

Learning Trivializing Flows

David Albandea,^{a,*} Luigi Del Debbio,^b Pilar Hernández,^a Richard Kenway,^b Joe Marsh Rossney^b and Alberto Ramos^a

^a*Instituto de Física Corpuscular (CSIC – University of Valencia), Parque Científico, C/Catedrático José Beltrán, 2, 46980, Paterna, Valencia, Spain*

^b*Higgs Centre for Theoretical Physics, School of Physics and Astronomy, The University of Edinburgh, Edinburgh EH9 3FD, UK*

E-mail: david.albandea@ific.uv.es

The recent introduction of machine learning techniques, especially normalizing flows, for the sampling of lattice gauge theories has shed some hope on improving the sampling efficiency of the traditional HMC algorithm. Naive use of normalizing flows has been shown to lead to bad scaling with the volume. In this talk we propose using local normalizing flows at a scale given by the correlation length. Even if naively these transformations have a small acceptance, when combined with the HMC algorithm lead to algorithms with high acceptance, and also with reduced autocorrelation times compared with HMC. Several scaling tests are performed in the ϕ^4 theory in 2D.

*The 39th International Symposium on Lattice Field Theory (Lattice2022),
8-13 August, 2022
Bonn, Germany*

*Speaker




Figure 1: Extracted from [1]: normalizing flow sketch. A set of configurations z_i is generated from a trivial probability distribution $r(z)$; a neural network model f is used to generate a new set of configurations via $\phi = f^{-1}(z)$, which follow $\tilde{p}_f(\phi) \approx p(\phi)$, with \tilde{p}_f the model's output distribution and p the target distribution.

1. Introduction

1.1 Normalizing flows

Normalizing flows are a machine learning sampling technique introduced to lattice field theories in [1]. As sketched in Fig.1, they propose a flow-based Markov Chain Monte Carlo algorithm whose workflow is

1. Generate a set of configurations z_i following a trivial probability distribution $r(z)$.
2. Apply a function f^{-1} to all configurations z_i , obtaining a new set of configurations ϕ_i via $\phi = f^{-1}(z)$, following a new probability distribution $\tilde{p}_f(\phi)$. f is a neural network which has been trained so that the new probability distribution \tilde{p}_f is as similar as possible to our target distribution $p \propto e^{-S}$, the distribution of the theory to be studied.
3. Use a Metropolis–Hastings accept-reject to correct for the bias in the approximation $\tilde{p}_f \approx p$.

The training of the neural network f is done by minimizing the Kullback-Leibler (KL) divergence,

$$D_{\text{KL}}(\tilde{p}_f || p) = \int \mathcal{D}\phi \tilde{p}_f(\phi) \log \frac{\tilde{p}_f(\phi)}{p(\phi)}, \quad (1)$$

which is a statistical distance satisfying $D_{\text{KL}}(\tilde{p}_f || p) \geq 0$ and $D_{\text{KL}}(\tilde{p}_f || p) = 0 \iff \tilde{p}_f = p$. Therefore, after training the network f is expected to be an approximate trivializing map [2].

As shown in Fig.2 (left), the good thing about this algorithm is that autocorrelation times in the resulting Markov chain do not scale when taking the continuum limit if the neural networks are trained up to the same reference acceptance, meaning that the cost of producing configurations does not scale; on the other hand, autocorrelations with the Hybrid Monte Carlo (HMC) algorithm are expected to scale as $\sim \xi^2$.

However, if one wants to study the scaling of the total computational cost, one needs to analyze the training costs as well, since

$$\text{total cost} = \text{configuration production cost} + \text{network training cost}. \quad (2)$$

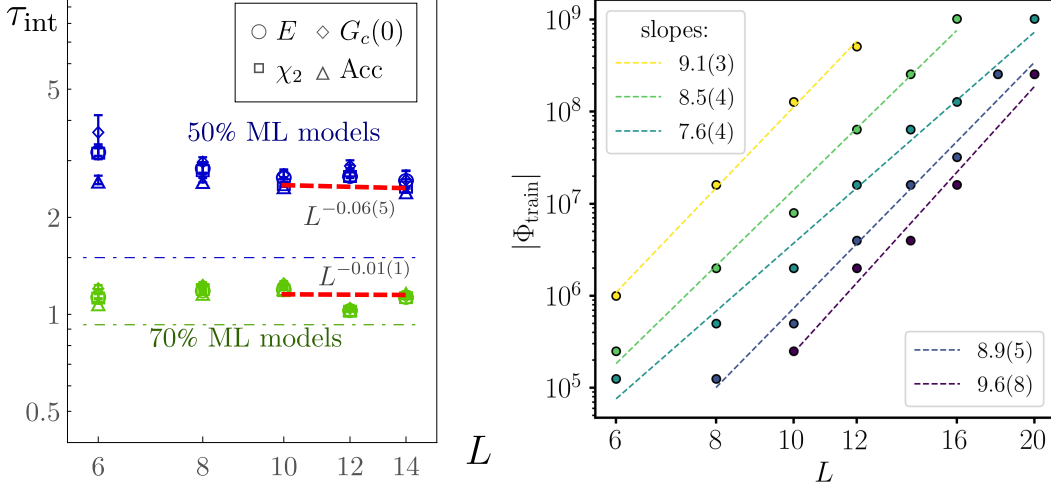


Figure 2: (Left) Extracted from [1]: autocorrelation time scaling for different observables using normalizing flows with networks trained to have Metropolis acceptances of 50% and 70%. (Right) Extracted from [3]: scaling of the number of configurations needed to train a network up to an acceptance of 70%.

It was shown in [3] that the cost of keeping a reference Metropolis acceptance of 70% seems to scale approximately as $\sim \xi^8$ (see Fig.2), indicating a transfer of the critical slowing down problem from the production of configurations to the training cost of the networks.

In what follows we want to address whether one can still benefit from normalizing flows keeping the training costs low.

2. Learning Trivializing Flows

2.1 The algorithm

The main idea is to use Lüscher’s trivializing flows [2], so that we can use the normalizing flows to *help* the HMC algorithm, rather than replacing it. For example, let’s consider the partition function of our target theory,

$$Z = \int \mathcal{D}\phi e^{-S[\phi]}. \quad (3)$$

We can make a change of variables $\phi \rightarrow \tilde{\phi}$ using our trained network f ,

$$\tilde{\phi} = f(\phi), \quad \phi = f^{-1}(\tilde{\phi}), \quad \mathcal{D}\phi = \left| \frac{\partial f^{-1}(\tilde{\phi})}{\partial \tilde{\phi}} \right| \mathcal{D}\tilde{\phi} \equiv \det J[f^{-1}(\tilde{\phi})] \mathcal{D}\tilde{\phi}, \quad (4)$$

so that the partition function becomes

$$Z = \int \mathcal{D}\tilde{\phi} e^{-S[f^{-1}(\tilde{\phi})] + \log \det J[f^{-1}(\tilde{\phi})]} \equiv \int \mathcal{D}\tilde{\phi} e^{-\tilde{S}[\tilde{\phi}]}, \quad (5)$$

where we have defined the new action

$$\tilde{S}[\tilde{\phi}] \equiv S[f^{-1}(\tilde{\phi})] - \log \det J[f^{-1}(\tilde{\phi})]. \quad (6)$$

This new action is a combination of the old action $S[\phi]$ and the logarithm of the Jacobian of the variable transformation of Eq.(4). If the Jacobian cancels out part of the action, then the probability distribution $e^{-\tilde{S}[\tilde{\phi}]}$ might be easier to sample from than $e^{-S[\phi]}$, and using HMC with the new action might yield lower autocorrelation times. The workflow of the algorithm would then be

1. Train the network f by minimizing the KL divergence.
2. Run the HMC algorithm to build a Markov chain of configurations following $\tilde{p}(\tilde{\phi}) = e^{-\tilde{S}[\tilde{\phi}]}$,

$$\{\tilde{\phi}_1, \tilde{\phi}_2, \tilde{\phi}_3, \dots, \tilde{\phi}_N\} \sim e^{-\tilde{S}[\tilde{\phi}]}.$$

3. Apply the inverse transformation f^{-1} to every configuration in the Markov chain to undo the variable transformation. This way we obtain a Markov chain of configurations following the target probability distribution $p(\phi) = e^{-S[\phi]}$,

$$\{f^{-1}(\tilde{\phi}_1), f^{-1}(\tilde{\phi}_2), f^{-1}(\tilde{\phi}_3), \dots, f^{-1}(\tilde{\phi}_N)\} = \{\phi_1, \phi_2, \phi_3, \dots, \phi_N\} \sim e^{-S[\phi]}.$$

The important point is that the acceptance of this algorithm does not depend on the transformation f that we are doing: it only depends on how well you integrate the HMC equations of motion. This means that the algorithm will work, no matter how well one trains the network f .

Lüscher proposed this algorithm using the Wilson flow as an approximate trivializing map [2], but it was not good enough to improve the scaling of autocorrelation times towards the continuum in a CP^{N-1} theory with topology [4]. The hope is that normalizing flows can play a better role as approximate trivializing maps; indeed, this idea has already been tested [5, 6], but we want to focus on the scaling of cheap training setups.

2.2 The model

As in [1, 3], we worked with a ϕ^4 theory with a massive scalar field in 2 dimensions,

$$S[\phi] = \sum_x \left[-\beta \sum_{\mu=1}^2 \phi_{x+\mu} \phi_x + \phi_x^2 + \lambda(\phi_x^2 - 1)^2 \right]. \quad (7)$$

Among its features, it has a \mathbb{Z}_2 symmetry, since the action is invariant under a sign flip of the scalar field, $\phi \rightarrow -\phi$; the probability density of the model has two modes, which correspond to positive and negative magnetization $M = \frac{1}{V} \sum_x \phi_x$; and it has a non-trivial correlation length ξ , yielding autocorrelations when building the Markov chain of configurations.

We will use these autocorrelations to benchmark our new algorithm against HMC, since the autocorrelation times of HMC are expected to scale as $\tau_{\text{int}} \propto \xi^2$. Also, this model does not have topology, as opposed to QCD, so we will not suffer from topology freezing.

2.3 Keeping training costs low

Our intention is to keep training costs as low as possible so that we can say that the total computational cost is essentially given by the cost of producing the Markov chain of configurations,

$$\text{total cost} \approx \text{configuration production cost} \quad (8)$$

For this to happen it is helpful to consider the information we know about the model:




Figure 3: (Left) History of the KL divergence during the training from independent Gaussians to a theory with parameters $\beta = 0.641$, $\lambda = 0.5$, lattice size $L = 18$ and $\xi = L/4$. (Right) History of the magnetization for a simulation with HMC (blue) and trivializing flows (orange).

- The action has translational symmetry, so there should be no difference between the transformation of the field at x and the transformation at any other point x' . This indicates that one should use convolutional neural networks (CNN) instead of fully connected networks, so that the same transformation kernel is applied at every point.
- The relevant physics is contained within the correlation length ξ of the system, so the transformation of ϕ_x should not depend on $\phi_{x'}$ if $|x - x'| \gg \xi$. This indicates that one should make the footprint of the network as small as possible according to the correlation length ξ .

With this in mind we studied very simple network architectures with only one affine coupling layer and no hidden layers. The affine transformation is defined as

$$\phi_x \rightarrow e^{s(\phi)} \phi_x + t(\phi), \quad (9)$$

where $s(\phi)$ and $t(\phi)$ are CNNs with kernel size k , which can be varied to control the footprint of the transformation. For the simplest case of $k = 3$ and a pair of checkerboard-masked affine layers [1], the transformation would only couple next-to-nearest neighbors and the neural network would have only 37 parameters¹.

With such a simple network the training cost is negligible with respect to the HMC simulation, and one could compare both algorithms just by studying the scaling of the autocorrelation time τ_{int} towards the continuum.

3. Results

3.1 Reduction of τ_{int} with minimal network architecture

A first thing to check is if a network with only 37 parameters can learn something. In Fig.3 (left) we show the evolution of the KL divergence for the training of a network from independent

¹A CNN with 1 layer has k^2 parameters. The transformation in Eq.(9) with N_l pairs of checkerboard-masked affine layers has $2 \times N_l$ different CNNs, and therefore $2 \times N_l \times k^2$ parameters. We also add a global rescaling parameter as a final layer of our network, so our models have $N_p = 2 \times N_l \times k^2 + 1$ parameters.

L	Acc. at L	Acc. at $2L$
3	0.3	0.2
4	0.04	0.001
5	0.002	0.00003
6	0.002	0.000007
7	0.0001	$< 10^{-7}$
8	0.0001	-
9	0.00007	-
10	0.00004	-

Table 1: Metropolis acceptances at lattice sizes L and $2L$ of networks trained at lattice size L .




Figure 4: Autocorrelation time of the magnetization at lattice size $2L$ using trivializing flows. In circles, the networks used were trained at lattice size $2L$; in triangles, they were trained at L and used at $2L$.

Gaussians to a theory with parameters $\beta = 0.641$, $\lambda = 0.5$, lattice size $L = 18$ and $\xi = L/4$. We see that D_{KL} saturates fast, and this is because the network is very simple: this is the best thing one can do with it, no matter how much longer is trained.

Having trained the network, one can run a simulation with the two algorithms: one with plain HMC on the action $S(\phi)$ —denoted as *HMC* from now on— and another with HMC on the transformed action $\tilde{S}(\tilde{\phi})$ —denoted as *trivializing flow*. The Monte Carlo history of the magnetization for both algorithms is shown in Fig.3 (right), and their respective autocorrelation times are

$$\tau_{\text{flow}} = 74.4(3), \quad \tau_{\text{HMC}} = 100.4(2). \quad (10)$$

Since the trivializing flows reduced the autocorrelation time of HMC, this indicates that a simple network can indeed learn something.

3.2 Infinite volume limit

Another thing to note is that CNNs allow to take a network trained at a lattice size L and reuse it for a bigger lattice size $L' > L$. In Tab.1 we trained a different network for every value of L and checked the Metropolis acceptance that they would have using Metropolis reweighting directly, as in [1, 3]. Then, we reused the network —without further retraining— for a lattice size $2L$ and checked again their Metropolis acceptance.

One can note that the acceptances are low, meaning that the output distribution \tilde{p}_f of our network is not a good approximation of the target distribution p , which was expected because the networks are very simple; and that reusing the network for a larger lattices decreases the acceptance a lot, which is also expected because the action is an extensive quantity and we are just doing reweighting for a bigger volume.

However, none of this matters if the network is used as a trivializing flow. We can see this in Fig.4, where we show the autocorrelation time of the magnetization versus different lattice sizes;




Figure 5: (Left) Scaling of the autocorrelation time of the magnetization towards the continuum for HMC (filled, blue circles) and trivializing flows with kernel sizes $k = 3, 5, 7$ (open circles, triangles and hexagons). (Right) Scaling of the ratio of autocorrelation times of the magnetization of HMC with respect to trivializing flows.

whether you train directly at a lattice size $2L$ or you train at L and reuse it for $2L$, you get the same autocorrelation times. This indicates that the training should be done at a lattice size similar to the correlation length of the system, $L \sim \xi$, where the physical information is contained, to reduce training costs.

3.3 Computational cost scaling

3.3.1 The setup

Finally, we study how the computational cost of the machine-learned trivializing flow scales towards the continuum, comparing it with the costs of HMC. The theory parameters used are in the table below, where β was tuned to fix the physical size of the lattice to $\xi = L/4$.

L	6	8	10	12	14	16	18	20	40	80
β	0.537	0.576	0.601	0.616	0.626	0.634	0.641	0.645	0.667	0.677
λ	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5

A few things are worth noting:

- For each of the columns we train a different network from independent Gaussians to the target theory.
- We always used simple network architectures with one coupling layer and no hidden layers, and we trained them until saturation of the KL divergence. This saturation happens very fast and the training costs are negligible with respect to the cost of the production of configurations. Therefore we only need to compare the scaling of the autocorrelation times of both algorithms.
- The integration step for the molecular dynamics evolution is tuned so that the Metropolis–Hastings acceptance of HMC and the trivializing flows is approximately 90%.




Figure 6: (Left) Scaling of the autocorrelation time of the magnetization with $k \sim \xi$. (Right) Scaling of the autocorrelation time of the flowed one-point susceptibility with $k \sim \xi$. The fits correspond to the assumption $\tau_{\text{int}} \propto \xi^z$.

3.3.2 Scaling with fixed architecture

In Fig.5 (left) we compare the autocorrelation times of the magnetization from simulations with both algorithms: in blue circles we have the autocorrelations of HMC, and the rest of points are from trivializing flow simulations with kernel sizes $k = 3, 5, 7$.

First, we see that all autocorrelations from trivializing flows are better than the ones of HMC. Also, for a fixed lattice size L , increasing the kernel size of the network decreases the autocorrelation, but it does not change the scaling of HMC.

The latter can be seen better if we plot the ratio $\tau_{\text{HMC}}/\tau_{\text{flow}}$, as we do in Fig.5 (right): as we go to the continuum the ratio tends to a constant, indicating that the scaling of both algorithms is the same for a fixed network architecture.

3.3.3 Scaling increasing the kernel size

The correlation length ξ increases in lattice units when taking the continuum limit, so one should also scale the kernel size of our transformations so that the transformation acts on the same physical region, $k \sim \xi$. When one chooses the optimum kernel size k of the CNNs for each lattice size, one gets the results displayed in Fig.6 (left), where now the gap between HMC and the trivializing flow algorithm seems to increase towards the continuum.

Assuming that the autocorrelation length scales as $\tau \propto \xi^z$ we can fit our results and compare the exponent of the two algorithms, obtaining

$$z_{M,\text{HMC}} = 2.20(4), \quad z_{M,\text{flow}} = 1.97(7). \quad (11)$$

Therefore, scaling the kernel size of the CNNs leads to a slight improvement in the autocorrelation scaling of the magnetization $M = 1/V \sum_x \phi_x$, which is a local operator. In Fig.6 (right) we plot instead the one-point susceptibility $\chi_t = \frac{1}{V} \sum_x \phi_{t,x}^2$ measured on configurations flowed up to flow time t with a smearing radius $R \sim \xi$. Assuming the same scaling we again find a slight improvement with respect to HMC:

$$z_{\chi_t,\text{HMC}} = 2.20(2), \quad z_{\chi_t,\text{flow}} = 1.92(4). \quad (12)$$

4. Conclusions and outlook

We have shown that, even working with very simple architectures, using neural networks as trivializing flows can improve the autocorrelation times of HMC, although the scaling is the same as HMC for a fixed network architecture.

Also, the networks trained at a small lattice size can be reused for larger volumes without further training. Focusing on topology freezing, this could be useful in QCD: one could train at the size of Λ_{QCD} for large values of the quark masses and then reuse the network for larger volumes and small masses, thus reducing training costs.

Finally, scaling the kernel size of the networks slightly improves the scaling of autocorrelations. An interesting question would then be if machine-learned trivializing flows could improve a much worse kind of continuum scaling, such as the one related to topological freezing in theories with topology.

Acknowledgments

We acknowledge support from the Generalitat Valenciana grant PROMETEO/2019/083, the European project H2020-MSCA-ITN-2019//860881-HIDDeN, and the national project PID2020-113644GB-I00. AR acknowledges financial support from Generalitat Valenciana through the plan GenT program (CIDEAGENT/2019/040). DA acknowledges support from the Generalitat Valenciana grant ACIF/2020/011. JMR is supported by STFC grant ST/T506060/1.

This work has been performed under the Project HPC-EUROPA3 (INFRAIA-2016-1-730897), with the support of the EC Research Innovation Action under the H2020 Programme; in particular, we gratefully acknowledge the support of the computer resources and technical support provided by EPCC. This work used the ARCHER2 UK National Supercomputing Service (<https://www.archer2.ac.uk>). We also acknowledge the computational resources provided by Finis Terrae II (CESGA), Lluís Vives (UV), Tirant III (UV). The authors also gratefully acknowledge the computer resources at Artemisa, funded by the European Union ERDF and Comunitat Valenciana, as well as the technical support provided by the Instituto de Física Corpuscular, IFIC (CSIC-UV).

References

- [1] M. S. Albergó, G. Kanwar, and P. E. Shanahan. Flow-based generative models for markov chain monte carlo in lattice field theory. *Phys. Rev. D*, 100:034515, Aug 2019.
- [2] M. Lüscher. Trivializing maps, the wilson flow and the hmc algorithm. *Commun. Math. Phys.*, 293, 2010.
- [3] Luigi Del Debbio, Joe Marsh Rossney, and Michael Wilson. Efficient modeling of trivializing maps for lattice ϕ^4 theory using normalizing flows: A first look at scalability. *Phys. Rev. D*, 104(9):094507, 2021.
- [4] Georg P. Engel and Stefan Schaefer. Testing trivializing maps in the hybrid monte carlo algorithm. *Computer Physics Communications*, 182(10):2107–2114, 2011.

- [5] Xiao-yong Jin. Neural Network Field Transformation and Its Application in HMC. *PoS, LATTICE2021:600*, 2022.
- [6] Sam Foreman, Taku Izubuchi, Luchang Jin, Xiao-yong Jin, James C. Osborn, and Akio Tomiya. HMC with Normalizing Flows. *PoS, LATTICE2021:073*, 2022.