
Equivariant Hamiltonian Flows

Danilo J. Rezende*

Sébastien Racanière*

Irina Higgins*

Peter Toth*

*{danilor, sracaniere, irinah, peterth}@google.com

Abstract

This paper introduces equivariant hamiltonian flows, a method for learning expressive densities that are invariant with respect to a known Lie-algebra of local symmetry transformations while providing an equivariant representation of the data. We provide proof of principle demonstrations of how such flows can be learnt, as well as how the addition of symmetry invariance constraints can improve data efficiency and generalisation. Finally, we make connections to disentangled representation learning and show how this work relates to a recently proposed definition.

1 Introduction

Learning generative models with structured latent representations is important for interpretability, data efficiency, and generalisation [2, 7]. One kind of structural bias that has been shown to work well in the past is that of invariance or equivariance with respect to a group of transformations. For example, convolutional neural networks [21] are invariant to the group of translations by construction, while disentangled representations learn equivariance to more general transformations [10]. In order to learn such structured latent representations, however, often a trade-off has to be made in terms of latent expressivity. For example, most of the disentangling models to date use a unit Gaussian prior [9, 14, 3]. Normalizing flows [27, 11, 5, 18, 4, 25, 16] provide a simple mechanism to build expressive density estimators in problems without much domain knowledge. But it remains a challenge to embed flow-based models with domain knowledge such as data on specific manifolds [6] and known symmetries or factorisation of the target density. In this paper we focus on problems where there is domain knowledge in the form of known invariances of a target density which we wish to learn. This is the typical case in many modelling problems, such as density over molecular structures, lattice-QCD [1] with internal gauge symmetries (e.g. $U(n)$ and $SU(n)$), multi-particle systems [20], etc. The challenge we address is the following: suppose that we start from a base distribution $\pi(z)$ that is invariant with respect to a group of transformations. If we transform this density via a generic flow $f(z)$, there will be no guarantees that the transformed density $p(z) = \pi(f^{-1}(z))/\det|Jac|$ will be invariant as well.

Our main contributions are: (i) Introduce equivariant Hamiltonian flows; (ii) Propose a general algorithm for enforcing equivariance in learned Hamiltonian flows with respect to any connected Lie-Group; and (iii) Prove a simple lemma that allow us to construct invariant densities from equivariant flows. The result of this are flows that transform a simple density that is invariant with respect to the actions of a known symmetry group, into another invariant density that is arbitrarily complex. We demonstrate that learning latent representations with the known equivariance structure helps with data efficiency and generalisation. Finally, we connect this work to the recent definition of disentangled representations [10].

2 Methods

We first revise the notion of Hamiltonian dynamics from physics [8]. In physics the Hamiltonian formalism is used to model energy conserving continuous dynamics in an abstract state space $s = (q, p)$ of generalised position q and momentum p . Formal connections can be made between Hamiltonian




Figure 1: Model diagrams. From left to right: (i) Equivalence between flows with swapped T_H^{dt} and T_g^ϵ , as a result of the equivariance of Hamiltonian flows with respect to generators of symmetries of the Hamiltonian function; (ii) Generative model diagram. The last momentum variable p_n (shaded node) is dropped during sampling; (iii) Inference mode diagram for computing the ELBO, where $h(p_n|q_n)$ is an encoder model. The last momentum variable p_n is inferred during training.

dynamics and the actions of continuous symmetry generators – those transformations that leave the dynamics of the system modelled by the Hamiltonian unchanged throughout the path traversed. Indeed, the roles of the Hamiltonian and its symmetry generators can be interchangeable, whereby a symmetry generator of one system can be a Hamiltonian for a different system, and both Hamiltonian and symmetry transformations commute with each other. The commutativity of transformations is measured by an operator known as the Poisson bracket. Hence, we use the Poisson bracket between the Hamiltonian and the known symmetry generators as a regularizer to restrict the form of the learnt Hamiltonian flow, whereby it models the final data density well, while also being invariant to the known symmetries.

Hamiltonian Flow Generative Model A Hamiltonian flow is a continuous-time normalizing flow induced by the Hamiltonian dynamics in the state-space $s = (q, p) \in \mathbb{R}^{2d}$ via the ODE $(\dot{q}, \dot{p}) = \{(q, p), H(q, p)\} = (\frac{\partial H}{\partial p}, -\frac{\partial H}{\partial q})$, where $\{, \}$ is a skew-symmetric differential operator known as the Poisson bracket or commutator and $H: \mathbb{R}^{2d} \rightarrow \mathbb{R}$ is a scalar function known as Hamiltonian. For two scalar functions f, g we set $\{f(q, p), g(q, p)\} = \sum_i \frac{\partial f(q, p)}{\partial q_i} \frac{\partial g(q, p)}{\partial p_i} - \frac{\partial f(q, p)}{\partial p_i} \frac{\partial g(q, p)}{\partial q_i}$. If f or g is a vector valued function, we extend the previous definition component wise so that $\{f, g\}$ is also a vector value function.

In what follows, we introduce the notation $T_f^\epsilon(x) = x + \epsilon\{x, f(x)\}$ to indicate an infinitesimal transformation induced by the Poisson bracket with a function f . With this notation, the Euler discretization of the Hamiltonian ODE is written as $s_{t+dt} = T_H^{dt}(s_t) = s_t + dt\{s_t, H(s_t)\}$. We can construct a more complex flow by chaining several transformations $s' = T_{H_n}^{dt} \circ \dots \circ T_{H_1}^{dt}(s)$. The inverse of this flow is obtained by replacing dt by $-dt$ and reversing the order of the transformations, $s = T_{H_1}^{-dt} \circ \dots \circ T_{H_n}^{-dt}(s') + O(dt^2)$.

We can create an expressive probability density $p_\theta(s_n)$ on $s_n = (q_n, p_n)$ by starting from an initial distribution $\pi(s_0)$, and transforming this density via the flow $s_n = T_{H_n}^{dt} \circ \dots \circ T_{H_1}^{dt}(s_0)$, where the Hamiltonians H_i are scalar functions parametrized by neural networks with parameters θ . This constitutes a more structured form of Neural ODE flow, [4]. There are a few notable differences between Hamiltonian flows and general ODE flows: (i) The Hamiltonian ODE is volume-preserving¹, which makes the computation of log-likelihood cheaper than for a general ODE flow. We show in Figure 2 that this does not necessarily reduce the expressivity of $p(q_n)$ even if $p(s_n)$ may be restricted; (ii) General ODE flows are only invertible in the limit $dt \rightarrow 0$, whereas for some Hamiltonians, we can use symplectic integrators (such as Leap-Frog, [23]) that are both invertible and volume-preserving for any $dt > 0$. The resulting density $p_\theta(s_n)$ is given by $\ln p(s_n) = \ln \pi(s_0) = \ln \pi(T_{H_1}^{-dt} \circ \dots \circ T_{H_n}^{-dt}(s_n)) + O(dt^2)$.

The structure $s = (q, p)$ on the state-space imposed by the Hamiltonian dynamics can be constraining from the point of view of density estimation, since this would require an artificial split of the data vectors in two disjoint sets. We could consider different mechanisms to address this: (i) Exploit the splits and use alternating masks similar to [5]; (ii) Treat the momentum variables p as latent variables. The latter is the same interpretation as in HMC [23, 26, 22]. It is also more elegant as it does not require an artificial split of the data. This results in a density $p(q_n)$ of the form $p(q_n) = \int dp_n p(q_n, p_n) = \int dp_n \pi(T_{H_1}^{-dt} \circ \dots \circ T_{H_n}^{-dt}(q_n, p_n))$. This integral is intractable, but the model can still be trained via variational methods where we introduce a variational density $h_\phi(p_n|q_n)$ with parameters ϕ and optimise the ELBO,

$$\text{ELBO}(q_n) = \mathbb{E}_{h_\phi(p_n|q_n)} [\ln \pi(T_{H_1}^{-dt} \circ \dots \circ T_{H_n}^{-dt}(q_n, p_n)) - \ln h_\phi(p_n|q_n)] \leq \ln p(q_n), \quad (1)$$

¹ This can be seen by computing the determinant of the Jacobian of the infinitesimal transformation T_H^{dt} , $\det(\text{Jac}) = \det \left[\mathbb{I} + dt \begin{pmatrix} \frac{\partial^2 H}{\partial q_i \partial p_j} & -\frac{\partial^2 H}{\partial q_i \partial q_j} \\ \frac{\partial^2 H}{\partial p_i \partial p_j} & -\frac{\partial^2 H}{\partial p_i \partial q_j} \end{pmatrix} \right] = 1 + dt \text{Tr} \begin{pmatrix} \frac{\partial^2 H}{\partial q_i \partial p_j} & -\frac{\partial^2 H}{\partial q_i \partial q_j} \\ \frac{\partial^2 H}{\partial p_i \partial p_j} & -\frac{\partial^2 H}{\partial p_i \partial q_j} \end{pmatrix} + O(dt^2) = 1 + O(dt^2)$.

instead. Note that, in contrast to VAEs ([17, 13]), the ELBO (1) is not explicitly in the form of a reconstruction error term plus a KL term.

Equivariant Hamiltonian Flows and Symmetries Let’s say we want to learn a density $p_\theta(x)$ and we also want it to be invariant with respect to a set of transformations $x' = T_\omega(x)$ induced by the elements ω of a symmetry group G . That is, $p(x') = p(T_\omega(x)) = p(x), \forall \omega \in G$. In the following we will assume that G is a connected Lie group². This will allow us to introduce infinitesimal transformations, and being invariant with respect to transformations in G will be equivalent to being invariant with respect to those infinitesimal transformations.

The Hamiltonian formalism provides a natural language to manipulate symmetries and, in particular, it provides a connection between invariance of the Hamiltonian function to the equivariance of the Hamiltonian flow. We can parametrize the action of an infinitesimal symmetry transformation on the state-space $s = (q, p)$ by a scalar function $g(s)$ via the Poisson bracket as $s' = s + \epsilon \{s, g\}$, where $\epsilon > 0, \epsilon \ll 1$. This can also be represented as $s' = T_g^\epsilon(s)$. In this context, a transformation T_g^ϵ is defined to be an infinitesimal symmetry if $T_g^\epsilon \circ T_H^{dt} = T_H^{dt} \circ T_g^\epsilon + O(\epsilon dt^2 + \epsilon^2 dt)$. That is, we can interchange the composition order of the Hamiltonian flow by the symmetry transformation. This is illustrated in Figure 1(left). This means that all intermediate states $s_{1, \dots, n}$ of the Hamiltonian flow will also transform in the same manner as the starting state, that is, they will form an *equivariant* or *covariant* representation of the starting state. When T_g^ϵ is a symmetry, the function g is called a *symmetry generator*. By Noether’s first Theorem [24], if T_g^ϵ is a symmetry then the generator g commutes with the Hamiltonian function, that is $\{H, g\} = 0$ ³. This implies that the numerical value of the function g will *remain constant throughout the Hamiltonian flow*, for this reason the symmetry generators g are also referred to as *conserved charges* in physics. Based on this reasoning, we say that a Hamiltonian flow is equivariant if $\{H, g\} = 0$. The set of all symmetry generators is closed under the Poisson bracket. That is, if two functions f and g are symmetry generators then $\{f, g\}$ is a symmetry generator. This means that the set of all generators forms a *Lie algebra*.

If we can express our domain knowledge about the target density as: (i) a set of symmetry generators $g_k, k = 1, \dots, N_g$ and (ii) an initial simple invariant density π , we can construct a new density p via Hamiltonian flow that is also invariant by learning a Hamiltonian such that $\{g_k, H\} = 0$. This is formalised in Lemma 1.

Lemma 1. *Given a Hamiltonian function $H : \mathbb{R}^{2d} \rightarrow \mathbb{R}$, a set of symmetry generators $g_k : \mathbb{R}^{2d} \rightarrow \mathbb{R}$ and a base density $\pi : \mathbb{R}^{2d} \rightarrow \mathbb{R}^+$, the density $p : \mathbb{R}^{2d} \rightarrow \mathbb{R}^+$ induced by the Hamiltonian flow T_H^{dt} will be invariant with respect to the generators g_k if $\{g_k, H\} = \{g_k, \pi\} = 0, \forall k$.*

Proof. From Noether’s theorem, $\{g_k, H\} = 0$ implies the equivariance of the flow $T_g^\epsilon \circ T_H^{dt} = T_H^{dt} \circ T_g^\epsilon + O(\epsilon dt^2 + \epsilon^2 dt)$. It remains to prove that $\pi(T_g^\epsilon(s)) = \pi(s) + O(\epsilon^2)$. Expanding to first order in ϵ , we have $\pi(T_g^\epsilon(s)) = \pi(s) + \epsilon \nabla \pi^T \{s, g_k(s)\} + O(\epsilon^2) = \pi(s) + \epsilon \{\pi, g_k(s)\} + O(\epsilon^2) = \pi(s) + O(\epsilon^2)$. \square

We can enforce $\{g_k, H\} = 0$ via constrained optimisation, where we perform a min-max optimisation of the Lagrangian $L(\theta, \phi, \lambda)$ instead,

$$(\theta^*, \phi^*, \lambda^*) = \min_{\theta, \phi} \max_{\lambda \geq 0} L(\theta, \phi, \lambda); \quad L(\theta, \phi, \lambda) = - \sum_{x \in \text{Data}} \text{ELBO}(x) + \sum_k \lambda_k \mathbb{E}_\pi [\{g_k(s), H(s)\}^2 - \kappa],$$

where λ_k is a Lagrange multiplier for the k th generator and κ controls the precision of the constraint.

In lemma 1, we found how to get the density over the full state s to be invariant under the symmetries. What we want though, is to have the marginal distribution over q to be invariant. This will happen when the Hamiltonian $H(q, p)$ is factored as a "kinetic" $K(p)$, and "potential" $U(q)$ energy terms, $H(q, p) = K(p) + U(q)$.⁴

²See [19] for an introduction to Lie groups and their infinitesimal counterparts, the Lie algebras. Readers unfamiliar with these concepts should think of Lie groups as abstract equivalent of groups of matrices. For example, the group of orientation preserving rotations $\text{SO}(n)$ is a Lie group, and its infinitesimal counterpart (its Lie algebra) is the space of anti-symmetric matrices.

³This can be easily proven by observing that $T_g^\epsilon \circ T_H^{dt} - T_H^{dt} \circ T_g^\epsilon = \epsilon dt \{g, H\} + O(\epsilon dt^2 + \epsilon^2 dt)$.

⁴This condition puts a constraint on the Hamiltonian to ensure that any learnt generator will induce an action on q such that the marginal distribution is invariant under that action. Alternatively, we could have put a constraint on




Figure 2: Learning with Hamiltonians $H(q, p) = K(p) + U(q)$. **A**: Learning a density with known $SO(2)$ symmetry. First two columns: training and test log-likelihoods of models with different constraint precision κ in the regime of infinite and finite data. Adding the symmetry constraint increases both train and test data efficiency. In the finite data regime it prevents over-fitting. Second two columns: KDE estimators of the target and learnt densities; learned kinetic energy $K(p)$ and potential energy $U(q)$. **B**: Multimodal density learning. First two columns: KDE estimators of the target and learnt densities; learned kinetic energy $K(p)$ and potential energy $U(q)$. Last column: single Leap-Frog step and integrated flow. The potential energy learned multiple attractors, also clearly visible in the integrated flow plot. The basins of attraction are centred at the modes of the data.

Connections to disentangling: Given a decomposition of a continuous group into a direct product of subgroups, the actions of the generators of these subgroups can be modelled by independent generators T_w^ϵ and T_g^ϵ such that $\{T_w^\epsilon, T_g^\epsilon\} = 0$. By definition in [10], these will act on independent subspaces of a disentangled q . Since equivariant flows preserve the action of the generators, they will preserve the disentanglement of the representation.

3 Experiments

For all experiments the Hamiltonian was of the form $H(q, p) = K(p) + U(q)$. The kinetic energy term K and the potential energy term U are soft-plus MLPs⁵ with layer-sizes $[d, 128, 128, 1]$ where d is the dimension of the data. The encoder network was parametrized as $h(p|q) = N(p; \mu(q), \sigma(q))$, where μ and σ are relu MLPs with size $[d, 128, 128, d]$. The Hamiltonian flow T_H^{dt} , was approximated using a Leap-Frog integrator [23] since it preserves volume and is invertible for any dt . We found that only two Leap-Frog steps were sufficient for all our examples. Parameters were optimised using Adam [15] (learning rate $3e-4$) and Lagrange multipliers were optimised using the same method as [12]. The initial density π was chosen separately for each dataset.

Domain knowledge via generators: In this experiment, we test the main idea of this paper where we want to learn a target density which is symmetric with respect to the $SO(2)$ group of 2D rotations. In Figure 2A we show the target density. The Lie-algebra of the $SO(2)$ group has a single generator which we express as $g(q, p) = q_1 p_2 - q_2 p_1$. For this experiment, the base density is a spherical normal $\pi(s) = N(s; 0, \mathbb{I})$ since it has $SO(2)$ symmetry. Our results demonstrate that the inclusion of the symmetry constraint increases the data-efficiency of the model in the infinite data regime and, most importantly, substantially reduces overfitting in the finite data-regime.

Multimodal density Learning: In this experiment we to assess the expressivity of Hamiltonian flows. We demonstrate in Figure 2B that it can transform a soft-uniform⁶ $\pi(s; \sigma, \beta)$ into a new density with arbitrary number of modes. Furthermore, the resulting model is nicely interpretable: The learned potential energy $U(q)$ learned to have several local minima, one for each mode of the data. As a consequence, the trajectory of the initial samples through the flow has attractors at the modes of the data.

4 Conclusions

We have demonstrated an effective method for learning invariant probability densities using Hamiltonian flows. The proposed method opens the doors to many potential applications of ML in physics, which we plan to explore in future work.

the type of generators (namely only consider those g such that $\frac{\partial^2 g}{\partial p_i \partial p_j} = 0, \forall i, j$) such that any learnt Hamiltonian will do. This condition is also satisfied in our experiments.

⁵Due to the Poisson bracket, optimisation of Hamiltonian flows involves second-order derivatives of the MLPs used for Hamiltonians and generators, so relu non-linearities are not suitable.

⁶The soft-uniform density $\pi(s; \sigma, \beta) \propto f(\beta(s + \sigma \frac{1}{2}))f(-\beta(s - \sigma \frac{1}{2}))$, where f is the sigmoid function was chosen to make it easier to visualise the learned attractors. The experiment also work if we start from other densities such as a Normal.

References

- [1] M. Albergo, G. Kanwar, and P. Shanahan. Flow-based generative models for markov chain monte carlo in lattice field theory. *arXiv preprint arXiv:1904.12072*, 2019.
- [2] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [3] T. Q. Chen, X. Li, R. Grosse, and D. Duvenaud. Isolating sources of disentanglement in variational autoencoders. *NIPS*, 2018.
- [4] T. Q. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud. Neural ordinary differential equations. In *Advances in Neural Information Processing Systems 31*, 2018.
- [5] L. Dinh, J. Sohl-Dickstein, and S. Bengio. Density estimation using real nvp. *arXiv preprint arXiv:1605.08803*, 2016.
- [6] M. C. Gemici, D. Jimenez Rezende, and S. Mohamed. Normalizing flows on riemannian manifolds. *arXiv preprint arXiv:1611.02304*, 2016.
- [7] R. Gens and P. M. Domingos. Deep symmetry networks. *NIPS*, 2014.
- [8] H. Goldstein. *Classical Mechanics*. Addison-Wesley Pub. Co, Reading, Mass., 1980.
- [9] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. β -vae: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2017.
- [10] I. Higgins, D. Amos, D. Pfau, S. Racaniere, L. Matthey, D. J. Rezende, and A. Lerchner. Towards a definition of disentangled representations. *arxiv*, 2018.
- [11] D. Jimenez Rezende and S. Mohamed. Variational inference with normalizing flows. *arXiv preprint arXiv:1505.05770*, 2015.
- [12] D. Jimenez Rezende and F. Viola. Taming vaes. *arXiv preprint arXiv:1810.00597*, 2018.
- [13] D. Jimenez Rezende, S. Mohamed, and D. Wierstra. Stochastic backpropagation and approximate inference in deep generative models. *arXiv preprint arXiv:1401.4082*, 2014.
- [14] H. Kim and A. Mnih. Disentangling by factorising. *ICLR*, 2018.
- [15] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [16] D. P. Kingma and P. Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. In *Advances in Neural Information Processing Systems*, pages 10215–10224, 2018.
- [17] D. P. Kingma and M. Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [18] D. P. Kingma, T. Salimans, R. Jozefowicz, X. Chen, I. Sutskever, and M. Welling. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pages 4743–4751, 2016.
- [19] A. A. Kirillov. *An introduction to Lie groups and Lie algebras*, volume 113. Cambridge University Press, 2008.
- [20] J. Köhler, L. Klein, and F. Noé. Equivariant flows: sampling configurations for multi-body systems with symmetric energies. 2019.
- [21] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1(4): 541–551, 1989.
- [22] D. Levy, M. D. Hoffman, and J. Sohl-Dickstein. Generalizing hamiltonian monte carlo with neural networks. *arXiv preprint arXiv:1711.09268*, 2017.

- [23] R. M. Neal et al. Mcmc using hamiltonian dynamics. *Handbook of markov chain monte carlo*, 2(11):2, 2011.
- [24] E. Noether. Invariant variation problems. *Transport Theory and Statistical Physics*, 1(3): 186–207, 1971.
- [25] G. Papamakarios, T. Pavlakou, and I. Murray. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pages 2338–2347, 2017.
- [26] T. Salimans, D. Kingma, and M. Welling. Markov chain monte carlo and variational inference: Bridging the gap. In *International Conference on Machine Learning*, pages 1218–1226, 2015.
- [27] E. G. Tabak, E. Vanden-Eijnden, et al. Density estimation by dual ascent of the log-likelihood. *Communications in Mathematical Sciences*, 8(1):217–233, 2010.