A generic thermostat for smooth vector fields and smooth target densities

FRANK, J. E.[1], LEIMKUHLER, B. J.[2], and MYERSCOUGH, K. W.[3]

1 Mathematical Institute, Utrecht University, P.O. Box 80010,
3508 TA Utrecht, THE NETHERLANDS
j.e.frank@uu.nl

2 School of Mathematics, University of Edinburgh, James Clerk Maxwell Building,
Kings Buildings Edinburgh EH9 3JZ, SCOTLAND
b.leimkuhler@ed.ac.uk

3 Department of Computer Science, KU Leuven, 3001 Leuven, BELGIUM
keith@myerscough.nl

Thermostat methods are routinely employed in molecular dynamics to simulate a system of particles at constant temperature. Molecular dynamics models are typically formulated as a classical mechanical n-body problem in high dimensions. The equations of motion constitute a Hamiltonian system

$$\frac{dy}{dt} = J\nabla H(y), \qquad H : \mathcal{R}^d \to \mathcal{R}, \quad J = -J^T, \tag{1}$$

with preserved total energy H. The motion is thus constrained to a surface of constant H (or the intersection of the level sets if more conserved quantities are present). If the motion is ergodic, then the flow samples the invariant measure

$$\mu(y) \propto \delta(H(y) - H_0), \qquad H(y(0)) = H_0.$$

On the other hand, a system that evolves in thermal equilibrium with respect to a large temperature reservoir of inverse temperature β does not evolve at constant energy. Instead, the states of the system are distributed according to the Gibbs canonical distribution

$$\mu(y) \propto \exp(-\beta H(y)).$$

To simulate a system at constant temperature, it is necessary to introduce some dynamic mechanism to perturb trajectories such that they ergodically sample the canonical distribution.

One well-known technique for canonical sampling is Langevin dynamics [3]. Here the Hamiltonian equations on $\mathcal{R}^d$ are equipped with a stochastic diffusion process. An alternative approach are the Nosé-Hoover type methods [10, 11, 6, 2, 12] where the phase space is augmented by one or more additional thermostat variables such that the projected motion on $\mathcal{R}^d$ again ergodically samples the canonical distribution. Construction of the

Nosé-Hoover dynamics makes explicit use of two properties: (1) the original dynamics is Hamiltonian (specifically, divergence free and Hamiltonian-conserving), and (2) the target distribution (Gibbs) is a smooth function of the conserved quantity H.

In a recent paper [9] we show how information theory can be used in combination with Nosé-Hoover type methods to correct dynamics for observations of the mean values of conserved quantities. Again, the approach of [9] is restricted to unperturbed systems with Hamiltonian structure, and observables that are functions of the conserved quantities. The methodology could be made significantly more generic if these restrictions on the dynamics and probability distributions could be removed. The purpose of this note is to describe a method that provides for this.

Specifically, we derive a thermostat that can be applied to an arbitrary, smooth differential equation to perturb its orbits such that they ergodically sample a generic, smooth target distribution. The target distribution can in principle be any distribution of the form $\rho(y) \sim \exp(-A(y))$, where $A : \mathcal{R}^d \to \mathcal{R}$ is bounded and differentiable. However, the thermostat is most effective when this distribution is 'close' in some sense to the invariant distribution of the unperturbed dynamics. In the §1 we review briefly an information theoretic approach to correcting a prior distribution for a set of observed expectations. In §2 we describe the new thermostat. In §2.1 we discuss ergodicity considerations. The new thermostat is ineffective in the classical setting of a Hamiltonian system and Gibbs distribution. Therefore, in Section §2.2 we describe necessary modifications for this case. Finally we demonstrate the new thermostats for some simple examples in Section §3.

1 BAYESIAN MODELLING

For the purpose of this section, suppose $y \in \mathcal{R}^d$ is a random variable with distribution (law) $y \sim \rho$, where $\rho : \mathcal{R}^d \to \mathcal{R}$ is unknown. Suppose further, that we are given a prior distribution $\pi : \mathcal{R}^d \to \mathcal{R}$, assumed to be close to ρ.

The Kullback-Leibler divergence, or relative entropy,

$$\mathcal{S}[\rho(y)] = \int \rho(y) \ln \frac{\rho(y)}{\pi(y)}\, dy$$

represents a (non-symmetric) distance between measures. In information theory it gives the information lost in approximating $\rho(y)$ by $\pi(y)$.

Next, suppose we are given a set of K observations of y in the form of expectations

$$\mathbb{E}_\rho C_k(y) = \int C_k(y)\rho(y)\, dy = c_k, \quad k = 1, \ldots, K. \tag{2}$$

Then the least biased distribution ρ consistent with the observations c_k and prior π is given by the solution of the constrained minimization problem

$$\rho = \arg\min_\rho\ \mathcal{S} - \lambda_0 \left(1 - \int \rho(y)\, dy\right) - \sum_{k=0}^{K} \lambda_k \left(c_k - \int C_k(y)\rho(y)\, dy\right),$$

where the λ_k are Lagrange multipliers associated with the observations (2). Solving the minimization problem is an exercise in variational calculus. One finds

$$\rho(y) = \lambda_0 \exp\left(-\lambda_1 C_1(y) - \cdots - \lambda_K C_K(y)\right) \pi(y), \tag{3}$$

where the λ_k are chosen such that the observations (2) are satisfied.

2 THERMOSTATS FOR THE POSTERIOR MEASURE

Now suppose that we are given a dynamical system, defined by the solution of a differential equation,

$$\frac{dy}{dt} = F(y),$$

which may be subject to model error. Further suppose we are given a prior distribution $\pi(y)$ that we believe to be close to the invariant distribution of the true dynamics, and a set of K observations of the system of the form (2). We construct a thermostat on the dynamics of y that samples the posterior distribution (3). To do this, let us write $\rho(y) = \exp(-A(y))$, and define the extended distribution $\hat{\rho}(y, \xi) = \rho(y) \exp(-\xi^2/2)$. Then we consider a thermostat of the form

$$dy = F(y)\, dt + \xi^2 G(y)\, dt \tag{4}$$

$$d\xi = \xi X(y)\, dt - \gamma \xi\, dt + \sqrt{2\gamma}\, dw, \tag{5}$$

where $\gamma > 0$ is a diffusion parameter.

The distribution $\hat{\rho}$ is stationary under the Fokker-Planck equation associated to this system if

$$\mathcal{L}^* \hat{\rho} = 0 = -\nabla \cdot \hat{\rho} F(y) - \xi^2 \nabla \cdot \hat{\rho} G(y) - \partial_\xi \xi \hat{\rho} X(y), \tag{6}$$

since additional terms in the Fokker-Planck operator cancel automatically due to fluctuation-dissipation balance in the last two terms of (5) (an Ornstein-Uhlenbeck process). A possible solution of this equation is given by

$$X(y) = F \cdot \nabla A - \nabla \cdot F, \qquad \nabla \cdot (F + G) - (F + G) \cdot \nabla A = 0. \tag{7}$$

Hence, defining $X(y)$ by the first condition above and choosing a G to satisfy the second condition ensures stationarity of $\hat{\rho}$.

One possible choice (which we will not use) for G is $G = J\nabla A - F$, where J is any skew-symmetric matrix. Intuitively, since $\mathbb{E}_{\hat{\rho}} \xi^2 = 1$, this choice just replaces the dynamics F with the Hamiltonian dynamics $J\nabla A$ on average. This can obviously have dire consequences for the thermostated dynamics, unless the vector field $G(y)$ is small in some sense.

Having found a G that satisfies the above condition, any other vector field $\tilde{G} = G + B\nabla A$ for any skew-symmetric matrix B also satisfies the condition. This can be used to find an optimal skew-symmetric B, for instance, such that the norm of $\tilde{g}$ is minimized.

2.1 Ergodicity

In the previous section we have formally constructed a dynamics under which the target distribution is stationary. It is also necessary to prove that this distribution is unique and attracting. Because the distributions we consider have global support, we will see that it is sufficient to show a Hörmander condition on the vector fields F and G (see related proofs in [1, 8]). Establishing this condition is problem dependent.

By assumption the desired density $\hat{\rho}(y) > 0$ for all y. Since $\hat{\rho}$ is stationary under the Fokker-Planck operator, ergodicity of $\hat{\rho}$ can be established under the ergodic decomposition theorem if the Hörmander condition holds [7, 4, 5]. Consider the deterministic and stochastic vector fields

$$U(y,\xi) = \begin{pmatrix} F(y) + \xi^2 G(y) \\ \xi X(y) - \gamma\xi \end{pmatrix}, \quad V(y,\xi) = \begin{pmatrix} 0 \\ \sqrt{2\gamma} \end{pmatrix}$$

The Hörmander condition requires that the Lie algebra generated by U and V span $\mathcal{R}^{n+1}$:

$$\mathcal{R}^{n+1} \subset \text{Lie}\{U,V\} = \text{span}\{U, V, [U,V], [U,[U,V]], [V,[U,V]], \dots\}.$$

Let us suppose that $G(y)$ is chosen such that the vector fields $F(y)$ and $G(y)$ satisfy the Hörmander condition on $\mathcal{R}^n$. Define vector fields $\hat{F} = (F(y), 0)$, $\hat{G} = (G(y), 0)$ in $\mathcal{R}^{n+1}$. We show that

$$\mathcal{R}^{n+1} \subset \text{Lie}\{\hat{F}, \hat{G}, e_{n+1}\} \subset \text{Lie}\{U,V\},$$

where $e_{n+1} = (0,\dots,0,1)$ is a canonical unit vector in the auxiliary variable direction. The first inclusion follows from the Hörmander condition on $\mathcal{R}^n$ and is immediate. Since $\gamma > 0$, it follows that V is proportional to e_{n+1}. We compute

$$U_1 = [U, e_{n+1}] = \begin{pmatrix} 2\xi\,[G(y) - F(y)] \\ X(y) - \gamma \end{pmatrix},$$

and

$$U_2 = \frac{1}{2}[U_1, e_{n+1}] = \begin{pmatrix} G(y) - F(y) \\ 0 \end{pmatrix} = \hat{G} - \hat{F}.$$

Next, define

$$V_1 = U - \frac{\xi^2}{2} U_1 - \frac{\xi}{2}(X(y) - \gamma)e_{n+1} = \begin{pmatrix} F(y) \\ 0 \end{pmatrix} = \hat{F}$$

Clearly, U_2 and V_1 are contained in Lie$\{U,V\}$, as are their higher order commutators. But $V_1 = \hat{f}$ and $U_2 + V_1 = \hat{g}$, combined with e_{n+1}, form the basis for the intermediate Lie algebra, from which the inclusion follows.

2.2 A double thermostat for Hamiltonian systems

The approach of the previous section can fail in the standard canonical thermostating situation when the vector field $F = J\nabla H$ is divergence-free and the posterior measure is the Gibbs measure, i.e. $A(y) = \beta H(y)$. Here it can be checked that $X(y)$ in (7) is identically zero, and hence there is no feedback. We can extend the above approach with a Nosé-Hoover-Langevin thermostat to ensure that in the absence of observations, the system samples a prior $\pi \propto \exp(-\beta H(y))$. To do so, let us take $F(y) = J\nabla H(y)$, $A(y) = \beta H(y) + \lambda C(y)$, and introduce a second auxiliary variable η, with dynamics

$$dy = F(y)\,dt + \eta g(y)\,dt + \xi^2 G(y,\eta)\,dt, \tag{8}$$

$$d\eta = (\nabla \cdot g(y) - g(y) \cdot \nabla A(y))\,dt - \gamma_H \eta\,dt + \sqrt{2\gamma_H} dw_H, \tag{9}$$

$$d\xi = \xi X(y,\eta)\,dt - \gamma_A \xi\,dt + \sqrt{2\gamma_A}\,dw_2. \tag{10}$$

It can be checked that the composite measure $\rho \propto \exp(-\beta H(y) - \lambda C(y) - \eta^2/2 - \xi^2/2)$ is stationary under the associated Fokker-Planck equation if we define $X(y,\eta)$ by

$$X(y,\eta) = \nabla \cdot G(y,\eta) - G(y,\eta) \cdot \nabla A(y),$$

and ensure that $G(y,\eta)$ satisfies

$$\nabla \cdot G(y,\eta) - G(y,\eta) \cdot \nabla A - \lambda f(y) \cdot \nabla C.$$

We give an example below.

3 NUMERICAL EXPERIMENTS

In this section we present some specific examples.
Example 1 Consider a Harmonic oscillator, $y \in \mathcal{R}^2$,

$$y' = F(y) = J\nabla H(y), \qquad H(y) = \frac{1}{2}(y_1^2 + y_2^2), \quad J = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix},$$

and suppose we wish to enforce the invariant measure $\rho(y) = \exp(-\frac{1}{2}(d_1 y_1^2 + y_2^2))$ following an observation of the variance of y_1. We use the method (4)–(5). Taking $A = (d_1 y_1^2 + y_2^2)/2$, we may choose G in the direction of the gradient of the observable $C(y) = y_1^2$ by taking

$$G(y) = (\alpha(y_2), 0)^T.$$

Recalling that $\nabla \cdot F \equiv 0$, the function α must satisfy

$$\nabla \cdot G - (F + G) \cdot \nabla A = 0 = (y_2 + \alpha(y_2), -y_1) \cdot (d_1 y_1, y_2) = d_1 y_1 y_2 + \alpha(y_2) d_1 y_1 - y_1 y_2 = 0$$

which we can solve to obtain

$$\alpha(y_2) = \frac{1 - d_1}{d_1} y_2.$$

We expect this to be a minimally intrusive perturbation. Figure 1 illustrates short trajectories for $d_1 = \{0.9, 0.75, 0.5, 0.25\}$. For the case $d_1 = 0.5$, Figure 2 illustrates the histograms of y_1 and y_2. We see that the variances of y_1, $\sigma_1^2 = 1.8$, and y_2, $\sigma_2^2 = 0.9$ are close to the target values of 2 and 1, respectively.

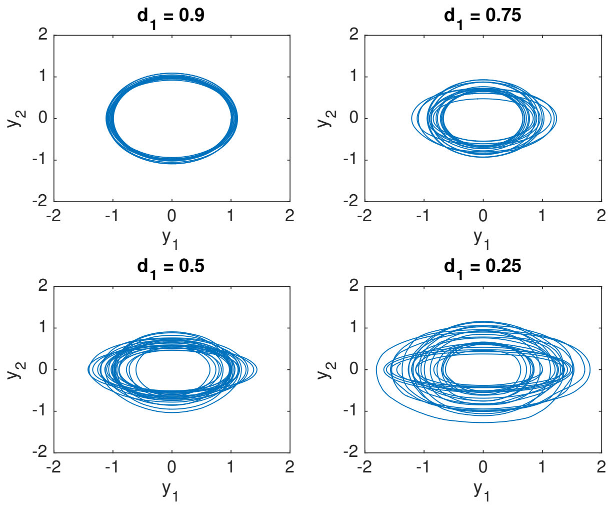

Figure 1: Simulation of harmonic oscillator with $A = (d_1 y_1^2 + y_2^2)/2$ and $\gamma = 0.1$.

Example 2 As a second example, we take a Hamiltonian system in $\mathcal{R}^2$ with double well potential, given by Hamiltonian:

$$H(q, p) = \frac{p^2}{2} + \frac{q^4}{4} - \frac{q^2}{2}. \tag{11}$$

We thermostat this system using (8)–(10). We choose the parameters as follows (note that these satisfy the necessary conditions)

$$g(q, p) = \begin{pmatrix} 0 \\ -p \end{pmatrix}, \quad G(q, p, \eta) = \begin{pmatrix} 0 \\ \frac{\lambda}{\beta}(q - 1) \end{pmatrix}, \quad X(q, p, \eta) = -\gamma p(q - 1).$$

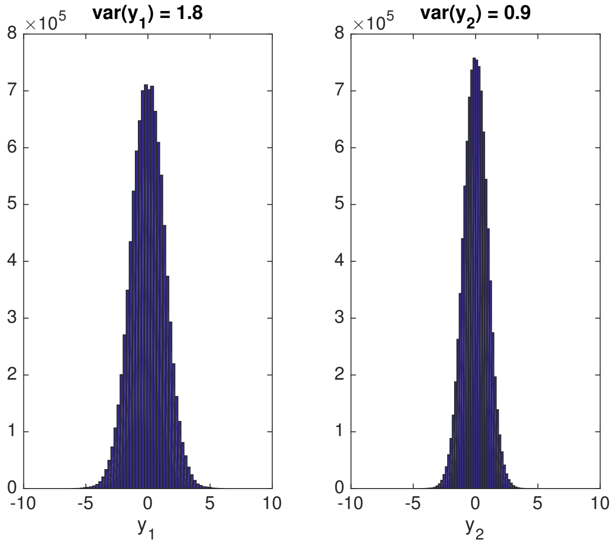

Figure 2: Thermostated harmonic oscillator with $A = (0.5y_1^2 + y_2^2)/2$ and $\gamma = 0.1$. Left: histogram of y_1; right: histogram of y_2.

To sample just the Gibbsian prior distribution we take $\beta = 10$, $\lambda = 0$. We obtain the dynamics and time series labeled *Prior* in Figure 3. The trajectory exhibits transition behavior, spending most of its time in the neighborhood of the fixed points $q = \pm 1$, and occasionally switching between these.

Suppose, now, we enforce the observation $\mathbb{E}(q-1)^2 = 0$. Constructing the posterior distribution as in §1, we take $A(y) = \beta H(y) + \lambda(q-1)^2/2$. In this case, the ratio λ/β can also be thought of as expressing our relative certainty between the prior and posterior distributions π and ρ, or put another way, a measure of the degree of confidence in our observation.

Figure 3 plots the phase trajectory of the dual thermostat (labelled *Posterior*) on top of the canonically thermostated trajectory for the case $\lambda = \beta = 10$. The trajectory now spends all of its time in the potential well around $q = 1$.

References

[1] Bajars, J., Frank, J., and Leimkuhler, B. Weakly coupled heat bath models for Gibbs-like invariant states in nonlinear wave equations. *Nonlinearity*, 26(7):1945, 2013.

[2] Bulgac, A. and Kusnezov, D. Canonical ensemble averages from pseudomicrocanonical dynamics. *Phys. Rev. A*, 42:5045–5048, Oct 1990.

[3] Frenkel, D. and Smit, B. *Understanding Molecular Simulation.* Academic Press, San Diego, second edition edition, 2002.

[4] Hairer, M. Convergence of Markov processes. Lecture notes, University of Warwick, 2010.

[5] Hairer, M. and Mattingly, J. Yet another look at Harris' ergodic theorem for Markov chains. In *Seminar on Stochastic Analysis, Random Fields and Applications VI*, volume 63 of *Progress in Probability*, pages 109–117, 2011.

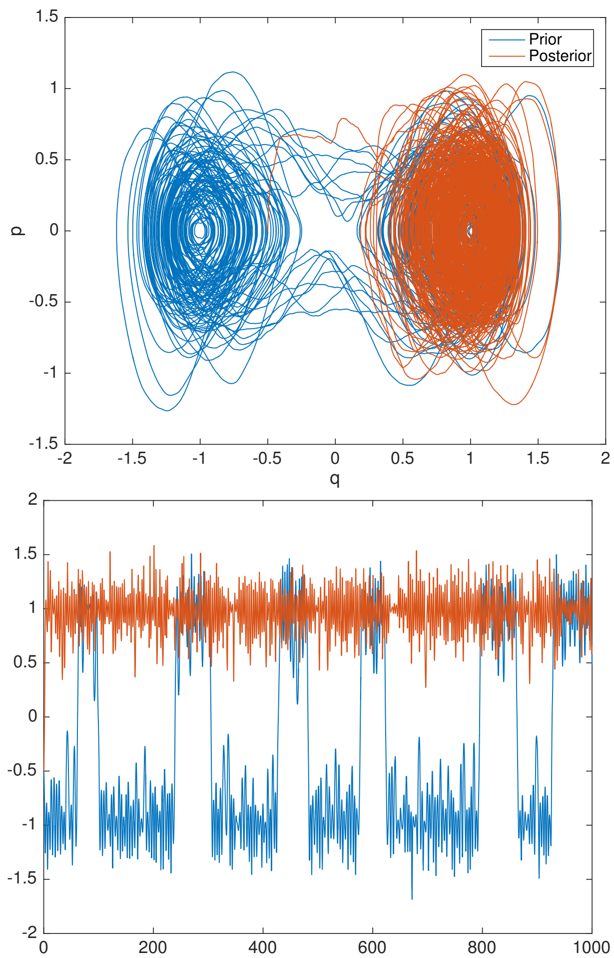

Figure 3: In red, a simulation of the double-well Hamiltonian (11) using the dual thermostat (8)–(9) with parameters $\beta = 10$ and $\lambda = 10$ and observation $\mathbb{E}(q-1)^2 = 0$. We observe that the transition behavior evident in the prior, Gibbs distribution (in blue) is suppressed. Upper plot: phase space orbits; lower plot: time series of $q(t)$.

[6] Hoover, W. Canonical dynamics: equilibrium phase-space distributions. *Physical Review A*, 31(3):1695, 1985.

[7] Kliemann, W. Recurrence and invariant measures for degenerate diffusions. *Annals of Probability*, 15(2):690–707, 1987.

[8] Leimkuhler, B., Noorizadeh, E., and Theil, F. A gentle thermostat for molecular dynamics. *J. Stat. Phys.*, 135(2):261–277, 2009.

[9] Myerscough, K. W., Frank, J., and Leimkuhler, B. Observation-based correction of dynamical models using thermostats. *Proc. R. Soc. A*, 2017.

[10] Nosé, S. A molecular dynamics method for simulations in the canonical ensemble. *Molecular physics*, 52(2):255–268, 1984.

[11] Nosé, S. A unified formulation of the constant temperature molecular dynamics methods. *The Journal of Chemical Physics*, 81(1):511–519, 1984.

[12] Samoletov, A., Dettmann, C., and Chaplain, M. Thermostats for “slow” configurational modes. *Journal of Statistical Physics*, 128(6):1321–1336, 2007.