

Deep Learning Hamiltonian Monte Carlo^[1]

Building topological samplers for lattice gauge theories

Sam Foreman¹, Xiao-Yong Jin², James Osborn³
Argonne National Laboratory

bit.ly/l2hmc-qcd

Abstract

We generalize the Hamiltonian Monte Carlo (HMC) algorithm with a stack of trainable neural network (NN) layers and evaluate its ability to sample from different topologies in a two-dimensional lattice gauge theory. We demonstrate that our model is able to successfully mix between modes of different topologies, significantly reducing the computational cost required to generate independent gauge field configurations.

Motivation

- Markov Chain Monte Carlo (MCMC) methods are pervasive throughout science and are used in applications ranging from epidemiological modeling to election forecasting.
- Recent developments in ML, together with ever-more-capable hardware has led to a resurgence in developing faster, more efficient simulation techniques.
- In particular, the development of invertible NN architectures has opened the flood-gates for new approaches that are capable of outperforming traditional techniques on particularly challenging distributions.
- Simulations in lattice gauge theory / lattice QCD are limited by our ability to generate independent configurations, making it a prime target for testing novel approaches.
- We propose a generalized version of the L2HMC algorithm [2], and look at applying it to generate configurations for a two-dimensional $U(1)$ lattice gauge theory.

Method

We begin by introducing distinct neural networks, called *leapfrog layers* for each leapfrog step of the HMC update, as shown in Figure (a).

- Denote the leapfrog step (layer) by a discrete index $k \in \{0, 1, \dots, N_{LF}\} \in \mathbb{N}$ where N_{LF} is the total number of leapfrogs.
- Introduce $d \sim \mathcal{U}(\pm, -)$ (direction—forward/backward) and denote the complete state $\xi = (x, v, d)$, then the *target distribution* is given by $p(\xi) = p(x) \cdot p(v) \cdot p(d)$.
- Each leapfrog step transforms $\xi_k \equiv (x_k, v_k, d_k) \rightarrow (x'_k, v'_k, d'_k) = \xi'_k$ by passing it through k^{th} leapfrog layer (note the direction, d_k is persistent).
- Consider the forward $d = +1$ direction² and introduce the notation:

$$v'_k \equiv \Gamma_k^+(v_k, \zeta_{v_k}) = v_k \odot \exp\left(\frac{\zeta_{v_k}^k}{2} \zeta_{v_k}^k\right) - \frac{\zeta_{v_k}^k}{2} \left[\partial_x S_\beta(x_k) \odot \exp\left(\frac{\zeta_{v_k}^k}{2} \zeta_{v_k}^k\right) + t_{v_k}^k(\zeta_{v_k}^k) \right]$$

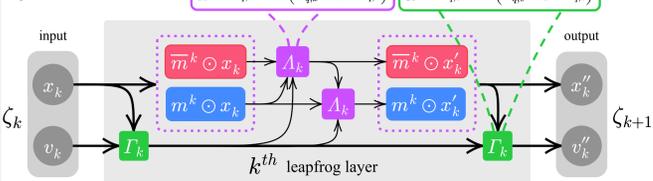
$$x'_k \equiv \Lambda_k^+(x_k, \zeta_{x_k}) = x_k \odot \exp\left(\frac{\zeta_{x_k}^k}{2} \zeta_{x_k}^k\right) + e_{x_k}^k \left[v'_k \odot \exp\left(\frac{\zeta_{x_k}^k}{2} \zeta_{x_k}^k\right) + t_{x_k}^k(\zeta_{x_k}^k) \right]$$

We can write a complete leapfrog update as: (compare with HMC (i.))

Generalized leapfrog update

- Half-step (v): $v'_k = \Gamma_k^+(v_k, \zeta_{v_k})$
- Full-step (1st half of x): $x'_k = \tilde{m}^k \odot x_k + m^k \odot \Lambda_k^\pm(x_k; \zeta_{x_k})$
- Full-step (2nd half of x): $x''_k = \tilde{m}^k \odot \Lambda_k^\pm(x'_k; \zeta_{x_k}) + m^k \odot x'_k$
- Half-step (v): $v''_k = \Gamma_k^\pm(v'_k, \zeta_{v_k})$

Figure (a.) Illustration of the k^{th} leapfrog layer.



- For reversibility we split the x update into two sub-updates using a binary mask, $m^k = 1 - \tilde{m}^k$ and update each half of x sequentially.
- We've introduced the shorthand notation for the networks' inputs: $\zeta_{v_k} \equiv (x_k, \partial_x S_\beta(x))$, $\zeta_{x_k} \equiv (m^k \odot x_k, v_k)$
- the *Jacobian factor* of the update $\xi \rightarrow \xi'$ can be easily computed to give: $\left| \frac{\partial v'_k}{\partial v_k} \right| = \exp\left(\frac{1}{2} \zeta_{v_k}^k \zeta_{v_k}^k\right)$, $\left| \frac{\partial x'_k}{\partial x_k} \right| = \exp\left(\zeta_{x_k}^k \zeta_{x_k}^k\right)$.

Method (contd.)

- For target distributions in Euclidean space, $x \in \mathbb{R}^n$, we define a loss function that encourages our sampler to move large distances in the phase space.
- To do this, we can maximize the *expected squared jump distance* (ESJD):

$$\mathcal{L}(\theta) \equiv \mathbb{E}_{p(\xi)} [\delta(\xi, \xi') \cdot A(\xi' | \xi)] \quad , \quad \delta(\xi, \xi') \equiv \|x - x'\|^2$$

where θ are (collectively) the weights in the NN

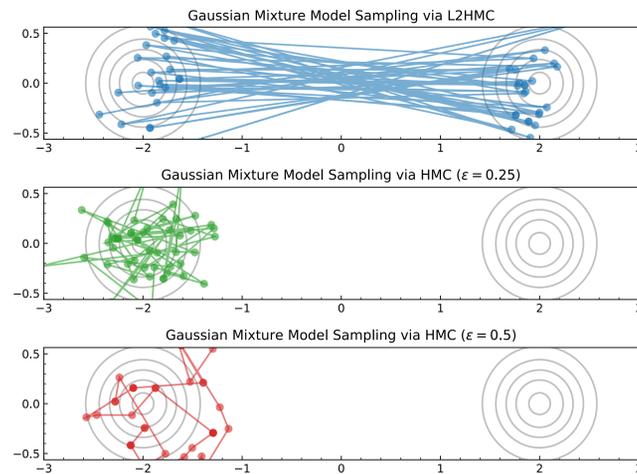


Figure (b.) Illustration of contours from the true target distribution vs samples obtained from both the trained model and generic HMC. We can see that HMC fails to mix between the two modes of the distribution whereas the trained model efficiently jumps between them.

Algorithm 1: Training Procedure

```

input:
1. Loss function,  $\mathcal{L}_\theta(\xi', \xi, A(\xi' | \xi))$ 
2. Batch of initial states,  $x$ 
3. Learning rate schedule,  $\{\alpha_t\}_{t=0}^{N_{train}}$ 
4. Annealing schedule,  $\{\gamma_t\}_{t=0}^{N_{train}}$ 
5. Target distribution,  $p_t(x) \propto e^{-\gamma_t S_\beta(x)}$ 

Initialize weights  $\theta$ 
for  $0 \leq t < N_{train}$ :
  resample  $v \sim \mathcal{N}(0, \mathbb{I})$ 
  resample  $d \sim \mathcal{U}(\pm, -)$ 
  construct  $\xi_0 \equiv (x_0, v_0, d_0)$ 
  for  $0 \leq k < N_{LF}$ :
    propose (leapfrog layer)  $\xi'_k \leftarrow \xi_k$ 
    compute  $A(\xi'_k | \xi_k) = \min \left\{ 1, \frac{p(\xi'_k)}{p(\xi_k)} \left| \frac{\partial \xi'_k}{\partial \xi_k} \right| \right\}$ 
    update  $\mathcal{L} \leftarrow \mathcal{L}_\theta(\xi'_k, \xi_k, A(\xi'_k | \xi_k))$ 
    backprop  $\theta \leftarrow \theta - \alpha_t \nabla_\theta \mathcal{L}$ 
  assign  $x_{t+1} \leftarrow \begin{cases} x' & \text{with probability } A(\xi'_k | \xi_k) \\ x & \text{with probability } (1 - A(\xi'_k | \xi_k)) \end{cases}$ 

```

Application to Lattice Gauge Theory

- Let $U_\mu(n) = e^{ix_\mu(n)} \in U(1)$, with $x_\mu(n) \in [-\pi, \pi]$ denote the *link variables*, where $x_\mu(n)$ is a link at the site n oriented in direction $\hat{\mu}$.
- Target distribution:** $p_t(x) \propto e^{-\gamma_t \cdot S_\beta(x)}$ • Introduce annealing factor γ_t , with $\|\gamma_t\| \leq 1$
- Wilson action:** $S_\beta(x) = \beta \sum_p 1 - \cos x_p$
- Topological Charge:** Each lattice has a charge $Q_Z \in \mathbb{Z}$ $Q_Z = \frac{1}{2\pi} \sum_p [x_p] \in \mathbb{Z}$, where: $[x_p] = x_p - 2\pi \left\lfloor \frac{x_p + \pi}{2\pi} \right\rfloor$
- Loss function:** We maximize the expected squared charge difference $\mathcal{L}(\theta) = \mathbb{E}_{p(\xi)} [-\delta(\xi', \xi) \cdot A(\xi' | \xi)]$ where $A(\xi' | \xi)$ is given in HMC (j.) $\delta(\xi', \xi) = [Q_R(x') - Q_R(x)]^2$

Results

- To measure the computational cost of our approach we use the *integrated autocorrelation time* $\tau_{int}^{Q_Z}$, which can be interpreted as the number of trajectories before an *independent sample* is drawn.
- We can see in Figure (e.) that the trained model consistently outperforms generic HMC across $\beta = 2, 3, \dots, 7$.

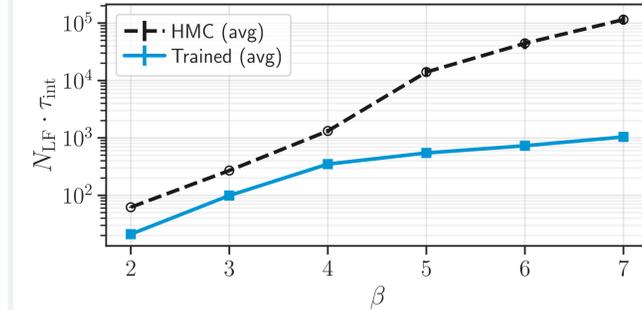


Figure (e.) Estimate of the integrated autocorrelation time $N_{LF} \cdot \tau_{int}^{Q_Z}$ vs β , scaled by N_{LF} to account for simulation cost.

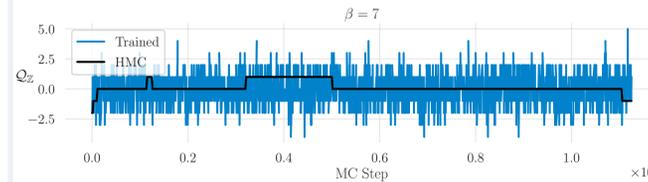


Figure (f.) Illustration of the integer valued topological charge Q_Z vs MC step for both HMC (black) and the trained model (blue).

- In order to understand the mechanism driving this improved behavior, we looked at how different physical quantities evolve during a single trajectory in the trained model as shown in Figures (g.1,2,3).
- We see that our sampler artificially *increased the energy* of the physical system during the first half of the trajectory, before returning back to its original physical value.
- We believe that this ability to vary the energy during the trajectory helps the sampler to overcome energy barriers between topological sectors whereas HMC remains stuck.

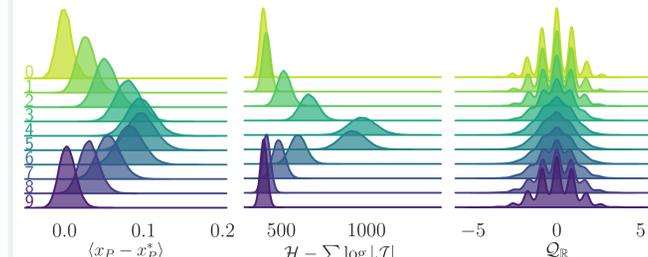


Figure (g.1) The variation in the average plaquette $\langle x_p - x_p^* \rangle$ at intermediate leapfrog layers. Figure (g.2) The adjusted energy $H - \log |J|$ at intermediate leapfrog layers. Figure (g.3) The real-valued topological charge Q_R at intermediate leapfrog layers.

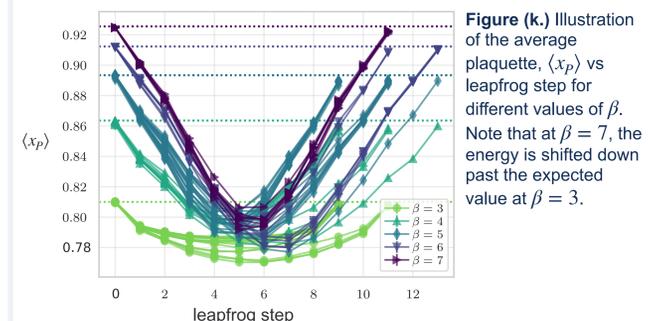


Figure (k.) Illustration of the average plaquette, $\langle x_p \rangle$ vs leapfrog step for different values of β . Note that at $\beta = 7$, the energy is shifted down past the expected value at $\beta = 3$.

Next Steps

- Going forward we plan to continue development of this approach towards more complex theories in higher space-time dimensions (e.g. 2D, 4D SU(3)).

Training Costs

- Our models were trained using Horovod on the ThetaGPU supercomputer at the Argonne Leadership Computing Facility (ALCF). A typical training run on 1 node ($8 \times$ NVIDIA A100 GPUs) using a batch size $M = 2048$, hidden layer shapes $[256, 256, 256]$ for each of the $N_{LF} = 10$ leapfrog layers, on a 16×16 lattice for 5×10^5 training steps takes roughly 24 hours to complete.

Conclusion

- Presented a generalized version of the L2HMC algorithm—consisting of a stack of leapfrog layers—that improves the existing approaches' flexibility while remaining statistically exact.
- Shown that our trained model successfully mixes between modes of a two-dimensional Gaussian Mixture Model while HMC remains stuck in a local mode.
- Looked at applying the described approach to a two-dimensional $U(1)$ lattice gauge theory.
- Saw that for this lattice gauge model, our trained sampler is capable of significantly outperforming traditional HMC across a range of coupling constants.

References

- Sam Foreman, Xiao-Yong Jin, James Osborn. [Deep Learning Hamiltonian Monte Carlo](#). (code: github.com/saforem2/l2hmc-qcd)
- Daniel Lévy, M. Hoffman, and Jascha Sohl-Dickstein. Generalizing Hamiltonian Monte Carlo with Neural Networks. [abs/1711.09268](https://arxiv.org/abs/1711.09268), 2018.
- Michael S Alberg, Denis Boyda, Daniel C Hackett, Gurtej Kanwar, Kyle Cranmer, Sébastien Racanière, Danilo Jimenez Rezende, and Phiala E Shanahan. Introduction to Normalizing Flows for Lattice Field Theory [arXiv:2101.08176](https://arxiv.org/abs/2101.08176), 2021.
- MS Alberg, G Kanwar, and PE Shanahan. Flow-based generative models for Markov Chain Monte Carlo in Lattice Field Theory. Physical Review D, [100\(3\):034515](https://doi.org/10.1103/PhysRevD.100.034515), 2019.
- Denis Boyda, Gurtej Kanwar, Sébastien Racanière, Danilo Jimenez Rezende, Michael S Alberg, Kyle Cranmer, Daniel C Hackett, and Phiala E Shanahan. Sampling using SU(n) Gauge Equivariant Flows. [arXiv:2008.05456](https://arxiv.org/abs/2008.05456), 2020.
- Gurtej Kanwar, Michael S Alberg, Denis Boyda, Kyle Cranmer, Daniel C Hackett, Sébastien Racanière, Danilo Jimenez Rezende, and Phiala E Shanahan. Equivariant Flow Based Sampling for Lattice Gauge Theory Physical Review Letters, [125\(12\):121601](https://doi.org/10.1103/PhysRevLett.125.121601), 2020

Hamiltonian Monte Carlo (HMC)

- Goal:** Sample from (difficult) target distribution: $p(x) \propto e^{-S(x)}$
- Method:**
 - For $x \in U(1)^n$, build chain $x_0 \rightarrow x_1 \rightarrow \dots \rightarrow x_N$ such that $x_N \sim p(x)$ as $N \rightarrow \infty$
 - Introduce $v \sim \mathcal{N}(0, I_n) \in \mathbb{R}^n$, write joint distribution: $p(x, v) = p(x)p(v) \propto e^{-S(x)} e^{-\frac{1}{2}v^T v} = e^{-H(x, v)}$
 - Evolve the system of equations $\dot{x} = \frac{\partial H}{\partial v}$, $\dot{v} = -\frac{\partial H}{\partial x}$ using the **leapfrog integrator (i.)** along $H = \text{const}$: $\xi \equiv (x, v) \rightarrow (x', v') = \xi'$
 - Accept or reject proposal configuration ξ' using **Metropolis-Hastings test (j.)**

(i.) Leapfrog update

- Half-step (v): $\tilde{v} = v - \frac{\epsilon}{2} \partial_x S(x)$
- Full-step (x): $x' = x + \epsilon \tilde{v}$
- Half-step (v): $v' = \tilde{v} - \frac{\epsilon}{2} \partial_x S(x')$

(j.) Metropolis-Hastings

$$x_{i+1} = \begin{cases} x' & \text{w/ prob } A(\xi' | \xi) \\ x & \text{w/ prob } 1 - A(\xi' | \xi) \end{cases}$$

where $A(\xi' | \xi) \equiv \min \left\{ 1, \frac{p(\xi')}{p(\xi)} \left| \frac{\partial \xi'}{\partial \xi} \right| \right\}$

¹foremans@anl.gov ²xjin@anl.gov ³osborn@alcf.anl.gov

²To obtain the expression for the reverse direction, we can invert each of the $\Gamma \equiv (\Gamma^+)^{-1}$, $\Lambda \equiv (\Lambda^+)^{-1}$ functions and perform the updates in the opposite order



Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.