Thermodynamic Cyclic Processes with Markov Samplers in Bayesian Inference

Anonymous authors
Paper under double-blind review

Abstract

The concept of Markov chain Monte Carlo (MCMC) cycles, an analogy to cyclic processes in heat engines, is presented in order to examine Bayesian inference problems. In this effort, we develop adaptive ensemble schedulers that allow the tuning of external parameters of a Bayesian canonical ensemble during an MCMC run. We apply our method to different statistical models. As a fundamental insight, we find that such systems produce a non-zero net work output if and only if the considered model is non-Gaussian.

1 Introduction

One of the prime motivations for the development of the theory of thermodynamics in the 19th century was the study of heat engines as the first apparatuses that could translate heat into work as a usable form of energy (Saslow, 2020). In the beginning, the nature of gases, their working substances, was poorly understood on the microscopic level. Thus, thermodynamics was developed as an *effective* theory to accurately capture their macroscopic behaviour nonetheless. Subsequently, thermodynamics became a powerful theory that could be applied to systems of few as well as many degrees of freedom. When the "molecular picture" of matter was established, statistical physics was developed by Boltzmann to bridge the gap to the macroscopic thermodynamic notions, and the relation to information theory was established by Jaynes (Jaynes, 1957).

Thermodynamics, in particular Markov Chain Monte-Carlo sampling, has been applied to the field of Bayesian inference (Hastings, 1970; Brooks et al., 2011; Lewis & Bridle, 2002). Bayes' theorem,

$$p(\theta \mid y) = \frac{\mathcal{L}(y \mid \theta)\pi(\theta)}{p(y)}\,, \tag{1}$$

with the likelihood $\mathcal{L}(y \mid \theta)$, parameter θ, the data y, and the prior $\pi(\theta)$, requires the evidence as a normalization constant

$$p(y) = \int \mathrm{d}^n\theta\, \mathcal{L}(y \mid \theta)\pi(\theta)\,. \tag{2}$$

Additional parameters, the temperature T and the sources J, extend this integral to a canonical partition sum in the sense of statistical physics

$$Z(T, J) = \int \mathrm{d}^n\theta\, (\mathcal{L}(y \mid \theta)\pi(\theta))^{1/T}\, e^{J\cdot\theta/T}\,. \tag{3}$$

To this one may then apply the tools from statistical mechanics and thermodynamics to obtain a theory of the Bayesian canonical ensemble and its extensions (Röver et al., 2023a; Kuntz et al., 2025; Röver et al., 2023b; Kuntz et al., 2024). In practice, the integrals often cannot be carried out analytically, which is why one resorts to Markov chain Monte Carlo (MCMC) sampling algorithms such as the Rosenbluth-Metropolis-Hastings (RMH) algorithm (Metropolis et al., 1953; Hastings, 1970). Much like in a physical gas, such algorithms work by simulating a random walk in the sample space. The samples of the posterior are then given by the trajectory of the sampling "particles".

Aside from sampling, a different use case for this algorithm is called *simulated annealing* with the aim of finding the global maximum of the posterior w. r. t. the position θ (van Laarhoven & Aarts, 1987). This is

accomplished by gradually decreasing the temperature T such that the sampler's motion is more and more restricted around the maximum. Our aim here is to generalize this mechanism in order to vary *both* the temperature T and the sources J during a simulation run. The sources J serve as a practical pathway to cumulants of the posterior distribution by logarithmic differentiation.

We establish a new method named MCMC cycles where both parameters are changed in turn, forming a closed loop in (T, J) space; this is analogous to cycles in heat engines. We tune the parameters in a thermodynamically consistent way by introducing *adaptive ensemble schedulers* inspired by the notion of quasi-staticity (see section 3). The thermodynamic theory derived from the partition sum (3) allows us to make predictions about the exchanged heat and performed work during the MCMC cycles (section 4). Our aim is to test whether these theoretical predictions may be confirmed in numerical experiments and what conclusions may be drawn from them (section 5). This new perspective on MCMC methods widens our understanding by verifying the validity of thermodynamics therein and characterising the underlying inference problems in a new manner. We will use index notation with Einstein summing convention.

2 Related work

MCMC methods were originally developed for use in statistical physics (Metropolis et al., 1953; Hastings, 1970) but have since been applied to many other fields, Lattice QCD (Creutz, 1980; Duane et al., 1987) and Bayesian inference (Gelfand & Smith, 1990; Geman & Geman, 1984) among them. The thermodynamic properties of heat engines have been studied since the 19th century (Carnot, 1872; Clausius, 1879; Thomson, 1852). While the general thermodynamic nature of Bayesian inference and MCMC methods has been studied extensively (Jaynes, 1957; Röver et al., 2023a;b; Herzog et al., 2024; Kuntz et al., 2024; 2025), there has been no work on thermodynamic cyclic processes in this context. While the work by Wang et al. (2024) is related by title, neither the aim nor the proposed method are connected to ours.

3 A tunable RMH sampling algorithm

An iteration of the aforementioned RMH algorithm is made up of (i) proposing a new position $\theta_{\text{prop}} \sim p_{\text{prop}}(\cdot \mid \theta_{\text{curr}})$ from a proposal distribution, typically a Gaussian centered around the current position $\mathcal{N}(\theta_{\text{curr}}, \Sigma_{\text{prop}})$; and (ii) accepting this proposal with probability

$$p_{\text{acc}}(\theta_{\text{prop}} \mid \theta_{\text{curr}}) = \min\left(1, e^{-\frac{1}{T}(V(\theta_{\text{prop}};J) - V(\theta_{\text{curr}};J))}\right), \tag{4}$$

where the potential corresponding to the Bayesian canonical ensemble above is $V(\theta; J) = -\log(\mathcal{L}\pi) - J \cdot \theta$. [1] Considering a potential that depends on additional external parameters J as above, our aim here is to develop a mechanism that allows tuning of both the temperature T and the external parameter J while the RMH algorithm runs.

The naive approach would be to implement two schedulers $\sigma_T(t)$ and $\sigma_J(t)$ that return the current values for the temperature T and the parameters J respectively at simulation step t. The acceptance probability at step t is then modified,

$$p_{\text{accept}}(\theta_{\text{prop}}, \sigma_T(t+1), \sigma_J(t+1) \mid \theta_{\text{curr}}, \sigma_T(t), \sigma_J(t)) = \min\left(1, e^{-\frac{V(\theta_{\text{prop}};\sigma_J(t+1))}{\sigma_T(t+1)} + \frac{V(\theta_{\text{curr}};\sigma_J(t))}{\sigma_T(t)}}\right), \tag{5}$$

taking into account the change in position θ, temperature T and external parameters J. The central question for this setup is then how slowly the external parameters should be changed. E. g., if one decreases the temperature very quickly between steps t and $t+1$, the first term in the exponential in (5) will dominate such that $p_{\text{accept}} \to 0$ and no further proposals will be accepted.

An excellent criterion for this is that of *quasi-staticity* from thermodynamics. It states that the time scale at which a thermodynamic system should be perturbed externally should be much larger than the time scale at

[1] Please note that the acceptance probability only takes the mentioned form if the proposal distribution is symmetric, $p_{\text{prop}}(\theta_{\text{prop}} \mid \theta_{\text{curr}}) = p_{\text{prop}}(\theta_{\text{curr}} \mid \theta_{\text{prop}})$. In our case this is given.

which the system equilibrates, such that at any point during the induced change of state, the system remains in perfect thermodynamic equilibrium. Note that this does not mean that the change of state must happen "slowly" by the standards of human perception (e.g., the cyclic processes in heat engines are accurately described by thermodynamics while they certainly do not look slow to the human eye). Furthermore, the notion of quasi-staticity is different from reversibility, which will be discussed below for the use case of sampling algorithms.

During a normal RMH step the scale of the proposal step size is typically chosen such that $V(\theta_{\text{prop}}) - V(\theta_{\text{curr}}) \propto T$. Thus, with the reasonable assumption that V is smooth in θ, J and T, one should choose the scale at which the external parameters are changed such that

$$V(\theta; \sigma_J(t+1)) - V(\theta; \sigma_J(t)) \ll T\,. \tag{6}$$

For a transition from $J = J_i$ to J_f, we do so by picking some small $\epsilon > 0$ and choosing the step size ΔJ such that $|V(J + \Delta J) - V(J)| \approx |\Delta J \cdot \nabla_J V| = \epsilon T$. This is achieved by

$$\Delta J = \min\left(\Delta J_{\max}, \frac{\epsilon T (J_f - J_i)}{|\nabla_J V \cdot (J_f - J_i)|}\right)\,, \tag{7}$$

where we introduced a maximum step size $\Delta J_{\max}$ to mitigate division by zero. By choosing ϵ sufficiently small, the acceptance probability remains uninhibited by the change in the external parameter J and the sampling continues adapting in real time and maintaining thermodynamic equilibrium. Additionally, instead of a single gradient we run multiple chains at the same time and use the average ensemble gradient $\langle \nabla_J V \rangle_{\text{chains}}$ as criterion for the step size. This is again inspired by physical heat engines where the combined motion of many particles drives the macroscopic motion of the piston. For the potential corresponding to the canonical partition sum (3), we find $\nabla_J V = -\theta$.

Similarly, for changes in temperature from $T = T_i$ to T_f, we pick $\Delta T = \epsilon T$ such that we find $\sigma_T(t+1) = (1+\epsilon)\sigma_T(t)$. This means that we may write

$$\sigma_T(t) = (1+\epsilon)^t T_i\,, \tag{8}$$

corresponding to an *exponential scheduler* in simulated annealing (van Laarhoven & Aarts, 1987).

4 Thermodynamic Concepts for an RMH cyclic process

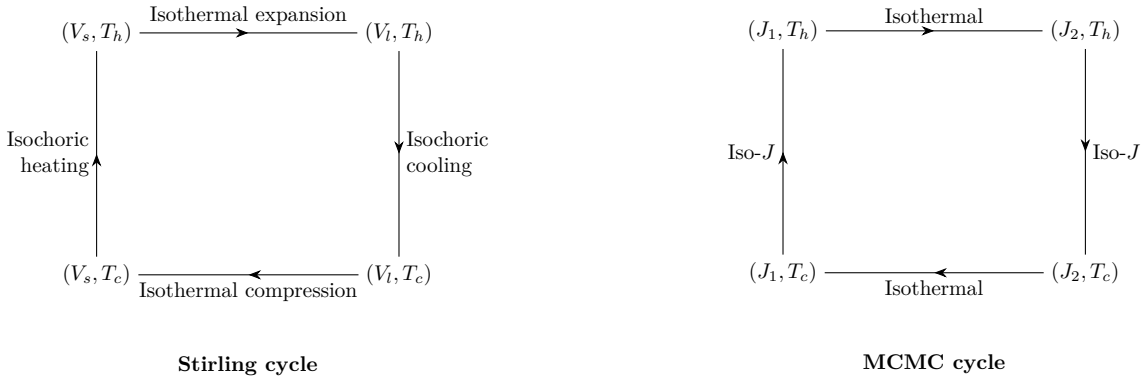

Figure 1: Comparison of a thermodynamic Stirling cycle (left) and its information theoretical analogue, the MCMC cycle (right).

With the canonical partition function defined above (3), one may define the free energy $F(T, J) = -T \log Z(T, J)$. Its total differential $\mathrm{d}F(T, J) = -S\mathrm{d}T + \Lambda \cdot \mathrm{d}J$ then allows the deduction of the thermodynamic entropy,

$$S = -\frac{\partial F}{\partial T}\,, \tag{9}$$

as well as the quantity $\Lambda_i = \partial F / \partial J_i$, which corresponds to volume or pressure for systems with gases as their working substance in their physical interpretation, but corresponds as well to the cumulants of the posterior distribution from which the Markov chain generates samples. For our particular case of a canonical partition function, the entropy S at unit temperature is in fact Shannon's entropy of the posterior distribution, making yet another link between thermodynamics and statistics.

The first law of thermodynamics relates the change of the inner energy $\mathrm{d}U$ to the heat exchanged with the system δQ and work performed on it δW for any quasi-static process; it states $\mathrm{d}U = \delta Q + \delta W$. Such a process is additionally called reversible if it may be undone without lasting change to the environment. (Intuitively, a process is reversible if one cannot decide whether a recording of that process is being played forward or backward.) In reversible processes, the heat form may be written as $\delta Q = T\mathrm{d}S$. Since both the chosen proposal and acceptance probabilities are symmetric,

$$\begin{aligned} &p_{\text{prop/acc}}(\theta_{\text{prop}}, \sigma_T(t+1), \sigma_J(t+1) \mid \theta_{\text{curr}}, \sigma_T(t), \sigma_J(t)) \\ &\qquad = p_{\text{prop/acc}}(\theta_{\text{curr}}, \sigma_T(t), \sigma_J(t) \mid \theta_{\text{prop}}, \sigma_T(t+1), \sigma_J(t+1))\,, \end{aligned}$$

the tunable RMH algorithm proposed above is generally reversible in the thermodynamic sense.

Thermodynamic cyclic processes in general are a series of changes of state that form a closed loop in the sense that the final state of the system at each cycle is the same as the initial one. As a classical example, consider the Stirling cycle (Moran et al., 2010). The idealized Stirling engine contains an ideal gas and the cycle is made up of four phases that are described in Figure 1 (left). In the first phase $(1 \to 2)$, the gas expands isothermally from volume V_s to V_l while in contact with a hot thermostat at temperature T_h. Then $(2 \to 3)$ the gas is cooled by contact with a colder thermostat from T_h to T_c while retaining the same volume V_l. The step $(1 \to 2)$ is then mirrored in $(3 \to 4)$ by a compression from V_l to V_s at the low temperature T_c. Finally, for $(4 \to 1)$, the contact switches again to the hot thermostat to heat up the gas from T_c to T_h while maintaining the volume V_s.

The Stirling cycle and engine have many interesting thermodynamic properties (Moran et al., 2010). For us, it is mainly an inspiration for the MCMC cyclic process depicted in Figure 1 (right). The sources J assume the role of the volume V here. Please note that the question of whether J are intensive (like the pressure P) or extensive (like the volume V) in nature is not straightforward to answer (Kuntz et al., 2025). Since it is not relevant to our discussion, we merely use the Stirling cycle as an analogy. (The corresponding "intensive" thermodynamic cycle where the changing parameters are (P, T) is in fact the Ericsson cycle.)

To implement the cycle, we use adaptive ensemble schedulers σ_T for the temperature and σ_J for the sources. We predefine the target temperatures (T_c, T_h) and sources (J_1, J_2) and let the scheduler run until the desired values are reached for each phase. The outputs of a run of such a cycle are the samples $\{\theta_{t,c} \in \mathbb{R}^n \mid t = 1, \ldots, N, c = 1, \ldots, N_{\text{chains}}\}$ as well as the trajectories of the temperatures $\{T_t \in \mathbb{R} \mid t = 1, \ldots, N\}$ and the sources $\{J_t \in \mathbb{R}^n \mid t = 1, \ldots, N\}$.

To analyze the data, we compute the change of the inner energy by finite differencing, $\dot{U}_t = \langle V(\theta_t; J_t) - V(\theta_{t-1}; J_{t-1}) \rangle_{\text{chains}}$, the rate of the work performed as $\dot{W}_t = \langle \Lambda_t \cdot (J_t - J_{t-1}) \rangle_{\text{chains}}$ and the rate of the exchanged heat as $\dot{Q}_t = \dot{U}_t - \dot{W}_t$ with the first law of thermodynamics. All of these may be integrated numerically to yield the change in inner energy ΔU_t, exchanged heat ΔQ_t, and performed work ΔW_t as a function of time.

Regarding heat engines, the most important properties of any cyclic process are the net work (Moran et al., 2010)

$$\Delta W_{\text{net}} = \oint_{\text{one cycle}} \mathrm{d}W \tag{10}$$

per cycle performed by the engine as well as the heat input ΔQ_{in}. They allow us to define the efficiency $\eta = \Delta W_{\text{net}} / \Delta Q_{\text{in}}$ which measures how well the engine "translates" the input energy in the form of heat into work. We will compute these quantities for MCMC cyclic processes with different statistical models.

To compare them to a theoretical prediction, we compute the entropy S from the free energy and its partial derivatives given by the heat capacity $C = T\partial S/\partial T$ as well as $h^i = T\partial S/\partial J_i$ to find

$$\dot{Q} = T\dot{S} = C\dot{T} + h \cdot \dot{J} \,. \tag{11}$$

5 Application

5.1 Gaussian Toy Model

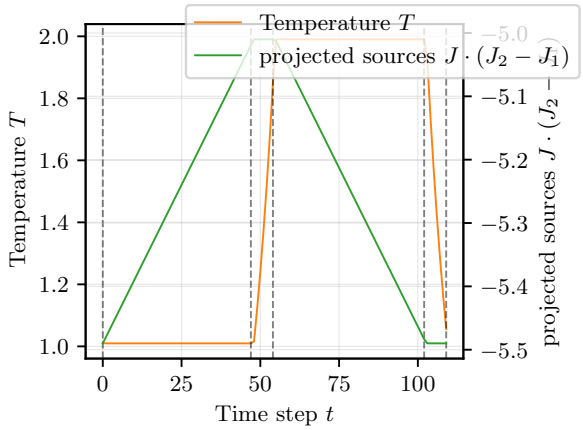

Figure 2: Trajectories of the parameters T and J during an MCMC cycle in an anisotropic Gaussian model as a function of simulation time.

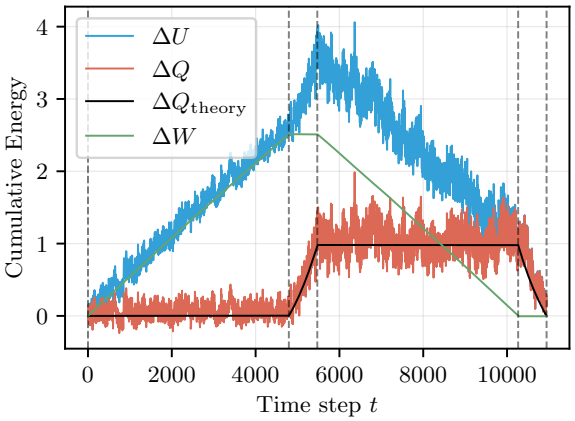

Figure 3: Energy changes during an MCMC cycle in an anisotropic Gaussian model along with a theoretical prediction for the heat change. The dotted vertical lines differentiate the four phases of the cycle.

As a first example, consider a Gaussian likelihood

$$\mathcal{L} \propto \exp\left(-\frac{1}{2}F_{ij}\theta^i\theta^j\right), \tag{12}$$

with Fisher information F and flat prior $\pi \equiv 1$. For such a model one may compute the partition function analytically (Röver et al., 2023a)

$$Z(T, J) = \sqrt{\frac{(2\pi T)^n}{\det F}} \exp\left(\frac{1}{2}TF^{ij}J_iJ_j\right), \tag{13}$$

with the components F^{ij} of the inverse Fisher matrix such that one finds the entropy (9)

$$S(T, J) = \frac{n}{2}\log(2\pi T) + \log(\det F)\,, \tag{14}$$

and thus the caloric coefficients $C(T, J) = n/2$ and $h^i(T, J) = 0$ (Röver et al., 2023a). The theoretical rate at which the heat is exchanged with the system is

$$\dot{Q} = \frac{n}{2}\dot{T}\,. \tag{15}$$

This result means that during the iso-J phase of the cycle the amount $\Delta Q = \frac{n}{2}(T_h - T_c)$ is exchanged with the environment. For both phases, it is the same amount with a flipped sign, implying $\Delta Q_{\text{net}} = 0$. And since $\Delta W_{\text{net}} = \Delta U_{\text{net}} - \Delta Q_{\text{net}} = 0$, there is also no work performed by the system. The inner energy is the same after the cycle as before, since the system returns to its initial state.

To test whether this theoretical prediction holds in practice, we run multiple cycles with this system. In this numerical experiment in two dimensions we set $F = \mathrm{diag}(1, 10)$, $J_1 = (5.5, 0)$, $J_2 = (5, 0)$, $T_c = 1.0$, $T_h = 2.0$.

In Figure 2, the trajectories of temperature and sources are plotted. One may observe that the exponential temperature scheduler does indeed deserve its name. While the iso-J phases look quite exponential in the temperature T, in contrast to this, the isothermal trajectories look rather linear in J. This is because the average gradient $\langle \nabla_J V \rangle_{\text{chains}}$ is too small such that the scheduler almost always takes the maximum step size $\Delta J_{\max}$.

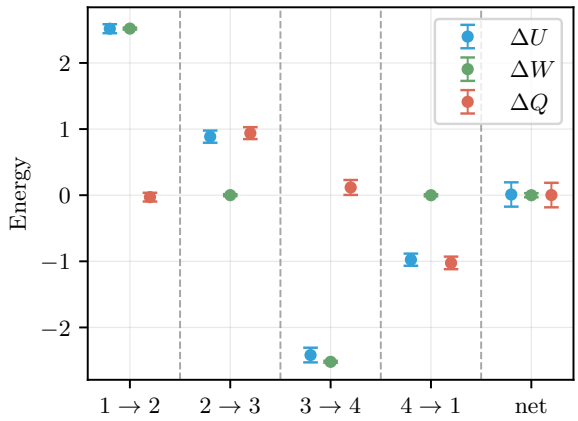

Figure 4: Total energy changes per phase of an MCMC cycle in an anisotropic Gaussian model. The results are averaged over 10^3 chains and 10 cycles.

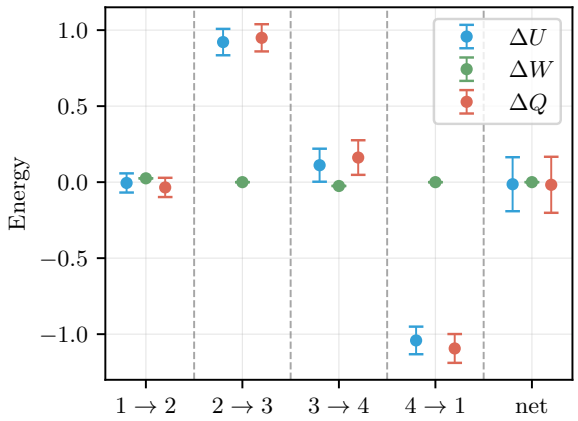

Figure 5: Results of an MCMC cycle in an anisotropic Gaussian model with different target sources. The performed work in phases $(1 \to 2)$ and $(3 \to 4)$ is much less than before.

In Figure 3 we plot the observed energy changes $\Delta U, \Delta W, \Delta Q$ along with the theoretical predictions ΔQ_{theory} made above. To reduce the considerable amount of noise we average over 10^3 chains. The work ΔW has much less noise than the energy change ΔU, since for an individual chain $\dot{W} \propto \theta \cdot \dot{J}$ (where $\dot{J}$ is the same for all chains) while $\dot{U} \propto \theta \cdot \dot{\theta}$. Thus we expect the noise on the energy change to be much larger compared to the noise of the work change.

It is clearly visible that $\Delta Q = 0$ in the isothermal phases, while $\Delta W = 0$ in the iso-J phases, confirming the theoretical predictions. Also, the noise of the energy change ΔU is greater for $T = T_h$ compared to $T = T_c < T_h$, which also makes sense since the sampling particles have more freedom to oscillate around the expectation value of θ in this phase. For comparison, we also plot the theoretical predictions (15) for the heat change ΔQ_{theory}, which nicely agrees with the observations. The fact that $\Delta U = 0$ after a full cycle is to be expected for any model since the system returns to its initial state. Additionally, this plot conveys that the system is not performing any net work, since $\Delta W = \Delta Q = 0$ after a full cycle.

The same phenomenon may be observed in Figure 6. On the right, we plot the temperature T against the entropy $S = Q/T$ (again we smooth over 10^3 cycles and compute a moving average to reduce noise). Since the net heat exchange over one cycle is given by $\Delta Q_{\text{net}} = \oint T \mathrm{d}S$, it is given in the plot as the area enclosed by the trajectory. On the left of Figure 6, we made a similar plot for the quantity $\Lambda = \nabla_J H$ against the sources J. Since the scheduled source $\sigma_J(t)$ always lies on a line connecting J_1 and J_2, we equivalently consider the projection of J and Λ on their difference $J_2 - J_1$. Again one observes that the net work $\Delta W_{\text{net}} = \oint \Lambda \cdot \mathrm{d}J$ vanishes. In Figure 4 all energy measurements concerning this experiment are summarized. Again, the higher temperature in phase $(3 \to 4)$ leads to larger fluctuations and the net performed work and exchanged heat are zero within tolerance.

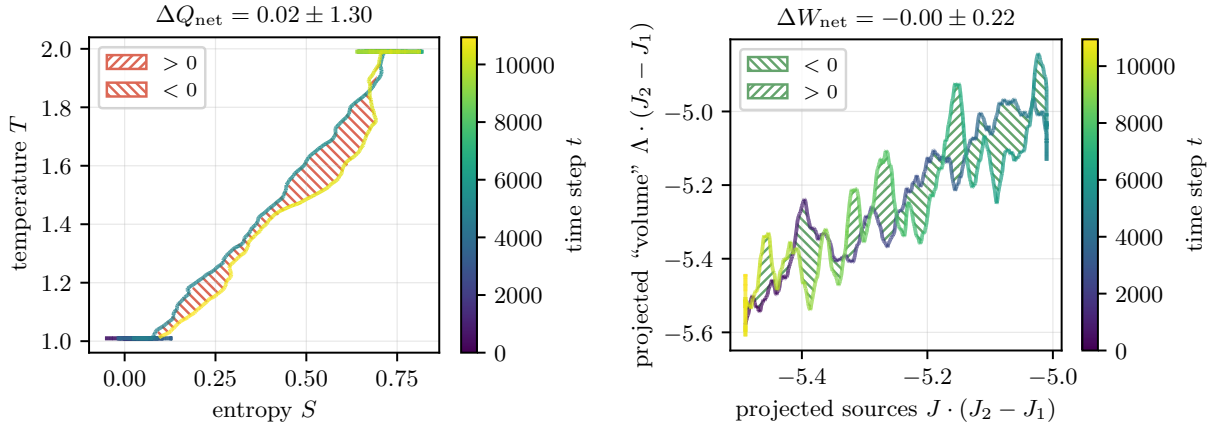

Figure 6: Illustration of the net exchanged heat (left) and net performed work (right) during one cycle. Since the heat and work forms are given by $\delta Q = T\mathrm{d}S$ and $\delta W = \Lambda \cdot \mathrm{d}J$ respectively, the areas enclosed by the trajectory correspond to the net exchanged heat and net performed work during one cycle. To reduce noise, the data are smoothed with a moving average with periodic boundary conditions of window size 100.

5.2 Gaussian Toy Model with different cycle trajectory

As a second example, we consider the same model as above (12), but this time we choose the sources $J_1 = (0, 5.5)$, $J_2 = (0, 5)$ to lie on a different axis. As above, the Gaussian is anisotropic with the covariance matrix $F = \mathrm{diag}(1, 10)$ such that the typical set has the shape of an ellipsoid which is elongated in θ_1-direction. In the previous example, we chose the sources along the elongated direction (J in θ_1). This time, they are pointed across it (J in θ_2).

The consequence of this may be observed in Figure 5, which corresponds to Figure 4 for the above experiment. The work performed in the isothermal phases $(1 \to 2)$ and $(3 \to 4)$ of the cycle is much smaller. This was to be expected since the system has to work against much less resistance in this case as the potential varies less in this direction.

5.3 Rosenbrock Toy Model

As an example of a non-linear and consequently non-Gaussian model, we consider the Rosenbrock likelihood

$$\mathcal{L} \propto \exp\left(-\frac{1}{2}\left((a\theta^0)^2 + (b(\theta^1 - (c\,\theta^0)^2))^2\right)\right), \tag{16}$$

with $a = 2$, $b = 0.3$ and $c = 5$ and a flat prior $\pi \equiv 1$. In this experiment we set $J_1 = (6, 0)$, $J_2 = (5, 0)$, $T_c = 1.0$, $T_h = 3.0$. In contrast to the Gaussian case above, the partition function may only be approximated this time. We employ a perturbative method (Röver et al., 2023a) to find

$$Z(T, J) = Z_{\text{G}} Z_{\text{NG}} :\approx \sqrt{\frac{(2\pi T)^n}{\det F}} \exp\left(\frac{1}{2} T F^{ij} J_i J_j\right) \cdot \left(1 - \frac{c^2}{4!a^2 b} J_2 + \frac{3b^2c^2}{96a^6}(J_1)^2 + \frac{3b^2c^2}{96a^4} T - \frac{c^2}{4a^4 b}\frac{1}{T^2}(J_1)^2 J_2\right). \tag{17}$$

From this, one may deduce the entropy (9)

$$S(T, J) = \frac{n}{2}\log(2\pi T) + \log(\det F) + \log(Z_{\text{NG}}) + T\frac{\partial \log(Z_{\text{NG}})}{\partial T}, \tag{18}$$

and subsequently the caloric coefficients $C(T, J)$ and $h^i(T, J)$. Since the coefficients in the higher-order terms of Z_{NG} are quite small, the overall prediction of the heat rate $\dot{Q}$ is dominated by the Gaussian part. Still,

since the entropy has an explicit dependence on both J and T through Z_{NG} in contrast to the pure Gaussian case, it is to be expected that the exchanged heat during the iso-J phases depends on J, thus $\Delta Q_{12} \neq -\Delta Q_{34}$. This means that $\Delta Q_{\mathrm{net}} \neq 0$ and consequently a non-zero amount of work $\Delta W_{\mathrm{net}} \neq 0$ is performed during a cycle. Please note that this perturbative approximation has limited validity (depending on the magnitude of c); it is, however, sufficient to demonstrate the point made above.

After a series of cycles with this model, the data are analyzed as above for the Gaussian case and the results are summarized in Figure 7. While the plots display similarities to the former case, all of them convey that the net performed work as well as the net heat input are non-zero this time. While the overall change in inner energy is zero after every cycle, the performed work rises gradually while heat input declines equivalently after every cycle.

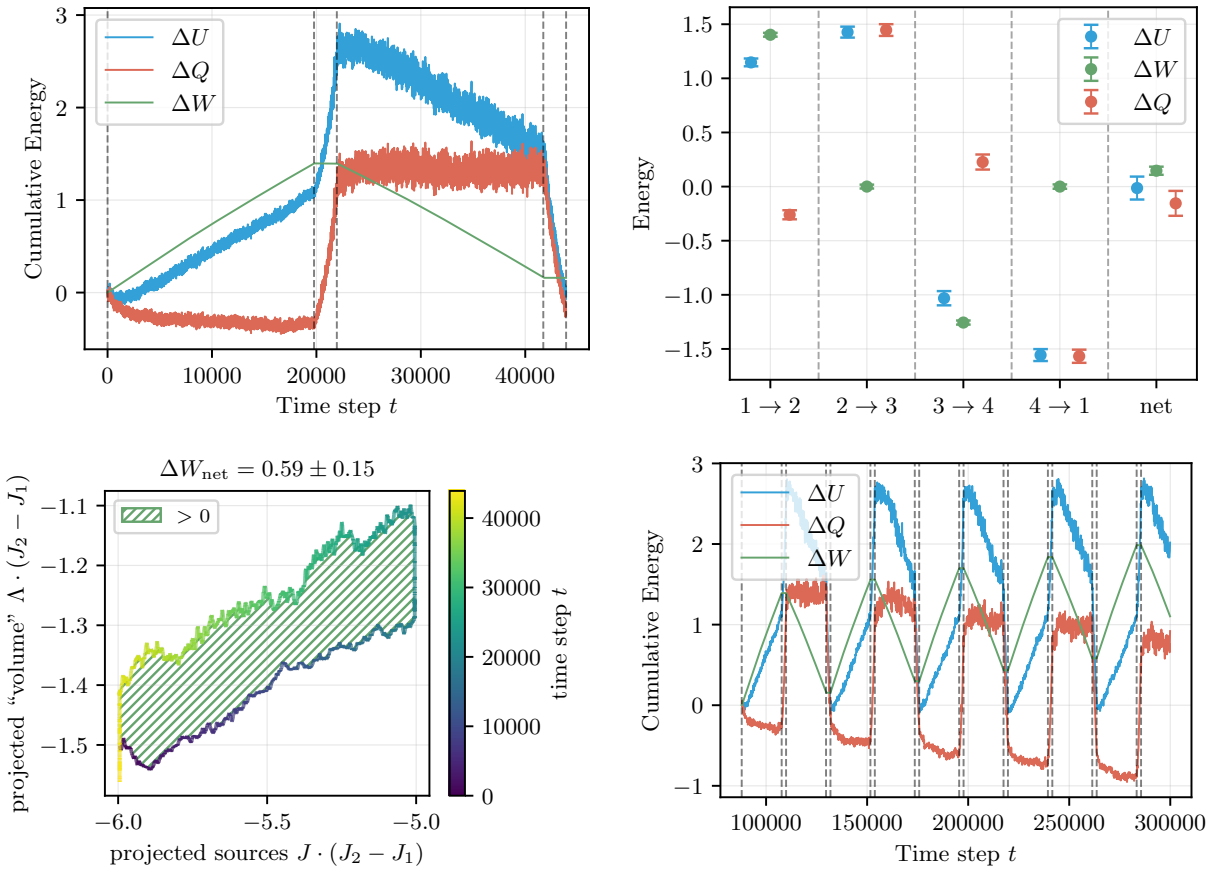

Figure 7: Results for the Rosenbrock model corresponding to Figures 3, 6 and 4 for the anisotropic Gaussian model. In contrast to before, a non-zero amount of net work is performed by the system.

To further study this phenomenon, we ran this experiment for different choices of the parameter c in the Rosenbrock likelihood. In Figure 8, the net performed work is plotted against it. At $c = 0$ the model is Gaussian, thus the output is zero within tolerance. The greater the parameter becomes, the more work is produced.

The effect of anharmonicities in an otherwise quadratic potential is well known in solid state physics, where it explains the thermal expansion of solids (Ziman, 1979). A rudimentary realisation of a heat engine might be a solid that is subjected to stresses and changes in temperature, emulating a cycle in J and T. If the atoms were bound in a mutual potential of quadratic shape, they would show vibrations about the equilibrium position. In an anharmonic potential, however, the centre of vibration would be shifted away from the minimum of the potential. This effectively increases the average distance between the atoms, leading to thermal expansion. From a thermodynamic point of view, the minimum in free energy determines the average distance between

the atoms. Then, thermal expansion can perform work against external stresses, in contrast to the case of a quadratic potential, where thermal expansion does not occur and mechanical work is not performed. This completes the argument why a quadratic potential representing Gaussian distributions is ineffective as a heat engine and leads to an efficiency of zero.

Effectively, the mechanical work performed in a Stirling-type MCMC cycle reflects the deviation from Gaussianity of the posterior distribution, and the efficiency of the cyclic MCMC engine serves as a measure of non-Gaussianity that is independent from the assumption of any particular functional shape.

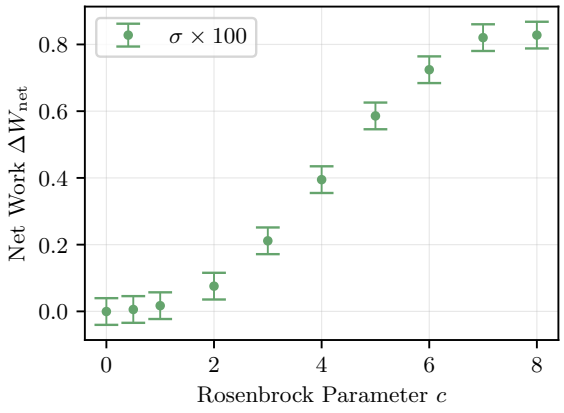

Figure 8: Net performed work as a function of the parameter c in the Rosenbrock model. The more non-Gaussian the model becomes, the greater is the work output.

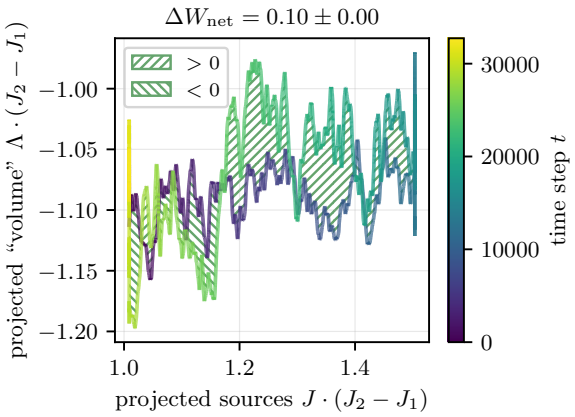

Figure 9: Results of an MCMC cycle for the cosmological inference problem. The performed net work is non-zero, indicating a non-Gaussian posterior.

5.4 Non-Gaussian cosmological inference problem

As a topical Bayesian inference problem we consider the determination of cosmological parameters in supernova cosmology (Brout et al., 2022), where the relationship between redshift z and apparent magnitudes m is probed. The Hubble function for a dark energy-dominated Universe is given by

$$H^2(z) = H_0^2 \left(\Omega_m (1+z)^3 + (1-\Omega_m)(1+z)^{3(1+w)} \right) . \tag{19}$$

from which the luminosity distance is obtained numerically by solving

$$\frac{\mathrm{d}d_L}{\mathrm{d}z} = \frac{d_L}{1+z} + \frac{c(1+z)}{H(z)} , \tag{20}$$

from which one may then compute the prediction for the observed apparent magnitude,

$$m_{\text{pred}} = M + 5 \log_{10} d_L + 10 , \tag{21}$$

with the absolute magnitude M as a nuisance parameter. The likelihood is Gaussian in the data,

$$\mathcal{L} \propto \exp \left(-\frac{1}{2} (m^i_{\text{obs}} - m^i_{\text{pred}}(\theta)) C_{ij} (m^j_{\text{obs}} - m^j_{\text{pred}}(\theta)) \right) , \tag{22}$$

with the inverse data covariance C_{ij}. To make this application comparable to the two examples above, we profile the likelihood, fixing the absolute magnitude $M = -19.25$ and the Hubble constant $H_0 = 73.28\,\mathrm{km/s/Mpc}$, thus obtaining a two-dimensional inference problem with $\theta = (\Omega_m, w)$.

Since the above likelihood is quite expensive to evaluate, we ran the experiment with a Gaussian surrogate first to tune the hyperparameters. For this, we approximated the likelihood around the maximum likelihood

estimate as

$$\mathcal{L} \propto \exp\left(-\frac{1}{2}(\theta^i - \theta^i_{\mathrm{MLE}})F_{ij}(\theta_{\mathrm{MLE}})(\theta^j - \theta^j_{\mathrm{MLE}})\right) , \tag{23}$$

with the Fisher information F. For the Bayesian RMH cycle, we chose $J_1 = (0.1, 1.0)$, $J_2 = (0.15, 1.5)$, $T_c = 1.0$, $T_h = 5.0$. We then used the full likelihood to produce the results in Figure 9. We see that the system has again produced a positive amount of net work. From this we conclude that the model is non-linear in its parameters, and consequently shows a non-Gaussian shape of the posterior distribution.

6 Conclusion

In a proof of concept, this paper demonstrates the feasibility of cyclic processes for MCMC algorithms. Our main technical contribution is the adaptive ensemble schedulers that facilitate a tunable version of the Rosenbluth-Metropolis-Hastings (RMH) algorithm. We apply the notion of quasi-staticity from thermodynamics to ensure that the sampling algorithm is affected at timescales large enough to ensure it remains in thermal equilibrium throughout any change of state. With these schedulers we implement an MCMC cycle for a Bayesian canonical ensemble made up of four phases alternately varying the temperature T and the sources J inspired by the Stirling cycle from thermodynamics. Such cycles are possible for any kind of statistical model, and generalise numerical methods such as simulated annealing by (i) introducing cyclic changes of state and by (ii) altering additional parameters, in our case control parameters used in the determination of cumulants or moments of the posterior distribution.

We first test them on an anisotropic Gaussian, confirming all predictions made by a corresponding thermodynamic theory. The most interesting fundamental insight surely is that the Bayesian canonical ensemble produces net work output if and only if the model under consideration is non-linear or non-Gaussian. This was found by comparing the Gaussian example to that of a Rosenbrock likelihood, and further in an example for Bayesian inference from cosmology with a topical data set. As such, the work performed in a cycle by a Bayesian engine serves as a measure of non-Gaussianity of the inferred posterior distribution, independent of its particular functional form. We give an example from solid state physics with an analogous phenomenology.

From a conceptual point of view, we follow the idea of Jaynes (1957) that thermodynamics at its core is a theory of information, and establish a thermodynamic interpretation of all quantities that appear in a Bayesian inference problem: These concern the partition functions, their associated thermodynamic potentials, their state variables and finally heat exchanged and mechanical work performed, for the definition of a thermodynamic efficiency.

The most obvious idea for further study is to apply these MCMC cycles to a wider range of statistical models. A more technical improvement would be to use more advanced sampling algorithms such as Hamiltonian Monte Carlo. Since it explores the typical set more quickly, it would adapt faster to changes in the parameters T and J, yielding a considerably more efficient way of performing MCMC cycles.

In conclusion, the application of concepts from statistical physics to canonical MCMC methods has again produced both fruitful technical and fundamental insights, providing a new angle on their thermodynamic nature.

References

Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng (eds.). *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC, may 2011. doi: 10.1201/b10905. URL `https://doi.org/10.1201%2Fb10905`.

Dillon Brout, Dan Scolnic, Brodie Popovic, Adam G Riess, Anthony Carr, Joe Zuntz, Rick Kessler, Tamara M Davis, Samuel Hinton, David Jones, et al. The pantheon+ analysis: cosmological constraints. *The Astrophysical Journal*, 938(2):110, 2022.

Sadi Carnot. Réflexions sur la puissance motrice du feu et sur les machines propres à développer cette puissance. In *Annales scientifiques de l'École normale supérieure*, volume 1, pp. 393–457, 1872.

Rudolf Clausius. *The mechanical theory of heat*. Macmillan, 1879.

Michael Creutz. Monte carlo study of quantized su (2) gauge theory. *Physical Review D*, 21(8):2308, 1980.

Simon Duane, Anthony D Kennedy, Brian J Pendleton, and Duncan Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987.

Alan E Gelfand and Adrian FM Smith. Sampling-based approaches to calculating marginal densities. *Journal of the American statistical association*, 85(410):398–409, 1990.

Stuart Geman and Donald Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*, PAMI-6(6):721–741, 1984.

W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57(1):97–109, 04 1970. ISSN 0006-3444. doi: 10.1093/biomet/57.1.97. URL https://doi.org/10.1093/biomet/57.1.97.

Maximilian Philipp Herzog, Heinrich von Campe, Rebecca Maria Kuntz, Lennart Röver, and Björn Malte Schäfer. Partition function approach to non-Gaussian likelihoods: macrocanonical partitions and replicating Markov-chains. *The Open Journal of Astrophysics*, 7, 2024. doi: 10.33232/001c.125132. URL https://doi.org/10.33232/001c.125132.

E. T. Jaynes. Information theory and statistical mechanics. *Physical Review*, 106(4):620–630, 1957. ISSN 0031-899X. doi: 10.1103/PhysRev.106.620. URL https://link.aps.org/doi/10.1103/PhysRev.106.620.

Rebecca Maria Kuntz, Maximilian Philipp Herzog, Heinrich von Campe, Lennart Röver, and Björn Malte Schäfer. Partition function approach to non-gaussian likelihoods: partitions for the inference of functions and the fisher-functional. *Monthly Notices of the Royal Astronomical Society*, 527(3):8443–8458, 2024.

Rebecca Maria Kuntz, Heinrich von Campe, Tobias Röspel, Maximilian Philipp Herzog, and Björn Malte Schäfer. Partition function approach to non-Gaussian likelihoods: information theory and state variables for Bayesian inference. *The Open Journal of Astrophysics*, 8, mar 5 2025. doi: 10.33232/001c.131858.

Antony Lewis and Sarah Bridle. Cosmological parameters from CMB and other data: a monte-carlo approach. *PRD*, 66(10), 2002. ISSN 0556-2821, 1089-4918. doi: 10.1103/PhysRevD.66.103511. URL http://arxiv.org/abs/astro-ph/0205436.

Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.

Michael J Moran, Howard N Shapiro, Daisie D Boettner, and Margaret B Bailey. *Fundamentals of engineering thermodynamics*. John Wiley & Sons, 2010.

Lennart Röver, Lea Carlotta Bartels, and Björn Malte Schäfer. Partition function approach to non-gaussian likelihoods: Formalism and expansions for weakly non-gaussian cosmological inference. *Monthly Notices of the Royal Astronomical Society*, 523(2):2027–2038, 2023a.

Lennart Röver, Heinrich von Campe, Maximilian Philipp Herzog, Rebecca Maria Kuntz, and Björn Malte Schäfer. Partition function approach to non-Gaussian likelihoods: physically motivated convergence criteria for Markov chains. *Monthly Notices of the Royal Astronomical Society*, 526(1):473–482, 09 2023b. ISSN 0035-8711. doi: 10.1093/mnras/stad2726. URL https://doi.org/10.1093/mnras/stad2726.

Wayne M Saslow. A history of thermodynamics: The missing manual. *Entropy*, 22(1):77, 2020.

William Thomson. Lxvii. on the dynamical theory of heat. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 4(27):424–434, 1852.

P.J.M. van Laarhoven and E.H.L. Aarts. *Simulated annealing : theory and applications*. Mathematics and its applications. Reidel, 1987. ISBN 90-277-2513-6.

Liwei Wang, Xinru Liu, Aaron Smith, and Aguemon Y Atchade. On cyclical mcmc sampling. In *International Conference on Artificial Intelligence and Statistics*, pp. 3817–3825. PMLR, 2024.

J. M. Ziman. *Principles of the Theory of Solids*. Cambridge University Press, 1979.