network for each trajectory (d) Average energy for a thermal ensemble of trajectories exposed to the actions chosen by the network and rescaled to give the value of the occupation number of the harmonic oscillator. (e) Zoom-in into the time dynamics of the cavity modulated actions $\Delta\omega = 2g^{(2)}|\alpha(t)|^2$ and quadratures for two distinct trajectories. (f) Phase space comparison of initial (orange) and final (blue) distributions. The 4×10^3 points of the distribution of final states are obtained by running each trajectory starting from its initial state under the actions of the policy for 2×10^4 time steps. (g) Corresponding histogram of energy distribution in the initial and final states. The parameters are $\gamma = 4 \times 10^{-5} \omega$, $g^{(2)} = 1 \times 10^{-8} \omega$, $|\alpha|^2 \approx 0.5 \times 10^7$ and $\Delta t = 0.05 \omega^{-1}$.

follow the time dynamics of the applied action and the effect onto both the position and momentum, which in total for all trajectories results in the reduction of the average energy as presented in Fig. 2d.

While the training of the neural network to produce an optimal policy is done with training batches of 80 trajectories each with 4000 time steps (already approaching a low energy steady state as presented in Fig. 2a,b), we test the stability of the cooling policy by applying the trained network on a sample of thousands of trajectories with an extended time range of 20000 time steps. These results of thousands of sample trajectories are presented as initial and final phase space distributions in Fig. 2f and as a histogram of the energy distributions in Fig. 2g. The injected thermal noise in all plots in Fig. 2 corresponds to a thermal occupation number of $\bar{n} \approx 100$ and a thermalization rate of $\gamma/\omega = 4 \times 10^{-5}$. The choice of the initial thermal state is however arbitrary and with equal computational power one can also describe the dynamics of oscillators initially populated with more than 10^5 quanta. These results show that the policy reduces the energy of the mechanical mode and converges at a steady state irrespective of the initial high energy state derived from the thermal distribution. Additionally, the larger data set does not show any divergent outliers suggesting that convergence has been reached.

Simultaneous cooling of many modes. We display the generality of this approach by applying the network to find a strategy to cool up to four modes simultaneously. As in the case of the single mode cooling, we apply a single force on the mirror: this poses a challenge as a good cooling strategy for a given mode might actually lead to the heating of the other modes. In general, owing to this challenge, a simultaneous cooling strategy has an increased complexity in the choice of the action sequences, which leads to overall slower cooling rates. The same principles for cooling a single resonator are applied to cooling four modes subjected to the same actions by the network as presented in Fig. 3. Here, we use the setup configuration presented in Fig. 1b. While the four modes have different frequencies ω_j and coupling strengths g_j they are subjected to the same time sequence of actions delivered by intensity variations of an impinging laser beam in free space or via the field intensity $|\alpha(t)|^2$ in the cavity.

The input to the network is given by $s_t = (Q(t), \dot{Q}(t))$ with $Q(t) = \sum_j q_j(t)$ and $\dot{Q}(t) = \sum_j \omega_j p_j(t)$ as a collective position coordinate and its derivative. The quantities can be obtained for example from an interferometer that is sensitive to the fluctuations on the membrane (see Fig. 1e) or via homodyne detection. The derivation of the initial condition is described in the Methods section.

Here, the agent needs to find a strategy that simultaneously cools the center of mass motion as well as all of the relative mode dynamics. As an example the trajectories of the four modes are presented in Fig. 3a,b, where the cooling results from the corresponding actions presented in Fig. 3c with magnified view shown in Fig. 3c. In contrast to the action sequence imposed on the trajectory of a single resonator presented in Fig. 2c, which basically shows a periodic signal matching the frequency of the oscillator, here we find a more complex signal with a quasi periodic pattern. The change of the average energy of the four resonators as a function of time is presented in Fig. 3f. In Fig. 3e the initial values obtained from a Boltzmann distribution and final phase space values of a thousand trajectories for all oscillators are presented. A histogram of the sum of their individual energies is given in Fig. 3g. These results show that all four resonators can be simultaneously cooled down to lower temperatures, that differ by orders of magnitude from their initial values and thereby exemplify the strength of this adaptive approach.

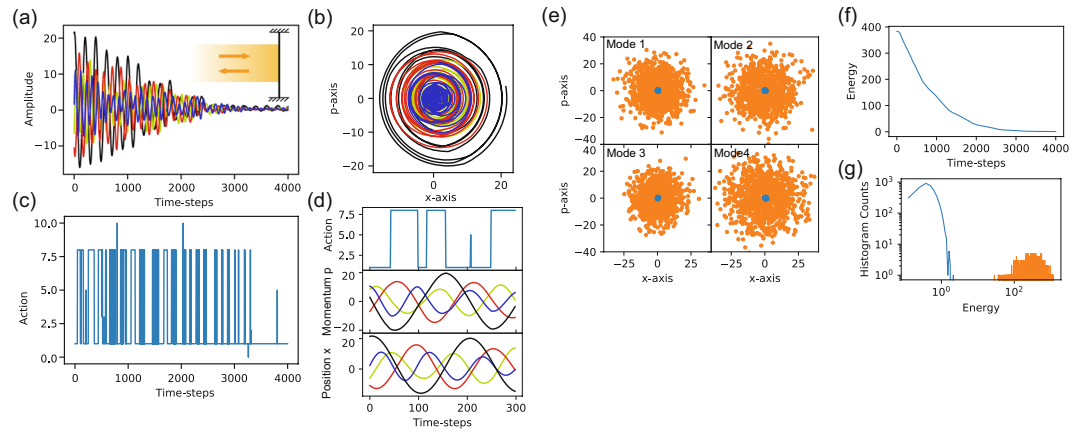


Figure 3. *Simultaneous cooling of 4 modes.* (a) Time dynamics of the oscillation amplitudes of four independent modes under the action of a collective force in the setup shown in the inset. (b) Corresponding time dynamics in phase space. (c) Neural network indicated sequence of actions leading to simultaneous cooling. (d) Magnification of the action and the momentum and position traces. (e) Reduction of the initial thermal distribution (orange) towards a low temperature distribution (blue) for each independent oscillation mode after 2×10^4 time steps. (f) Corresponding decrease of the total average energy of all four modes. (g) Histogram of initial (orange) and final (blue) energy distribution for 4×10^3 trajectories. The parameters for the simulation are given by $\omega_{2,3,4} = (0.8, 1.2, 0.6) \omega_1$, $\gamma_{1,2,3,4} = (4, 3, 5, 2) \times 10^{-5} \omega_1$, where the multiplication factors for the action are $g_{1,2,3,4} = (0.3, 0.2, 0.4, 0.3)$ to obtain F_j and we have $\Delta t = 0.05 \omega_1^{-1}$.

Discussions and Outlook

We have shown results of numerical simulations for the simultaneous cooling of a few degrees of freedom of a vibrating mechanical resonator. The feedback action has been realized via the reinforcement learning technique implemented on a neural network. There is a variety of other optimization methods that could obtain similar results. For example, stochastic optimization methods such as hill climbing, random walks or genetic algorithms³⁶. It has been recently shown that evolution strategies (ES) offer a similar performance and efficiency as RL⁴⁸. We have selected RL and especially the policy gradient method to approach this problem due to its efficiency when a large continuous or quasi continuous set of states is present⁴⁹.

Simultaneous cooling of a few modes indicates the possibility of partial or full refrigeration of mechanical resonators via optical control. It is remarkable that the network can perform efficient cooling of many modes while only being fed information of a time-evolving collective displacement quadrature. This is owed to the fact that the designed strategy optimizes single mode cooling at every instance in time while keeping the heating of all other modes to small values. The described procedure works both inside and outside optical cavities and both in linear or nonlinear regimes therefore being easily adaptable to new systems. The technique could be easily extended to cool a number of oscillators or a number of particles trapped inside optical cavities or with tweezers. While the present treatment considers classical stochastic dynamics, a full quantum theory of neural network aided cooling will be tackled in the future that might also lead towards feedback production of squeezed or squashed states. In this regard, for both single and many oscillation modes, we plan to analyze the efficiency of neural network cooling in comparison with standard cold-damping optical cooling. It is expected that a Fourier analysis of the action function indicated by the network could hint towards feedback implementations that could surpass existing techniques, especially at the level of many degrees of freedom.

Methods

Initial conditions. We assume an initial Boltzmann distribution of energies $P(E_1, \dots, E_n) = Z^{-1} e^{-\beta(E_1 + \dots + E_n)}$ where $Z = 1/\beta^n$, from which we extract the initial conditions by integration

$$\int_0^{\tilde{E}_1} \dots \int_0^{\tilde{E}_n} dE_1 \dots dE_n P(E_1, \dots, E_n) = \prod_{j=1}^n [1 - e^{-\beta \tilde{E}_j}] = s,$$

where s is a value between 0 and 1. We set $b_j = [1 - e^{-\beta \tilde{E}_j}]$ and define that $b_j \in [0, 1]$ for $0 \leq j \leq n$ which results in

$$\tilde{E}_j = -\frac{1}{\beta} \ln(1 - b_j). \quad (9)$$

Since $p_j(0)^2 + q_j(0)^2 = 2\tilde{E}_j/\hbar\omega_j$ we obtain p_j and q_j from

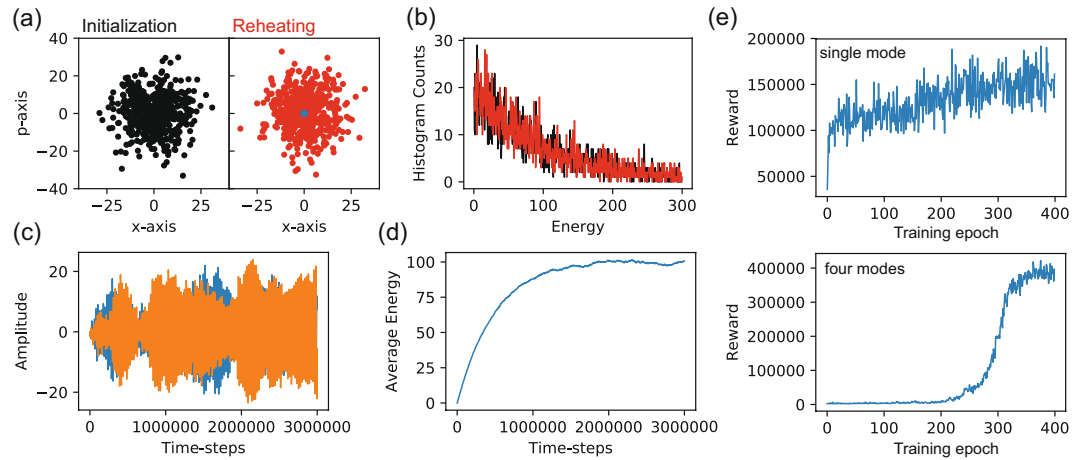


Figure 4. Thermalization dynamics of a single mode of vibration. (a) Phase space coordinates from initialization with a Boltzmann distribution with $\bar{n} = 100$ (black points) and obtained from the evolution equation containing thermal noise. An initial distribution around the phase space origin (blue points) is evolved up to the steady state (red points). In (b) the corresponding energy distributions following the same color coding are presented. (c) Two example trajectories with starting energies close to zero are following the evolution guided by thermal noise. In (d) the average energy is presented as a function of time. (e) The learning progress for the single and four-mode cooling presented in the main text. Here, the reward increases steadily with each training epoch until reaching a saturation.

$$q_j(0) = \sqrt{\frac{2\bar{E}_j}{\hbar\omega_j}} \cos(2\pi\phi_j), \tag{10a}$$

$$p_j(0) = \sqrt{\frac{2\bar{E}_j}{\hbar\omega_j}} \sin(2\pi\phi_j), \tag{10b}$$

where ϕ_j is a random number between 0 and 1 for all $j \in 1, \dots, n$.

Thermalization dynamics. Let us describe the numerical procedure for simulating the action of a thermal environment onto the state of the mechanical resonator. We consider the equations of motion

$$\dot{q} = \omega p, \tag{11a}$$

$$\dot{p} = -\omega q - \gamma p + \xi, \tag{11b}$$

which we rewrite as equations of differential forms

$$dq = \omega p dt, \tag{12a}$$

$$dp = -\omega q dt - \gamma p dt + \sqrt{(2\bar{n} + 1)\gamma} dW(t), \tag{12b}$$

where the noise $dW(t)$ is included as a Wiener process. Numerically, this can be realized by $\Delta W(t) \propto \sqrt{\Delta t} N(0, 1)$, where $N(0, 1)$ describes a normally distributed random variable of unit variance - consequently the Wiener increment is normally distributed with a variance equal to the numerical time increment Δt . As a numerical check we simulate the thermalization of an initially cold mechanical mode under the action of an environment with rate $\gamma = 4 \times 10^{-5} \omega$ and at an effective occupancy $\bar{n} = 100$. The analytical Boltzmann distribution for this occupancy is shown in Fig. 4a as black dots while the final state obtained from the numerical integration is represented by the red dots (blue dots in the middle are the initial state). In Fig. 4b the agreement between the numerical simulation and the Boltzmann distribution is illustrated as a histogram of energy states. The time evolution is shown in Fig. 4c as dynamics for the position and in Fig. 4d for the total energy showing it approaching $\bar{n} = 100$ in the long time limit.

Additionally, in Fig. 4e we present the learning process for the results presented in Fig. 2 for a single mode and for four modes as presented in Fig. 3, where the increase in the average reward over 400 epochs each with batches of 80 trajectories is shown.

	Network parameters
Training batch size	80
Learning rate η (Adam optimizer)	0.00008 (single oscillator)
	0.0006 (four oscillators)
Neurons per layer	(2, 60, 60, 11) (single oscillator)
	(2, 100, 100, 11) (four oscillators)
Reward scale	1.0

Table 1. List of neural network and training parameters.

Markov decision process. The reinforcement learning procedure presented above can be fully described by a discrete time stochastic control process. Since in this process the agent selects an action based on the stochastic policy $\pi_\theta(a|s)$ which is only dependent on the current observation of the given state s , the problem is described by a Markov decision process (MDP)⁴⁶. Formally, a Markov decision process is a 4-tuple (S, A, P, R) , where S is the set of states, A forms the set of possible actions, $P: S \times A \times S \rightarrow [0, 1]$ is the transition function between states and $R: S \times A \times S \rightarrow \mathcal{R}$ is the reward function as described above with $\mathcal{R} = [R_{\min}, R_{\max}] \subset \mathbb{R}$ being the continuous set of possible rewards.

In the case of a single resonator mode the state space is given by $S = \{s = (q, p) | p = \omega^{-1}\dot{q}, q, p \in \mathbb{R}\}$ while the action space is the finite set $A = \{0, 1, \dots, 10\}$ which allows the agent to choose between ten different force strength. The transition function $P(s_{t+\Delta t}|a_t, s_t)$ giving the probability for moving to the state $s_{t+\Delta t}$ from s_t under the action a_t can be obtained from the equations of motion $\dot{s} = Ms + \xi + \tilde{a}$ with $\xi = (0, \xi)^\top$ being the noise and $\tilde{a} = (0, F(a))^\top$ the force term. In the case the random noise contribution ξ is zero the transition function is deterministic and $P(s_{t+\Delta t}|a_t, s_t) = 1$ for $s_{t+\Delta t} = \tilde{M}s_t + \tilde{a}_t\Delta t$ with $\tilde{M} = 1 + M\Delta t$ and zero otherwise. The reward function is given by the expression defined above $R_t = R(s_{t+\Delta t}, a_t, s_t) = (E_0 - E_{t+\Delta t})\theta(E_t - E_{t+\Delta t})$. For multiple resonator modes where we observe $s_t = (Q(t), \dot{Q}(t))$ with $Q(t) = \sum_j q_j(t)$ and $\dot{Q}(t) = \sum_j \omega_j p_j(t)$, we only obtain partial information of the state which originally is described by the phase space vector $(q_1, \dots, q_n, p_1, \dots, p_n)$. Here we need a generalization of an MDP which is given by a partially observable Markov decision process (POMDP), where the agent cannot observe the full state. Here, we additionally have the sets O which describes the set of observations and O describing the set of conditional observation probabilities.

A pseudo code to implement reinforcement learning (RL) is presented in the following:

```

initialize  $\theta$ 
for iteration = 1, 2, ...,  $N$  epoch do
  initialize training batch of initial states  $s$ 
  for each time step  $t$  do
     $a_t = \arg \max_{a' \in A} \pi_\theta(a', s_t)$  sample action for the full batch
     $s_{t+\Delta t} = \tilde{M}s_t + d\tilde{W} + \tilde{a}_t\Delta t$  run dynamics for the full batch
  end for
   $R = \sum_t R_t$  total reward for each trajectory of the batch
   $\theta \leftarrow \theta + \eta \nabla_\theta \mathbb{E}(R)$  update network parameters
end for

```

Network parameters. The simulations were run on a standard Laptop computer (CPU, Intel Core i7 – 5500U @ 2.40 GHz). We use the Keras package for Python and the Theano framework⁵⁰ to realize the neural network and the reinforcement learning procedure⁵¹. The network and training parameters are presented in Table 1.

Received: 25 July 2019; Accepted: 28 January 2020;

Published online: 14 February 2020

References

- Aspelmeyer, M., Kippenberg, T. J. & Marquardt, F. Cavity optomechanics. *Rev. Mod. Phys.* **86**, 1391–1452, <https://doi.org/10.1103/RevModPhys.86.1391> (2014).
- Windey, D. *et al.* Cavity-based 3d cooling of a levitated nanoparticle via coherent scattering. *Phys. Rev. Lett.* **122**, 123601, <https://doi.org/10.1103/PhysRevLett.122.123601> (2019).
- Delic, U. *et al.* Cavity cooling of a levitated nanosphere by coherent scattering. *Phys. Rev. Lett.* **122**, 123602, <https://doi.org/10.1103/PhysRevLett.122.123602> (2019).
- Rossi, M. *et al.* Enhancing sideband cooling by feedback-controlled light. *Phys. Rev. Lett.* **119**, 123603, <https://doi.org/10.1103/PhysRevLett.119.123603> (2017).
- Clark, J. B., Lecocq, F., Simmonds, R. W., Aumentado, J. & Teufel, J. D. Sideband cooling beyond the quantum backaction limit with squeezed light. *Nature* **541**, 191 EP, <https://doi.org/10.1038/nature20604> (2017).
- Qiu, L., Shomroni, I., P. S. & Kippenberg, T. J. High-fidelity laser cooling to the quantum ground state of a silicon nanomechanical oscillator. arXiv:1903.10242, <https://arxiv.org/abs/1903.10242> (2019).

7. Asenbaum, P., Kuhn, S., Nimmrichter, S., Sezer, U. & Arndt, M. Cavity cooling of free silicon nanoparticles in high vacuum. *Nature Communications* **4**, 2743 EP, <https://doi.org/10.1038/ncomms3743> (2013).
8. Mancini, S., Vitali, D. & Tombesi, P. Optomechanical cooling of a macroscopic oscillator by homodyne feedback. *Phys. Rev. Lett.* **80**, 688–691, <https://doi.org/10.1103/PhysRevLett.80.688> (1998).
9. Schäfermeier, C. *et al.* Quantum enhanced feedback cooling of a mechanical oscillator using nonclassical light. *Nature Communications* **7**, 13628 EP, <https://doi.org/10.1038/ncomms13628> (2016).
10. Kiesel, N. *et al.* Cavity cooling of an optically levitated submicron particle. *Proceedings of the National Academy of Sciences* **110**, 14180–14185, <https://www.pnas.org/content/110/35/14180.full.pdf> (2013).
11. Millen, J., Fonseca, P. Z. G., Mavrogordatos, T., Monteiro, T. S. & Barker, P. F. Cavity cooling a single charged levitated nanosphere. *Phys. Rev. Lett.* **114**, 123602, <https://doi.org/10.1103/PhysRevLett.114.123602> (2015).
12. Genes, C., Vitali, D., Tombesi, P., Gigan, S. & Aspelmeyer, M. Ground-state cooling of a micromechanical oscillator: Comparing cold damping and cavity-assisted cooling schemes. *Phys. Rev. A* **77**, 033804, <https://doi.org/10.1103/PhysRevA.77.033804> (2008).
13. Steixner, V., Rabl, P. & Zoller, P. Quantum feedback cooling of a single trapped ion in front of a mirror. *Phys. Rev. A* **72**, 043826, <https://doi.org/10.1103/PhysRevA.72.043826> (2005).
14. Bushev, P. *et al.* Feedback cooling of a single trapped ion. *Phys. Rev. Lett.* **96**, 043003, <https://doi.org/10.1103/PhysRevLett.96.043003> (2006).
15. Rossi, M., Mason, D., Chen, J., Tsaturyan, Y. & Schliesser, A. Measurement-based quantum control of mechanical motion. *Nature* **563**, 53–58, <https://doi.org/10.1038/s41586-018-0643-8> (2018).
16. Cohadon, P. F., Heidmann, A. & Pinard, M. Cooling of a mirror by radiation pressure. *Phys. Rev. Lett.* **83**, 3174–3177, <https://doi.org/10.1103/PhysRevLett.83.3174> (1999).
17. Poggio, M., Degen, C. L., Mamin, H. J. & Rugar, D. Feedback cooling of a cantilever's fundamental mode below 5 mk. *Phys. Rev. Lett.* **99**, 017201, <https://doi.org/10.1103/PhysRevLett.99.017201> (2007).
18. Wilson, D. J. *et al.* Measurement-based control of a mechanical oscillator at its thermal decoherence rate. *Nature* **524**, 325 EP, <https://doi.org/10.1038/nature14672> (2015).
19. Tebbenjohanns, F., Frimmer, M., Militaru, A., Jain, V. & Novotny, L. Cold Damping of an Optically Levitated Nanoparticle to Microkelvin Temperatures. *Phys. Rev. Lett.* **122**, 223601, <https://doi.org/10.1103/PhysRevLett.122.223601> (2019).
20. Gigan, S. *et al.* Self-cooling of a micromirror by radiation pressure. *Nature* **444**, 67–70, <https://www.nature.com/articles/nature05273> (2006).
21. Braginsky, V. B., Strigin, S. E. & Vyatchanin, S. P. Parametric oscillatory instability in Fabry–Perot interferometer. *Phys. Lett. A* **287**, 331, <https://www.sciencedirect.com/science/article/pii/S0375960101005102?via%3Dihub> (2001).
22. Marquardt, F., Chen, J. P., Clerk, A. A. & Girvin, S. M. Quantum theory of cavity-assisted sideband cooling of mechanical motion. *Phys. Rev. Lett.* **99**, 093902, <https://doi.org/10.1103/PhysRevLett.99.093902> (2007).
23. Wilson-Rae, I., Nooshi, N., Zwerger, W. & Kippenberg, T. J. Theory of ground state cooling of a mechanical oscillator using dynamical backaction. *Phys. Rev. Lett.* **99**, 093901, <https://doi.org/10.1103/PhysRevLett.99.093901> (2007).
24. Teufel, J. D. *et al.* Sideband cooling of micromechanical motion to the quantum ground state. *Nature* **475**, 359–363, <https://doi.org/10.1038/nature10261> (2011).
25. Sommer, C. & Genes, C. Partial optomechanical refrigeration via multimode cold-damping feedback. *Phys. Rev. Lett.* **123**, 203605, <https://doi.org/10.1103/PhysRevLett.123.203605> (2019).
26. Chen, C., Dong, D., Li, H., Chu, J. & Tarn, T. Fidelity-based probabilistic q-learning for control of quantum systems. *IEEE Transactions on Neural Networks and Learning Systems* **25**, 920–933 (2014).
27. Carrasquilla, J. & Melko, R. G. Machine learning phases of matter. *Nat. Phys.* **13**, 431, <https://doi.org/10.1038/nphys4035> (2017).
28. van Nieuwenburg, E. P. L., Liu, Y.-H. & Huber, S. D. Learning Phase Transitions by Confusion. *Nat. Phys.* **13**, 435, <https://doi.org/10.1038/nphys4037> (2017).
29. Dunjko, V. & Briegel, H. J. Machine learning & artificial intelligence in the quantum domain: a review of recent progress. *Reports on Progress in Physics* **81**, 074001, [10.1088/1361-6633/2018/81/7/074001](https://doi.org/10.1088/1361-6633/2018/81/7/074001) (2018).
30. Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606, <https://science.sciencemag.org/content/355/6325/602.full.pdf> (2017).
31. Mills, K., Spanner, M. & Tamblyn, I. Deep learning and the Schrödinger equation. *Phys. Rev. A* **96**, 042113, <https://doi.org/10.1103/PhysRevA.96.042113> (2017).
32. Melnikov, A. A. *et al.* Active learning machine learns to create new quantum experiments. *Proceedings of the National Academy of Sciences* **115**, 1221–1226, <https://www.pnas.org/content/115/6/1221.full.pdf> (2018).
33. Bukov, M. *et al.* Reinforcement learning in different phases of quantum control. *Phys. Rev. X* **8**, 031086, <https://doi.org/10.1103/PhysRevX.8.031086> (2018).
34. Fösel, T., Tighineanu, P., Weiss, T. & Marquardt, F. Reinforcement learning with neural networks for quantum feedback. *Phys. Rev. X* **8**, 031084, <https://doi.org/10.1103/PhysRevX.8.031084> (2018).
35. Sweke, R., Kesselring, M. S., van Nieuwenburg, E. P. L. & Eisert, J. Reinforcement learning decoders for fault-tolerant quantum computation. 1810.07207 (2018).
36. Russel, S. & Norvig, P. *Artificial Intelligence: A Modern Approach*. (Pearson, Boston, 2018).
37. Nielsen, W. H. P., Tsaturyan, Y., Möller, C. B., Polzik, E. S. & Schliesser, A. Multimode optomechanical system in the quantum regime. *Proceedings of the National Academy of Sciences* **114**, 62–66, <https://www.pnas.org/content/114/1/62.full.pdf> (2017).
38. Piergentili, P. *et al.* Two-membrane cavity optomechanics. *New Journal of Physics* **20**, 083024, [10.1088/1367-2630/20/8/083024](https://doi.org/10.1088/1367-2630/20/8/083024) (2018).
39. Wei, X., Sheng, J., Yang, C., Wu, Y. & Wu, H. Controllable two-membrane-in-the-middle cavity optomechanical system. *Phys. Rev. A* **99**, 023851, <https://doi.org/10.1103/PhysRevA.99.023851> (2019).
40. Thompson, J. D. *et al.* Strong dispersive coupling of a high-finesse cavity to a micromechanical membrane. *Nature* **452**, 72–75, <https://doi.org/10.1038/nature06715> (2008).
41. Jayich, A. M. *et al.* Dispersive optomechanics: a membrane inside a cavity. *New Journal of Physics* **10**, 095008, [10.1088/1367-2630/10/9/095008](https://doi.org/10.1088/1367-2630/10/9/095008) (2008).
42. Asjad, M. *et al.* Robust stationary mechanical squeezing in a kicked quadratic optomechanical system. *Phys. Rev. A* **89**, 023849, <https://doi.org/10.1103/PhysRevA.89.023849> (2014).
43. Higham, D. An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Review* **43**, 525–546, <https://doi.org/10.1137/S003614450037830> (2001).
44. Sutton, R. S. & Barto, A. G. *Reinforcement learning: An introduction*. (MIT press, Cambridge, 1998).
45. Williams, R. J. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning* **8**, 229–256, <https://doi.org/10.1007/BF00992696> (1992).
46. Bellman, R. A markovian decision process. *Journal of Mathematics and Mechanics* **6**, 679–684, <http://www.jstor.org/stable/24900506> (1957).
47. Weaver, L. & Tao, N. The Optimal Reward Baseline for Gradient-Based Reinforcement Learning. *UAI P*, 538, <https://arxiv.org/abs/1301.2315> (2001).
48. Salimans, T., Ho, J., Chen, X., Sidor, S. & Sutskever, I. Evolution strategies as a scalable alternative to reinforcement learning. 1703.03864 (2017).

49. Dutta, S. *Reinforcement Learning with TensorFlow* (Packt Publishing Ltd., 2018).
50. Al-Rfou, R. *et al.* Theano: A Python Framework for Fast Computation of Mathematical Expressions. arXiv:1605.02688, <https://arxiv.org/abs/1605.02688> (2016).
51. Marquardt, F. Machine Learning for Physicists, <https://machine-learning-for-physicists.org/> (2017).

Acknowledgements

We acknowledge financial support from the Max Planck Society. We acknowledge fruitful discussions with Michael Reitz. We are thankful for the comprehensive lecture notes on machine learning for physicists presented by Florian Marquardt at the University of Erlangen-Nuremberg.

Author contributions

C.S. has been the main investigator in this project taking part in developing the idea, model and goals, performing calculations, numerical simulations, producing figures and taking part in writing. M.A. and C.G. participated in calculations and model building and C.G. supervised the project and took part in the writing process.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to C.S.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020