

Advanced Simulation Methods

Chapter 7 - Hamiltonian Monte Carlo

In this chapter, we are going to study an MCMC method called Hamiltonian Monte Carlo. This method is based on Hamiltonian dynamics, and allows efficient exploration of the state space by making large moves. It is known to be more efficient in high dimensions than Random Walk Metropolis-Hastings and MALA, due to the fact that it avoids random walk behaviour.

1 Hamiltonian mechanics

Let $\mathbf{q} \in \mathbb{R}^d$ denote the position, $\mathbf{p} \in \mathbb{R}^d$ denote the momentum, $U : \mathbb{R}^d \rightarrow \mathbb{R}$ be a continuously differentiable potential function, and $\mathbf{M} \in \mathbb{R}^{d \times d}$ be a symmetric positive definite matrix. Then Hamilton's equations are

$$\frac{d}{dt}\mathbf{q} = \mathbf{M}^{-1}\mathbf{p}, \quad (1)$$

$$\frac{d}{dt}\mathbf{p} = -\nabla U(\mathbf{q}). \quad (2)$$

In the case when $\mathbf{M} = m\mathbf{I}_d$, these equations correspond to Newtonian mechanics with a non-dissipative force arising from the potential field $U(\mathbf{q})$. Equations (1)-(2) are termed a Hamiltonian system, which has *Hamiltonian energy*

$$H(\mathbf{q}, \mathbf{p}) := \frac{\mathbf{p}^T \mathbf{M}^{-1} \mathbf{p}}{2} + U(\mathbf{q}) = K(\mathbf{p}) + U(\mathbf{q}). \quad (3)$$

An important property of the Hamiltonian dynamics is that the Hamiltonian energy is preserved, since

$$\frac{d}{dt}H(\mathbf{q}(t), \mathbf{p}(t)) = \mathbf{p}(t)^T \mathbf{M}^{-1} \frac{d}{dt}\mathbf{p}(t) + \left(\frac{d}{dt}\mathbf{q}(t)\right)^T \nabla U(\mathbf{q}) = 0. \quad (4)$$

It is going to be very useful in our analysis to rewrite Hamiltonian equations in their canonical form

$$\frac{d}{dt}\mathbf{q} = +\nabla_{\mathbf{p}}H(\mathbf{q}, \mathbf{p}), \quad (5)$$

$$\frac{d}{dt}\mathbf{p} = -\nabla_{\mathbf{q}}H(\mathbf{q}, \mathbf{p}). \quad (6)$$

It is also going to be convenient to use the notation $\mathbf{z} := \begin{pmatrix} \mathbf{q} \\ \mathbf{p} \end{pmatrix}$, then $\mathbf{z} \in \mathbb{R}^{2d}$. Let \mathbf{J} be the *canonical structure matrix*

$$\mathbf{J} := \begin{pmatrix} \mathbf{0} & \mathbf{I}_d \\ -\mathbf{I}_d & \mathbf{0} \end{pmatrix}, \quad (7)$$

then the canonical Hamiltonian equations (5)-(6) can be rewritten as

$$\frac{d}{dt}\mathbf{z} = \mathbf{J}\nabla_{\mathbf{z}}H(\mathbf{z}). \quad (8)$$

Example 1 (Harmonic oscillator) Let $\mathbf{L} = \begin{pmatrix} \omega^2 & 0 \\ 0 & 1 \end{pmatrix}$, and $H(\mathbf{z}) = \frac{1}{2}\mathbf{z}^T \mathbf{L} \mathbf{z} = \frac{1}{2}\omega^2 q^2 + \frac{1}{2}p^2$, and the Hamiltonian equations become

$$\frac{d}{dt}\mathbf{z} = \mathbf{J}\nabla_{\mathbf{z}}H(\mathbf{z}) = \mathbf{J}\mathbf{L}\mathbf{z}, \quad (9)$$

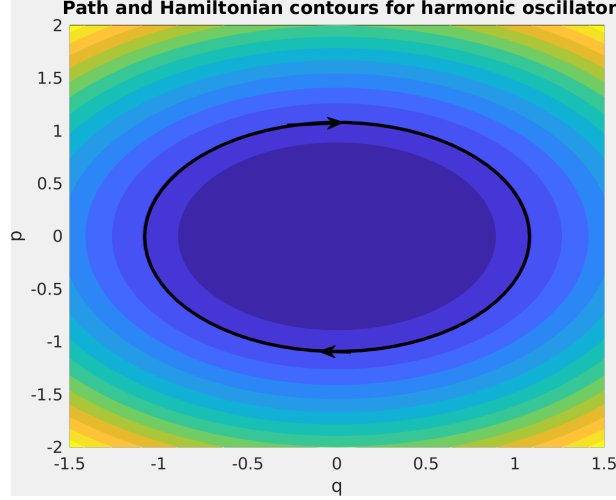
using the particular choice of \mathbf{L} , this is equivalent to

$$\frac{d}{dt}q = p \quad (10)$$

$$\frac{d}{dt}p = -\omega^2 q, \quad (11)$$

which corresponds to the harmonic oscillator. In general, (9) can be defined for arbitrary $2d \times 2d$ symmetric matrix \mathbf{L} , and this is called a linear Hamiltonian system.

The following figure illustrates the evolution of the Harmonic oscillator, and the preservation of the Hamiltonian energy $H(q, p) = \frac{1}{2}\omega^2 q^2 + \frac{1}{2}p^2$, corresponding to an ellipse in the phase space.



2 Symplecticity and other properties

A smooth map $\Psi : \mathbb{R}^{2d} \rightarrow \mathbb{R}^{2d}$ is called *symplectic* if its Jacobian $\nabla_z \Psi(z)$ satisfies that

$$[\nabla \Psi(z)]^T \mathbf{J}^{-1} \nabla \Psi(z) = \mathbf{J}^{-1} \quad (12)$$

for all \mathbf{z} in \mathbb{R}^{2d} , where $\mathbf{J} = \begin{pmatrix} \mathbf{0} & \mathbf{I}_d \\ -\mathbf{I}_d & \mathbf{0} \end{pmatrix}$ is the canonical structure matrix. This satisfies that $\mathbf{J}^{-1} = \mathbf{J}^T$.

An important property of symplectic maps is the volume preservation property.

Proposition 1 *Symplectic maps are volume preserving.*

Proof. It can be shown that the infinitesimal cube $[z_1, z_1 + \delta] \times \dots \times [z_{2d}, z_{2d} + \delta]$ will be mapped by Ψ into a $2d$ dimensional parallelepiped with volume $|\det(\nabla_z \Psi(\mathbf{z}))| \delta^{2d}$. Using the product rule for determinants on the symplectic condition (12), we have

$$\det(\nabla_z \Psi(\mathbf{z}))^2 \det(\mathbf{J}^{-1}) = \det(\mathbf{J}^{-1}),$$

which implies that $|\det(\nabla_z \Psi(\mathbf{z}))| = 1$ since $|\det(\mathbf{J}^{-1})| = 1$. Hence the infinitesimal volumes are preserved, and the same property can be obtained for non-infinitesimal volumes by integration. ■

A key property of Hamiltonian dynamics is that it is symplectic.

Proposition 2 *The Hamiltonian flow defined by (8) is symplectic and volume preserving.*

Proof. Let $\Psi_{t,H}(\mathbf{z})$ denote the flow-map of the Hamiltonian dynamics (8) with Hamiltonian $H(\mathbf{z})$, i.e. $\Psi_{t,H}(\mathbf{z}(0)) = \mathbf{z}(t)$ for a solution of (8). In order to show the symplecticity of the Hamiltonian dynamics, we need to understand the behaviour of the Jacobian $\frac{\partial}{\partial \mathbf{z}} \Psi_{t,H}(\mathbf{z})$. Now we are going to describe the evolution of this Jacobian in time. Let $\mathbf{z}(0)$ be an initial point, and $\bar{\mathbf{z}}(0) = \mathbf{z}(0) + \delta \mathbf{z}(0)$ be another nearby initial point. Then the differences of two paths of the Hamiltonian dynamics at time t can be written as

$$\delta \mathbf{z}(t) := \bar{\mathbf{z}}(t) - \mathbf{z}(t) = \Psi_{t,H}(\bar{\mathbf{z}}(0)) - \Psi_{t,H}(\mathbf{z}(0)) = \frac{\partial}{\partial \mathbf{z}(0)} \Psi_{t,H}(\mathbf{z}(0)) \cdot \delta \mathbf{z}(0) + o(\delta \mathbf{z}(0)),$$

and by differentiating the right hand side in t , we obtain that

$$\begin{aligned} \frac{d}{dt} \left[\frac{\partial}{\partial \mathbf{z}(0)} \Psi_{t,H}(\mathbf{z}(0)) \delta \mathbf{z}(0) \right] &= \frac{\partial}{\partial \mathbf{z}(0)} \left(\frac{d}{dt} \Psi_{t,H}(\mathbf{z}(0)) \right) \cdot \delta \mathbf{z}(0) \\ &= \frac{\partial}{\partial \mathbf{z}(0)} (\mathbf{J} \nabla H(\mathbf{z}(t))) \cdot \delta \mathbf{z}(0) = \frac{\partial}{\partial \mathbf{z}(0)} (\mathbf{J} \nabla H(\Psi_{t,H}(\mathbf{z}(0)))) \cdot \delta \mathbf{z}(0) \\ &= \mathbf{J} \nabla^2 H(\mathbf{z}(t)) \cdot \left[\frac{\partial}{\partial \mathbf{z}(0)} \Psi_{t,H}(\mathbf{z}(0)) \delta \mathbf{z}(0) \right]. \end{aligned}$$

The above equation describing the evolution of $\frac{\partial}{\partial \mathbf{z}(0)} \Psi_{t,H}(\mathbf{z}(0)) \cdot \delta \mathbf{z}(0)$, the infinitesimal distance between two paths. It is called the *variational equation*. Let $\mathbf{F}(t) := \frac{\partial}{\partial \mathbf{z}(0)} \Psi_{t,H}(\mathbf{z}(0))$ be the Jacobian matrix of the Hamiltonian flow at time t (for simplicity, the dependence on $\mathbf{z}(0)$ is suppressed in the notation). Using the fact that the variational equations hold for every direction $\delta \mathbf{z}(0)$, it follows that this satisfies the matrix valued variational equation

$$\frac{d}{dt} \mathbf{F} = \mathbf{J} \nabla^2 H(\mathbf{z}(t)) \mathbf{F}, \quad (13)$$

with initial condition $\mathbf{F}(0) = \mathbf{I}_{2d}$.

In order to establish that Hamiltonian flow is symplectic, it suffices to verify that (12) holds. In this case this is equivalent to

$$\mathbf{F}_t^T \mathbf{J}^{-1} \mathbf{F}_t = \mathbf{J}^{-1}. \quad (14)$$

Since $\mathbf{F}_0 = \mathbf{I}_{2d}$, (14) holds for $t = 0$. By taking the derivative, and using (13) we obtain that

$$\begin{aligned} \frac{d}{dt} (\mathbf{F}_t^T \mathbf{J}^{-1} \mathbf{F}_t) &= (\mathbf{J} \nabla^2 H(\mathbf{z}(t)) \mathbf{F}_t)^T \mathbf{J}^{-1} \mathbf{F}_t + \mathbf{F}_t^T \mathbf{J}^{-1} \mathbf{J} \nabla^2 H(\mathbf{z}(t)) \mathbf{F}_t \\ &= \mathbf{F}_t^T \nabla^2 H(\mathbf{z}(t)) \mathbf{J}^T \mathbf{J}^{-1} \mathbf{F}_t + \mathbf{F}_t^T \mathbf{J}^{-1} \mathbf{J} \nabla^2 H(\mathbf{z}(t)) \mathbf{F}_t = \mathbf{0}, \end{aligned}$$

since $\mathbf{J}^T \mathbf{J}^{-1} = -\mathbf{J} \mathbf{J}^{-1} = -\mathbf{I}_{2d}$ and $\mathbf{J}^{-1} \mathbf{J} = \mathbf{I}_{2d}$. Hence (14) holds for every $t \geq 0$ and the Hamiltonian flow is symplectic. Since every symplectic flow is volume preserving, the Hamiltonian flow is also volume preserving. ■

Finally, the next proposition states the stationarity of the target with respect to Hamiltonian flow.

Proposition 3 *The distribution $\pi(d\mathbf{z}) \propto \exp(-H(\mathbf{z}))d\mathbf{z}$ is stationary with respect to the Hamiltonian flow.*

Proof. Let the density of π be equal $\pi(\mathbf{z}) = \frac{\exp(-H(\mathbf{z}))}{C}$ for some normalising constant C . To see that this result holds, let $\mathbf{Z}_0 \sim \pi$, and $\mathbf{Z}_t = \Psi_{t,H}(\mathbf{Z}_0)$. Let $\mathbf{z} = (z_1, \dots, z_d)^T$, then for infinitesimally small δ ,

$$\begin{aligned} \mathbb{P}(\mathbf{Z}_t \in [z_1, z_1 + \delta] \times \dots \times [z_d, z_d + \delta]) &= \mathbb{P}(\mathbf{Z}_0 \in \Psi_{-t,H}([z_1, z_1 + \delta] \times \dots \times [z_d, z_d + \delta])) \\ &= \pi(\Psi_{-t,H}(\mathbf{z})) \text{Vol}(\Psi_{-t,H}([z_1, z_1 + \delta] \times \dots \times [z_d, z_d + \delta])) \\ &= \frac{\exp(-H(\Psi_{-t,H}(\mathbf{z})))}{C} \text{Vol}(\Psi_{-t,H}([z_1, z_1 + \delta] \times \dots \times [z_d, z_d + \delta])) \end{aligned}$$

using the preservation of the Hamiltonian, and the volume,

$$= \frac{\exp(-H(\mathbf{z}))}{C} \text{Vol}([z_1, z_1 + \delta] \times \dots \times [z_d, z_d + \delta]) = \pi([z_1, z_1 + \delta] \times \dots \times [z_d, z_d + \delta]),$$

hence the stationarity is established. ■

Hamilton's equations themselves do not define an ergodic Markov chain, as they keep the Hamiltonian invariant. Let \mathbf{P}_R be the Markov kernel that resample the momentum according to the Gaussian distribution with covariance matrix \mathbf{M} , and $\Psi_{T,H}$ be the Hamiltonian flow-map for time T . Then π is stationary with respect to the combination $\mathbf{P}_R \cdot \Psi_{T,H}$ (i.e. first resample momentum, then move according to the Hamiltonian dynamics). Moreover, it is easy to show that if we choose T to be a random variable chosen independently at each time step, then the resulting Markov kernel is a mixture of kernels whose stationary distribution is π , and therefore it also admits π as a stationary distribution. We call this Markov kernel *Randomized Hamiltonian Monte Carlo* with continuous dynamics. The following proposition states some conditions that guarantee π -irreducibility.

Proposition 4 Let T be distributed according to ν_T , which has positive density on an interval $[0, \tau]$ for some $\tau > 0$. Let \mathbf{K} denote the Markov kernel for the position variables on \mathbb{R}^d corresponding to sampling a random momentum \mathbf{p} , then running the Hamiltonian dynamics started at (\mathbf{q}, \mathbf{p}) up to time T sampled from ν_T (independently at each step), and finally discarding the momentum variable. Suppose that U is continuously differentiable on \mathbb{R}^d , and satisfies that $\sup_{\mathbf{q}} \|\nabla^2 U(\mathbf{q})\| \leq L$, and $\inf_{\mathbf{q} \in \mathbb{R}^d} U(\mathbf{q}) > -\infty$. Then \mathbf{K} is strongly ν -irreducible.

In order to show this result, we are going to first show two preliminary lemmas.

Lemma 1 In addition to the conditions of Proposition 4, suppose that $\mathbf{M} = \mathbf{I}_d$, and that $0 \leq t \leq \frac{1}{2(L+1)^3}$. For any starting point $\mathbf{q}(0), \mathbf{p}(0)$, let $(\mathbf{q}(s), \mathbf{p}(s))_{s \geq 0}$ be the solution of the Hamiltonian dynamics (1)-(2). Then the Jacobian of $\mathbf{q}(t)$ in terms of the initial momentum $\mathbf{p}(0)$ satisfies that

$$\frac{1}{2}t\mathbf{I}_d \preceq \frac{\partial \mathbf{q}(t)}{\partial \mathbf{p}(0)} \preceq \frac{3}{2}t\mathbf{I}_d,$$

where \preceq denotes the positive semidefinite order (i.e. $\mathbf{A} \preceq \mathbf{B}$ means that $\mathbf{B} - \mathbf{A}$ is positive semidefinite).

Proof. Let $\mathbf{F}(t) := \frac{\partial}{\partial \mathbf{z}(0)} \Psi_{t,H}(\mathbf{z}(0))$ denote the Jacobian matrix at time t , then from the evolution equation (13) it follows that

$$\begin{aligned} \frac{d}{dt}(\mathbf{F} - \mathbf{I}_{2d}) &= \mathbf{J}\nabla^2 H(\mathbf{z}(t))\mathbf{F} = \mathbf{J}\nabla^2 H(\mathbf{z}(t))(\mathbf{F} - \mathbf{I}_{2d}) + \mathbf{J}\nabla^2 H(\mathbf{z}(t)) \\ &= \begin{pmatrix} \mathbf{0} & \mathbf{I}_d \\ -\nabla^2 U(\mathbf{q}_t) & \mathbf{0} \end{pmatrix} ((\mathbf{F}(t) - \mathbf{I}_{2d}) + \mathbf{I}_{2d}). \end{aligned} \quad (15)$$

Using the assumption of the lemma we know that $\|\nabla^2 U(\mathbf{q}_t)\| \leq L$, hence $\left\| \begin{pmatrix} \mathbf{0} & \mathbf{I}_d \\ -\nabla^2 U(\mathbf{q}_t) & \mathbf{0} \end{pmatrix} \right\| \leq (L+1)$, and therefore by (15) we have $\frac{d}{dt}(\|\mathbf{F} - \mathbf{I}_{2d}\| + (L+1)) \leq (L+1)(\|\mathbf{F} - \mathbf{I}_{2d}\| + (L+1))$. Using Gronwall's lemma, and the fact that $\mathbf{F}(0) = \mathbf{I}_d$, we have $\|\mathbf{F}(t) - \mathbf{I}_{2d}\| + (L+1) \leq (L+1)\exp((L+1)t)$. Hence for $t \leq \frac{1}{L+1}$, we have

$$\|\mathbf{F}(t) - \mathbf{I}_{2d}\| \leq (L+1)(\exp((L+1)t) - 1) \leq 2(L+1)^2 t.$$

We can rewrite (15) as

$$\frac{d}{dt} \left(\mathbf{F} - \mathbf{I}_{2d} - \begin{pmatrix} \mathbf{0} & t\mathbf{I}_d \\ -\int_{s=0}^t \nabla^2 U(\mathbf{q}_s) ds & \mathbf{0} \end{pmatrix} \right) = \begin{pmatrix} \mathbf{0} & \mathbf{I}_d \\ -\nabla^2 U(\mathbf{q}_t) & \mathbf{0} \end{pmatrix} (\mathbf{F}(t) - \mathbf{I}_{2d}),$$

where $\left\| \begin{pmatrix} \mathbf{0} & \mathbf{I}_d \\ -\nabla^2 U(\mathbf{q}_t) & \mathbf{0} \end{pmatrix} (\mathbf{F}(t) - \mathbf{I}_{2d}) \right\| \leq 2(L+1)^3 t$ for $t \leq \frac{1}{L+1}$. Since $\frac{\partial \mathbf{q}(t)}{\partial \mathbf{p}(0)}$ is the upper-right block matrix of $\mathbf{F}(t)$, the above equation implies that $\left\| \frac{\partial \mathbf{q}(t)}{\partial \mathbf{p}(0)} - t\mathbf{I}_d \right\| \leq \int_{s=0}^t s \cdot 2(L+1)^3 ds = (L+1)^3 t^2$, and the result follows. ■

Lemma 2 Under the conditions of Proposition 4, and assuming that $\mathbf{M} = \mathbf{I}_d$, for any $\mathbf{q}(0), \mathbf{Q} \in \mathbb{R}^d$, any $0 < t \leq \tau$, if $\mathbf{p}(0) = \frac{\mathbf{Q} - \mathbf{q}(0)}{t}$, and $(\mathbf{q}(s), \mathbf{p}(s))_{s \geq 0}$ is the solution of the Hamiltonian dynamics (1)-(2), then

$$\|\mathbf{q}(t) - \mathbf{Q}\| \leq Ct^2$$

for a finite constant C depending only on $\mathbf{q}, \mathbf{Q}, \tau$ and U but independent of t .

Proof. Let $H(\mathbf{q}(0), \mathbf{p}(0)) = U(\mathbf{q}(0)) + \frac{\|\mathbf{p}(0)\|^2}{2}$ denote the initial Hamiltonian, and $H(\mathbf{q}(s), \mathbf{p}(s)) = U(\mathbf{q}(s)) + \frac{\|\mathbf{p}(s)\|^2}{2}$ be the current Hamiltonian. Since this is preserved by the Hamiltonian dynamics, for any $s \geq 0$, we have

$$\begin{aligned} \frac{\|\mathbf{p}(s)\|^2}{2} &= U(\mathbf{q}) + \frac{\|\mathbf{p}(0)\|^2}{2} - U(\mathbf{q}(s)) \leq \frac{\|\mathbf{p}(0)\|^2}{2} + U(\mathbf{q}) - \inf_{\mathbf{q}' \in \mathbb{R}^d} U(\mathbf{q}'), \\ \|\mathbf{p}(s)\| &\leq \sqrt{2(U(\mathbf{q}) - \inf_{\mathbf{q}' \in \mathbb{R}^d} U(\mathbf{q}')) + \|\mathbf{p}(0)\|^2} \leq \sqrt{2(U(\mathbf{q}(0)) - \inf_{\mathbf{q}' \in \mathbb{R}^d} U(\mathbf{q}'))} + \|\mathbf{p}(0)\|. \end{aligned}$$

Since $\frac{d\mathbf{q}}{ds} = \mathbf{p}(s)$, we have

$$\|\mathbf{q}(s) - \mathbf{q}(0)\| \leq \int_{r=0}^s \|\mathbf{p}(r)\| dr \leq s \cdot \left(\sqrt{2(U(\mathbf{q}(0)) - \inf_{\mathbf{q}' \in \mathbb{R}^d} U(\mathbf{q}'))} + \|\mathbf{p}(0)\| \right).$$

Hence for any $0 < t \leq \tau$, for $\mathbf{p}(0) = \frac{\mathbf{Q} - \mathbf{q}(0)}{t}$, for any $0 \leq s \leq t$, we have

$$\|\mathbf{q}(s) - \mathbf{q}(0)\| \leq D \text{ for } D := \tau \sqrt{2(U(\mathbf{q}) - \inf_{\mathbf{q}' \in \mathbb{R}^d} U(\mathbf{q}'))} + \|\mathbf{Q} - \mathbf{q}(0)\|. \quad (16)$$

Let $E := \sup_{\mathbf{q}: \|\mathbf{q} - \mathbf{q}(0)\| \leq D} \|\nabla U(\mathbf{q})\|$, then using the assumption that $\sup_{\mathbf{q} \in \mathbb{R}^d} \|\nabla^2 U(\mathbf{q})\| \leq L$, it follows that E is finite. Since by the Hamiltonian equations $\frac{d\mathbf{p}}{ds} = -\nabla U(\mathbf{q}(s))$, it follows that for $0 \leq s \leq t \leq \tau$, we have $\|\frac{d\mathbf{p}}{ds}\| \leq E$, and therefore

$$\|\mathbf{p}(s) - \mathbf{p}(0)\| = \left\| \int_{r=0}^s \frac{d\mathbf{p}}{ds} ds \right\| \leq \int_{r=0}^s \left\| \frac{d\mathbf{p}}{ds} \right\| ds \leq s \cdot E.$$

Now using the Hamiltonian equations again in q , we have that for $0 \leq t \leq \tau$,

$$\begin{aligned} \mathbf{q}(t) &= \mathbf{q}(0) + \int_{s=0}^t \mathbf{p}(s) ds = \mathbf{q}(0) + \int_{s=0}^t [p(0) + (p(s) - p(0))] ds \\ &= \mathbf{q}(0) + \mathbf{p}(0)t + \int_{s=0}^t (\mathbf{p}(s) - \mathbf{p}(0)) ds \end{aligned}$$

using the fact that $\mathbf{p}(0) = \frac{\mathbf{Q} - \mathbf{q}(0)}{t}$

$$= \mathbf{Q} + \int_{s=0}^t (\mathbf{p}(s) - \mathbf{p}(0)) ds.$$

Therefore we have

$$\|\mathbf{q}(t) - \mathbf{Q}\| = \left\| \int_{s=0}^t (\mathbf{p}(s) - \mathbf{p}(0)) ds \right\| \leq \int_{s=0}^t \|\mathbf{p}(s) - \mathbf{p}(0)\| ds \leq \int_{s=0}^t s \cdot E ds = \frac{E}{2} t^2,$$

hence the claim of the lemma holds for $C = E/2$. ■

Proof of Proposition 4. Let $\nu(\mathbf{q}) \propto \exp(-U(\mathbf{q}))$ be the target distribution, and μ be the multivariate normal distribution on \mathbb{R}^d with mean 0 and covariance matrix \mathbf{M} . Assume without loss of generality that $\mathbf{M} = \mathbf{I}_d$ (the general case follows by a simple modification of this argument).

Our goal is to show that for any starting point $\mathbf{q}(0)$, any measurable set $A \subset \mathbb{R}^d$ with $\nu(A) > 0$, we have $K(\mathbf{q}(0), A) > 0$. Since we can fill \mathbb{R}^d into countably many balls of radius δ (for any $\delta > 0$), we can assume without loss of generality that A is a subset of a ball of radius δ centered at some point $\mathbf{Q} \in \mathbb{R}^d$. We will denote this closed ball by $B_\delta(\mathbf{Q}) = \{\mathbf{q} \in \mathbb{R}^d : \|\mathbf{q} - \mathbf{Q}\| \leq \delta\}$.

Then by Lemma 2, it follows that if $t \leq \frac{\sqrt{\delta}}{\sqrt{C}}$, and

$$\mathbf{p}(0) = \mathbf{p}^* = \frac{\mathbf{Q} - \mathbf{q}(0)}{t},$$

then $\mathbf{q}(t) \in B_\delta(\mathbf{Q})$.

Let $\Psi(\mathbf{p}(0))$ denote the position (\mathbf{q}) component of $\Psi_t(\mathbf{q}(0), \mathbf{p}(0))$ ($\mathbf{q}(0)$ is assumed fixed). Ψ maps the initial momentum $\mathbf{p}(0)$ to the position at time t .

It follows from Lemma 1 that for any vector \mathbf{v} , we have

$$\|\Psi(\mathbf{p}^*) - \Psi(\mathbf{p}^* + \mathbf{v})\| = \left\| \left(\int_{s=0}^1 \frac{\partial \mathbf{q}(t)}{\partial \mathbf{p}(0)} \Big|_{\mathbf{p}(0) = \mathbf{p}^* + \mathbf{v}s} ds \right) \mathbf{v} \right\| \geq \frac{t \|\mathbf{v}\|}{2}.$$

Therefore, with the choice $r = \frac{6\delta}{t}$ the sphere $S_r(\mathbf{p}^*) = \{\mathbf{p}' \in \mathbb{R}^d : \|\mathbf{p}' - \mathbf{p}^*\| = r\}$ satisfies that for any $\mathbf{p}' \in S_r(\mathbf{p}^*)$,

$$\|\Psi(\mathbf{p}^*) - \Psi(\mathbf{p}')\| \geq 3\delta.$$




Figure 1: Effect of the map Ψ

Since $\Psi(\mathbf{p}^*)$ is contained in $B_\delta(\mathbf{Q})$, it follows that the map of the ball $B_r(\mathbf{p}(0))$ by Ψ will contain the ball $B_\delta(\mathbf{Q})$ (see Figure 1).

Hence for any point in $\mathbf{Q}' \in B_\delta(\mathbf{Q})$, there is at least one point $\mathbf{p}' \in B_r(\mathbf{p}(0))$ such that $\Psi(\mathbf{p}') = \mathbf{q}'$. Denote the set of points in $B_r(\mathbf{p}(0))$ that get mapped into the set A by Ψ as B_t (preimage of A). Since $\nu(A) > 0$, and ν has a density, it follows that $\text{Vol}(A) > 0$, and by Lemma 1, it follows that if $0 < t \leq \frac{1}{2(L+1)^3}$, the determinant of the Jacobian of Ψ on B_t is finite and positive, so we also must have $\text{Vol}(B_t) > 0$, implying that $\mu(B_t) > 0$ (μ denotes the standard normal distribution on \mathbb{R}^d , corresponding to the distribution of the resampled momentum variable). Since this holds for any $t \leq \min\left(\frac{\sqrt{\delta}}{\sqrt{C}}, \frac{1}{2(L+1)^3}\right)$, and $T \sim \nu_T$ has positive density on $[0, \tau]$, it follows by integration that $K(\mathbf{q}(0), A) > 0$, thus K is strongly ν -irreducible. ■

In general the Hamiltonian flow cannot be simulated exactly due to the nonlinearity of the ODE. Hence we are going to approximate it by discretization.

3 Discretizing Hamilton's equations

We consider 3 simple discretization schemes.

Explicit scheme:

$$\begin{aligned}\mathbf{p}(t + \epsilon) &= \mathbf{p}(t) - \epsilon \nabla U(\mathbf{q}(t)) \\ \mathbf{q}(t + \epsilon) &= \mathbf{q}(t) + \epsilon \mathbf{M}^{-1} \mathbf{p}(t).\end{aligned}$$

Modified explicit scheme:

$$\begin{aligned}\mathbf{p}(t + \epsilon) &= \mathbf{p}(t) - \epsilon \nabla U(\mathbf{q}(t)) \\ \mathbf{q}(t + \epsilon) &= \mathbf{q}(t) + \epsilon \mathbf{M}^{-1} \mathbf{p}(t + \epsilon).\end{aligned}$$

Leapfrog (Störmer-Verlet) scheme

$$\begin{aligned}\mathbf{p}(t + \epsilon/2) &= \mathbf{p}(t) - \frac{\epsilon}{2} \nabla U(\mathbf{q}(t)) \\ \mathbf{q}(t + \epsilon) &= \mathbf{q}(t) + \epsilon \mathbf{M}^{-1} \mathbf{p}(t + \epsilon/2). \\ \mathbf{p}(t + \epsilon) &= \mathbf{p}(t + \epsilon/2) - \frac{\epsilon}{2} \nabla U(\mathbf{q}(t + \epsilon))\end{aligned}$$

The following figure shows these schemes implemented for the 1 dimensional standard Gaussian distribution corresponding to $H(p, q) = \frac{p^2}{2} + \frac{q^2}{2}$.

As Figure 2 shows, these 3 schemes behave very differently. The first explicit scheme diverges for step size $\epsilon = 0.3$. The modified explicit scheme does not diverge, but it also does not track the true trajectory closely. The Leapfrog scheme tracks the true trajectory very closely, and it does not diverge.

An important property of these discretizations is that they are symplectic, and hence volume preserving.

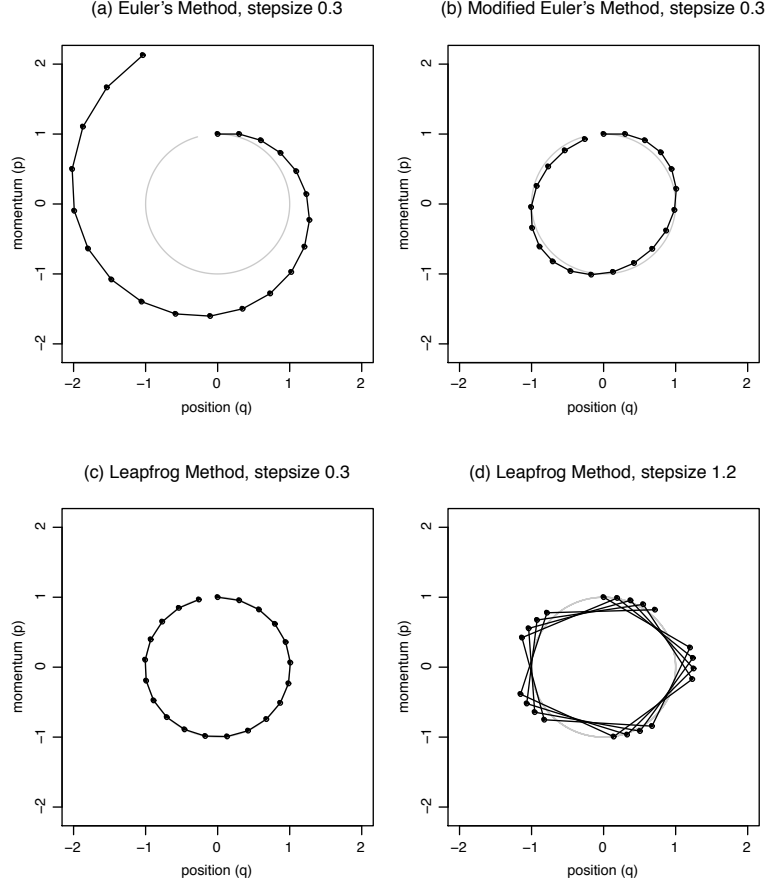


Figure 2: Approximation of Hamiltonian dynamics for $H(p, q) = \frac{p^2}{2} + \frac{q^2}{2}$ by 3 schemes. 20 steps in each case in black, the true trajectory in grey.

Proposition 5 *The above 3 schemes are symplectic, and hence volume preserving.*

Proof. First, note that by the chain rule for Jacobians, if $\Psi(\mathbf{z}) = \Psi_1(\Psi_2(\mathbf{z}))$, then $\nabla_{\mathbf{z}}\Psi(\mathbf{z}) = \nabla\Psi_1(\Psi_2(\mathbf{z})) \cdot \nabla\Psi_2(\mathbf{z})$. If Ψ_1 and Ψ_2 are symplectic, then

$$[\nabla\Psi(\mathbf{z})]^T \mathbf{J}^{-1} \nabla_{\mathbf{z}}\Psi(\mathbf{z}) = (\nabla\Psi_2(\mathbf{z}))^T (\nabla\Psi_1(\Psi_2(\mathbf{z})))^T \mathbf{J}^{-1} \nabla\Psi_1(\Psi_2(\mathbf{z})) \cdot \nabla\Psi_2(\mathbf{z}) = \mathbf{J}^{-1},$$

so Ψ is also symplectic. By induction, this also holds for the composition of more than two maps. Using this fact, it suffices to check symplecticity for a single step in the above schemes, modifying only one of p or q . For example, the step $\mathbf{p}(t + \epsilon) = \mathbf{p}(t) - \epsilon\nabla U(\mathbf{q}(t))$ corresponds to the map $\Psi(\mathbf{z}) = \begin{pmatrix} \mathbf{p} - \epsilon\nabla U(\mathbf{q}) \\ \mathbf{q} \end{pmatrix}$, which has Jacobian $\nabla\Psi(\mathbf{z}) = \begin{pmatrix} \mathbf{I}_d & \mathbf{0} \\ -\epsilon\nabla^2 U(\mathbf{q}) & \mathbf{I}_d \end{pmatrix}$. Now by direct calculation, one can see that

$$(\nabla\Psi(\mathbf{z}))^T \mathbf{J}^{-1} \nabla\Psi(\mathbf{z}) = \begin{pmatrix} \mathbf{I}_d & -\epsilon\nabla^2 U(\mathbf{q}) \\ \mathbf{0} & \mathbf{I}_d \end{pmatrix} \begin{pmatrix} \mathbf{0} & -\mathbf{I}_d \\ \mathbf{I}_d & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{I}_d & \mathbf{0} \\ -\epsilon\nabla^2 U(\mathbf{q}) & \mathbf{I}_d \end{pmatrix} = \begin{pmatrix} \mathbf{0} & -\mathbf{I}_d \\ \mathbf{I}_d & \mathbf{0} \end{pmatrix} = \mathbf{J}^{-1},$$

hence the symplecticity holds. This can be also checked for the other steps, and hence all 3 schemes are symplectic. The volume preservation follows by Proposition 1. ■

4 Hamiltonian Monte Carlo

In this section, we are going to introduce Hamiltonian Monte Carlo, originally proposed in [1]. This is an MCMC method that samples from a target distribution with density $\mu(\mathbf{q}) \propto \exp(-U(\mathbf{q}))$ on \mathbb{R}^d by extending the state space and sampling from the distribution π on \mathbb{R}^{2d} with density $\pi(\mathbf{z}) \propto \exp(-H(\mathbf{z})) =$

$\exp(-U(\mathbf{q}) - K(\mathbf{p})) = \exp(-U(\mathbf{q})) \exp\left(-\frac{\mathbf{p}^T \mathbf{M}^{-1} \mathbf{p}}{2}\right)$. From this decomposition, it is clear that the position and momentum variables are independent.

Let $\Psi_\epsilon : \mathbb{R}^{2d} \rightarrow \mathbb{R}^{2d}$ denote the Leapfrog map, and for some $L \in \mathbb{N}$, we denote by Ψ_ϵ^L the L times composition of the Leapfrog map. Let $\mathbf{N} : \mathbb{R}^{2d} \rightarrow \mathbb{R}^{2d}$ denote the map that negates the momentum, i.e. $\mathbf{N} \begin{pmatrix} \mathbf{q} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} \mathbf{q} \\ -\mathbf{p} \end{pmatrix}$, and finally \mathbf{P}_R be the Markov kernel that resamples the momentum component \mathbf{p} from a multivariate Gaussian distribution with covariance matrix \mathbf{M} . Hamiltonian Monte Carlo consists of the iteration of 2 steps,

1. resample the momentum component \mathbf{p} , i.e. apply the Markov kernel \mathbf{P}_R .
2. We propose a new position $\begin{pmatrix} \mathbf{q}^* \\ \mathbf{p}^* \end{pmatrix} = \Psi \begin{pmatrix} \mathbf{q} \\ \mathbf{p} \end{pmatrix} := \mathbf{N} \left(\Psi_\epsilon^L \begin{pmatrix} \mathbf{q} \\ \mathbf{p} \end{pmatrix} \right)$ by applying L Leapfrog steps and then flipping the momentum. This new position is accepted with probability

$$\min[1, \exp(H(\mathbf{q}, \mathbf{p}) - H(\mathbf{q}^*, \mathbf{p}^*))] = \min \left[1, \exp \left(U(\mathbf{q}) - U(\mathbf{q}^*) + \frac{1}{2} \mathbf{p}^T \mathbf{M}^{-1} \mathbf{p} - \frac{1}{2} (\mathbf{p}^*)^T \mathbf{M}^{-1} \mathbf{p}^* \right) \right]. \quad (17)$$

The following proposition shows the invariance of the target distribution for HMC.

Proposition 6 *π is invariant with respect to the Markov kernel proposed above. Moreover, π is reversible with respect to the second step of the Markov kernel.*

Proof. Since \mathbf{q} and \mathbf{p} are independent according to the target distribution π , it is clear that the first step keeps the target invariant. Let \mathbf{P}_2 denote the Markov kernel corresponding to the second step (a combination of a deterministic step and Metropolis-Hastings accept-reject step), then we are going to check that π is reversible with respect to \mathbf{P}_2 .

We have seen during the previous lectures that in general a Markov kernel K is reversible with respect to a distribution π on state space \mathbb{X} if for every bounded measurable function $f : \mathbb{X}^2 \rightarrow \mathbb{R}$, we have

$$\int \int f(\mathbf{x}, \mathbf{y}) \pi(d\mathbf{x}) K(\mathbf{x}, d\mathbf{y}) = \int \int f(\mathbf{x}, \mathbf{y}) \pi(d\mathbf{y}) K(\mathbf{y}, d\mathbf{x}). \quad (18)$$

In our case, $\mathbf{P}_2(\mathbf{x}, d\mathbf{y})$ is non-zero only for $\mathbf{y} = \Psi(\mathbf{x})$ and $\mathbf{y} = \mathbf{x}$, so we have

$$\begin{aligned} \int \int f(\mathbf{x}, \mathbf{y}) \pi(d\mathbf{x}) \mathbf{P}_2(\mathbf{x}, d\mathbf{y}) &= \int \int f(\mathbf{x}, \Psi(\mathbf{x})) \min[1, e^{H(\mathbf{x}) - H(\Psi(\mathbf{x}))}] \pi(d\mathbf{x}) \\ &\quad + \int \int f(\mathbf{x}, \mathbf{x}) \left(1 - \min[1, e^{H(\mathbf{x}) - H(\Psi(\mathbf{x}))}] \right) \pi(d\mathbf{x}) \end{aligned}$$

Let $\mathbf{y} = \Psi(\mathbf{x})$, then $\mathbf{x} = \Psi(\mathbf{y}) = \Psi(\Psi(\mathbf{x}))$, and by the volume preserving property of Ψ , we have

$$\pi(d\mathbf{y}) = \pi(d\mathbf{x}) \cdot \frac{\exp(-H(\mathbf{y}))}{\exp(-H(\mathbf{x}))} = \pi(d\mathbf{x}) \cdot e^{H(\mathbf{x}) - H(\mathbf{y})},$$

and the first part of the above sum can be written

$$\begin{aligned} &\int \int f(\mathbf{x}, \Psi(\mathbf{x})) \min[1, e^{H(\mathbf{x}) - H(\Psi(\mathbf{x}))}] \pi(d\mathbf{x}) \\ &= \int \int f(\Psi(\mathbf{y}), \mathbf{y}) \min[1, e^{H(\Psi(\mathbf{y})) - H(\mathbf{y})}] \pi(d\mathbf{x}) \\ &= \int \int f(\Psi(\mathbf{y}), \mathbf{y}) \min[1, e^{H(\Psi(\mathbf{y})) - H(\mathbf{y})}] \cdot e^{H(\mathbf{y}) - H(\Psi(\mathbf{y}))} \pi(d\mathbf{y}) \\ &= \int \int f(\Psi(\mathbf{y}), \mathbf{y}) \min[1, e^{H(\mathbf{y}) - H(\Psi(\mathbf{y}))}] \pi(d\mathbf{y}). \end{aligned}$$

The second part satisfies that

$$\begin{aligned} &\int \int f(\mathbf{x}, \mathbf{x}) \left(1 - \min[1, e^{H(\mathbf{x}) - H(\Psi(\mathbf{x}))}] \right) \pi(d\mathbf{x}) \\ &= \int \int f(\mathbf{y}, \mathbf{y}) \left(1 - \min[1, e^{H(\mathbf{y}) - H(\Psi(\mathbf{y}))}] \right) \pi(d\mathbf{y}). \end{aligned}$$

By combining these two equations, we can see that the reversibility condition (18) holds. ■

In practice, one does not need to flip the momentum at the end of the second step, since it will be resampled in the first step in the next iteration. This was nevertheless required for showing the reversibility.

The following two examples illustrate the behaviour of Hamiltonian Monte Carlo.

Example 2 *In this example, we did some simulations for the simple 2D Gaussian distribution with covariance matrix $\Sigma = \begin{pmatrix} 1 & 0.98 \\ 0.98 & 1 \end{pmatrix}$. We have used mass matrix $M = I_2$, and 25 leapfrog steps per iteration using stepsize 0.25. Figures 3 and 4 illustrate the behaviour of Hamiltonian Monte Carlo in this example, and compare it with random walk Metropolis.*




Figure 3: The position, momentum coordinates and the value of the Hamiltonian for 25 Leapfrog steps. As we can see, the Hamiltonian is kept approximately constant by the Leapfrog steps, hence the acceptance rate is close to 1.

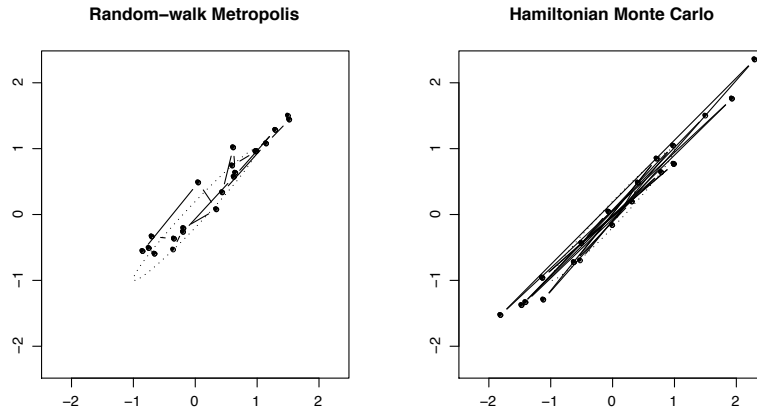


Figure 4: Twenty iterations of the random walk Metropolis method (20 updates per iteration) and the Hamiltonian Monte Carlo method (20 Leapfrog steps) for a highly correlated 2 dimensional Gaussian distribution. As we can see, Hamiltonian Monte Carlo is making much larger moves and mixes faster than random walk Metropolis.

Example 3 *In this example, we did some simulations for a 100 dimensional multivariate Gaussian distribution where the components are independent with standard deviations 0.01, 0.02, ..., 1.00. The parameter ϵ was randomly chosen at each iteration uniformly from (0.0104, 0.0156), and we selected $L = 150$. We compared HMC with the random walk Metropolis chain, counting 150 Metropolis updates per one iteration to be fair. The Figure 5 shows the last component (which had the largest standard deviation) for the two chains.*

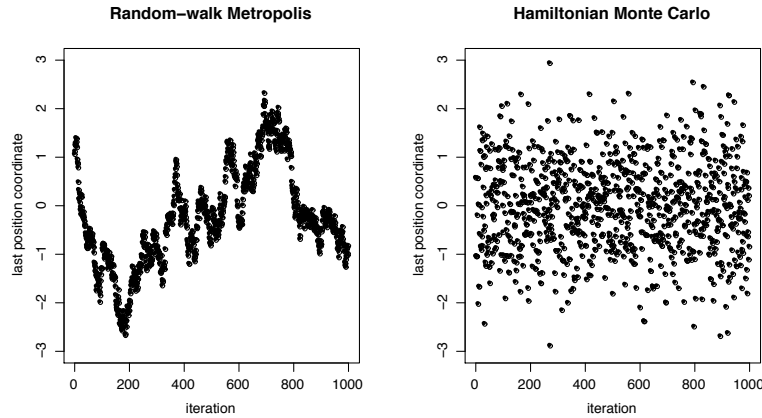


Figure 5: The last component in a 100 dimensional Gaussian distribution, random walk Metropolis versus Hamiltonian Monte Carlo. Hamiltonian Monte Carlo is mixing much faster.

5 Concentration of measure and high dimensional behaviour

As we have seen in the previous examples, HMC had performed much better than random walk Metropolis in high dimensions. To understand the reason for this, it is important to study the concentration of measure phenomenon that appears in high dimensions.

Suppose that $\mathbf{Z} \sim N(0, \mathbf{I}_d)$ is a d dimensional standard normal random vector. Then the Euclidean norm of \mathbf{Z} satisfies that for every $t \geq 0$,

$$P(|\|\mathbf{Z}\| - \sqrt{d}| \geq t) \leq C \exp\left(-\frac{t^2}{C}\right), \quad (19)$$

where C is an absolute constant independent of d . This means that with high probability, $\|\mathbf{Z}\| = \sqrt{d} + O(1)$, i.e. most of the probability is concentrated in a thin layer around the sphere of radius \sqrt{d} . Note that here $H(\mathbf{z}) = \frac{\|\mathbf{z}\|^2}{2}$, and since $H(\mathbf{z})$ is preserved by the Hamiltonian, Hamiltonian dynamics moves around in circular arcs. See Figure 6 for an illustration of this.

In general, if the Hessian of the target potential satisfies that $\mu \mathbf{I}_d \preceq \nabla^2 U(\mathbf{x}) \preceq L \mathbf{I}_d$ for some $0 < \mu < L < \infty$ (strongly convex and smooth potential), and we let $H_{\min} := \inf_{\mathbf{z}} H(\mathbf{z})$, then it is possible to show that

$$P\left(\left|\sqrt{H(\mathbf{z}) - H_{\min}} - \mathbb{E}\sqrt{H(\mathbf{z}) - H_{\min}}\right| \geq t\right) \leq C \exp\left(-\frac{t^2}{C}\right), \quad (20)$$

for some constant C that depends on μ and L but is independent of the dimension d . Hence the Hamiltonian is close to constant in the area of the space with high probability density, and therefore Hamiltonian dynamics can be very efficient in exploring this potentially complicated set automatically. In contrast to this, random walk Metropolis will need to take small moves of size $O(1)$ to keep acceptance rates reasonably high, and therefore it takes a long time to explore such high dimensional distributions.

References

- [1] Duane, Simon and Kennedy, Anthony D and Pendleton, Brian J and Roweth, Duncan, Hybrid Monte Carlo, *Physics letters B*, vol. 195, no. 2, pp. 216-222, 1987.
- [2] Neal, Radford M, MCMC using Hamiltonian dynamics, *Handbook of Markov Chain Monte Carlo*, vol. 2, no. 11, 2011
- [3] Leimkuhler, Benedict and Reich, Sebastian, *Simulating Hamiltonian dynamics*, Cambridge University Press, vol. 14, 2004
- [4] Betancourt, Michael, A conceptual introduction to Hamiltonian Monte Carlo, <https://arxiv.org/abs/1701.02434>




Figure 6: The density of the distribution π for a 100 dimensional standard Gaussian, when looking at the first components p_1, q_1 (other components have been integrated out). The dashed line shows a possible Hamiltonian path in this case.