Unbiased constrained sampling with Self-Concordant Barrier Hamiltonian Monte Carlo

Maxence Noble[1] Valentin de Bortoli[2] Alain Durmus[1]

[1]Centre de Mathématiques Appliquées, Ecole Polytechnique
Institut Polytechnique de Paris, France

[2]Département d'Informatique, École Normale Supérieure
CNRS, Université PSL, Paris, France

Constrained sampling & self-concordance

- $\mathrm{M} \subset \mathbb{R}^d$
- $\pi \in \mathscr{P}(\mathbb{R}^d)$, supported on M, known up to a normalising constant

Our goal: sample from π.

Assumption: M is convex, with a self-concordant barrier ϕ (Nesterov and Nemirovskii, 1994):

- ϕ is convex (*and regular*),
- $\phi(x) \to \infty$ as $x \to \partial\mathrm{M}$.

Example: polytope with its logarithmic barrier.

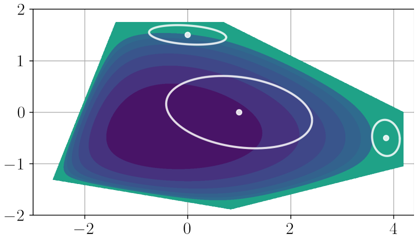

Traditional sampling with HMC

When $\mathrm{M} = \mathbb{R}^d$

- **Hamiltonian Monte Carlo** (HMC) (Duane et al., 1987)

Extended target probability distribution:

$$\begin{aligned} &\mathrm{d}\bar{\pi}(x,p) = \mathrm{d}\pi(x)\mathrm{d}\nu(p), \quad \nu = \mathrm{N}(0, \mathrm{I}_d) \\ \implies &\mathrm{d}\bar{\pi}(x,p) \propto \exp(-H(x,p))\mathrm{d}(x,p) \end{aligned}$$

- H: *Hamiltonian* function

Traditional sampling with HMC

When $\mathrm{M} = \mathbb{R}^d$

- **Hamiltonian Monte Carlo** (HMC) (Duane et al., 1987)

Extended target probability distribution:

$$\mathrm{d}\bar{\pi}(x, p) = \mathrm{d}\pi(x)\mathrm{d}\nu(p), \quad \nu = \mathrm{N}(0, \mathrm{I}_d)$$
$$\Longrightarrow \mathrm{d}\bar{\pi}(x, p) \propto \exp(-H(x, p))\mathrm{d}(x, p)$$

- H: *Hamiltonian* function
- *Hamiltonian* dynamics → de-coupled

$$\dot{x} = \partial_p H(x, p) , \qquad \dot{p} = -\partial_x H(x, p) .$$

$\Longrightarrow$ Explicit and **involutive** integrator F_h (h : step-size)

Markov chain step: given $z^0 = (x^0, p^0)$, compute the *proposal* state $z^1 = \mathrm{F}_h(z^0)$ and apply a *Metropolis-Hastings* filter.

🎯 Reversible scheme w.r.t. $\bar{\pi}$

Constrained sampling with RHMC

In our setting: M is a *Riemmanian submanifold* with metric $\mathfrak{g} = D^2\phi$

- **Riemannian HMC** (RHMC) (Girolami and Calderhead, 2011; Lee and Vempala, 2018; Kook et al., 2022)

Extended target probability distribution:

$$\begin{aligned} \mathrm{d}\bar{\pi}(x,p) &= \mathrm{d}\pi(x)\mathrm{d}\bar{\pi}(y|x), \quad \bar{\pi}(y|x) = \mathrm{N}(0,\mathfrak{g}(x)) \ , \\ \implies \mathrm{d}\bar{\pi}(x,p) &\propto \exp(-H(x,p))\mathrm{d}(x,p) \end{aligned}$$

- H: *Riemannian Hamiltonian* function

Constrained sampling with RHMC

In our setting: M is a *Riemmanian submanifold* with metric $\mathfrak{g} = D^2\phi$

- **Riemannian HMC** (RHMC) (Girolami and Calderhead, 2011; Lee and Vempala, 2018; Kook et al., 2022)

Extended target probability distribution:

$$\mathrm{d}\bar{\pi}(x,p) = \mathrm{d}\pi(x)\mathrm{d}\bar{\pi}(y|x), \quad \bar{\pi}(y|x) = \mathrm{N}(0, \mathfrak{g}(x)) \;,$$
$$\implies \mathrm{d}\bar{\pi}(x,p) \propto \exp(-H(x,p))\mathrm{d}(x,p)$$

- H: *Riemannian Hamiltonian* function
- *Hamiltonian* dynamics → coupled

$$\dot{x} = \partial_p H(x,p) \;, \qquad \dot{p} = -\partial_x H(x,p) \;.$$

$\implies$ Implicit integrator F_h → **numerical integrator** Φ_h

⛔ F_h: several (or none) solutions → Φ_h: no involution guarantee...

Markov chain step: given $z^0 = (x^0, p^0)$, compute the *proposal* state $z^1 = \Phi_h(z^0)$ and apply a *Metropolis-Hastings* filter.

⛔ No more reversibility w.r.t. $\bar{\pi}$

Our contribution: Barrier-HMC

Our idea: enforce an involution condition on Φ_h in RHMC.

💡 **"Involution Checking Step"**

1. Compute $z^1 = \Phi_h(z^0)$.
2. Check if $\|\Phi_h(z^1) - z^0\| \approx 0$.
3. If not, reject z^1.

In practice

- Easy to implement.
- $\|\cdot\|$ depends on z^0 and $\mathfrak{g}$.
- $\approx$ is controlled by some $\eta > 0$.

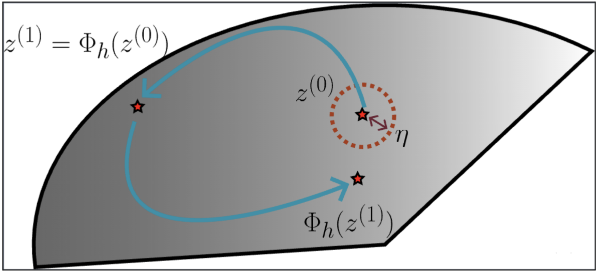

Used with *Metropolis-Hastings* filter, this new scheme is reversible w.r.t. $\bar{\pi}$!

Conclusion

- Constrained sampling for distributions supported on convex subsets endowed with a self-concordant barrier (including polytopes).
- Asymptotic bias fixed in practice by the "**involution checking step**" in our algorithm : **Barrier HMC**.
- Reversibility results & numerical experiments.

Unbiased constrained sampling with Self-Concordant Barrier Hamiltonian Monte Carlo

Maxence Noble, Valentin de Bortoli, Alain Durmus

Any questions ?

Send an e-mail to maxence.noble-bourillot@polytechnique.edu.

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

Mark Girolami and Ben Calderhead. Riemann manifold langevin and Hamiltonian Monte Carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(2):123–214, 2011.

Yunbum Kook, Yin-Tat Lee, Ruoqi Shen, and Santosh Vempala. Sampling with Riemannian Hamiltonian Monte Carlo in a constrained space. *Advances in Neural Information Processing Systems*, 35:31684–31696, 2022.

Yin Tat Lee and Santosh S Vempala. Convergence rate of Riemannian Hamiltonian Monte Carlo and faster polytope volume computation. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 1115–1121, 2018.

Yurii Nesterov and Arkadii Nemirovskii. *Interior-point polynomial algorithms in convex programming*. SIAM, 1994.