
Hamiltonian prior to Disentangle Content and Motion in Image Sequences

Asif Khan

School of Informatics
University of Edinburgh
asif.khan@ed.ac.uk

Amos Storkey

School of Informatics
University of Edinburgh
a.storkey@ed.ac.uk

Abstract

We present a deep latent variable model for high dimensional sequential data. Our model factorises the latent space into content and motion variables. To model the diverse dynamics, we split the motion space into subspaces, and introduce a unique Hamiltonian operator for each subspace. The Hamiltonian formulation provides reversible dynamics that learn to constrain the motion path to conserve invariant properties. The explicit split of the motion space decomposes the Hamiltonian into symmetry groups and gives long-term separability of the dynamics. This split also means representations can be learnt that are easy to interpret and control. We demonstrate the utility of our model for swapping the motion of two videos, generating sequences of various actions from a given image and unconditional sequence generation.

1 Introduction

The ability to learn to generate artificial image sequences has diverse uses, from animation, key frame generation, summarisation to restoration and has been explored in previous work over many decades [1, 2, 3, 4, 5]. However, learning to generate arbitrary sequences is not enough; to provide useful value, the user must be able to have control over aspects of the sequence generation, such as the motion being enacted, or the characteristics of the agent doing an action. To enable this, we must be able to learn to decompose image sequences into *content* and *motion* characteristics such that we can apply learnt motions to new objects or vary the types of motions being applied.

The dynamical processes creating the evolution of image sequences are highly constrained. Consider the simplistic case of a person walking in a scene with a camera moving around that individual. The walking pose will return to similar positions periodically, and likewise, the revolving camera will revisit previous positions. Even without strict periodicity, many dynamical processes are reversible. Any time that a dynamic could conceivably return to an earlier position suggests an implicit conservation law—the conservation of information in the underlying scene generator as it must be capable of returning to and regenerating the same scene with non-negligible probability. The critical observation we wish to capture in this paper is that understanding the conservation occurring in the context of a set of sequences is a vital ingredient to decompose content from motion.

Given any conserved quantity, any motion must be modelled in a way that maintains the conserved quantity. In physics, such a motion is called a *Hamiltonian* motion; it keeps the corresponding Hamiltonian function constant. Hence we argue that a flexible latent Hamiltonian model allows us to learn a representation that enables conservation of the right quantities (which are themselves learnt) and models the dynamic evolution. This is mathematically equivalent to learning to represent the underlying motions as combinations of differentiable symmetry groups; all differentiable symmetry transformations follow a conservation law [6].

Identifying fundamental symmetries is essential for developing expressive deep generative models (DGMs) that understand the motion constraints and can generalise beyond the training data [7, 8]. Existing DGMs [9, 10, 11] for sequences using RNNs suffer because they lack such constraint, and accumulate errors as the length of the sequence grows, ending up deviating from the relevant path [12, 13, 14, 15].

In this paper, we intimate the more general applicability of latent Hamiltonian model; previous applications have been limited to fairly constrained physical systems. We propose a VAE framework to model the dynamics of sequences using a collection of Hamiltonian operators. Specifically, for any motion sequence, we model the transition from time step t to step $t + 1$ using a group action of a Hamiltonian operator. The evolution of the dynamics of a sequence leaves certain information unchanged, identified as *content*, and certain properties that evolve in conjunction known as *motion*. Since different image sequences involve the evolution of various actions, we split the *motion space* into subspaces where each subspace is to model a unique action and is unaffected by other actions. We focus in this paper on a discrete, identified set of actions, that we can then compose at generation time. However this could be extended to learning from action compositions themselves in further work. Our main contributions are as follows,

- We propose learnable Hamiltonian operators which associate conserved quantities with latent dynamics. In contrast to existing Hamiltonian approaches, usually restricted to constrained physical systems, we extend the formulation to more natural image sequences. Furthermore, the explicit form of operator imposes the structure in the motion latent space. In this way, we simultaneously learn a structured representation along with symmetry transformations that act on the space.
- The high dimensional sequences are composed of various motions; therefore, we model each action in a separate subspace that further allows the separability of dynamics. It reduces the computational cost since the Hamiltonian of the whole space is now in a block diagonal form where each block is a Hamiltonian of a symmetry subgroup.
- We empirically demonstrate several benefits of our model, i) generation of diverse dynamics from a starting frame, ii) useful notion of disentanglement where the content representation is separated from motion representation and where each motion subspace only controls a single action, iii) we demonstrate its use for motion swapping as well as unconditional sequence generation.

2 Related Work

Hamiltonian Neural Networks In recent years, several deep learning (DL) methods have been proposed to learn the dynamics of physical systems using Hamiltonian mechanics. In [16], authors use NNs to predict Hamiltonian from phase-space coordinates $\mathbf{s} = (\mathbf{p}, \mathbf{q})$ and their derivatives. A similar work [17] used NNs to discover symmetries of Hamiltonian mechanical systems. Much recently, Hamiltonian NNs are used for simulating complex physical systems [18, 19]. The key idea of this work is to represent the states of particles as a graph and use a graph neural network (GNN) to predict the change from the current state to the next state. In a follow-up [20], authors introduce sparsity on the messages in a graph and use the symbolic regression method to search for physical laws that describe the messages in the graph. In a very recent work, Hamiltonian generative network (HGN) [8] proposed to learn Hamiltonian from image sequences. HGN maps a sequence to a latent representation and then projects it to the phase space to unroll the dynamics using an ODE integrator with Hamilton’s equation. In another work [14] use second-order ODE parameterised as a BNN for modelling dynamics of high dimensional sequence data in the latent space of VAE. Most of the developments are built on neural ODE [21] an idea to view layers of NNs as internal states of an ODE. These methods rely on the numerical integration scheme and the stability of the ODE solver. A Hamiltonian formalism is an additional requirement that the dynamics of an ODE should be volume-preserving and reversible.

Latent Space Models There is a long history of latent state space models for modelling sequences [22, 23, 24, 25, 26]. Much recently, these methods are combined with DGMs for generating high dimensional sequences as well as learning disentangled representation [12, 9, 27, 28, 29]. MoCoGAN [9] developed an adversarial framework of combining a random content noise with a sequence of random motion noise to generate videos. More recently, DSVAE [27] proposed to split latent space into time-variant and invariant representation and use LSTM [30] to model the prior on time-variant representation. S3VAE [29] improves disentanglement of DSVAE by minimising a mutual information loss between content and motion variables. Some Hamiltonian methods [8, 14]

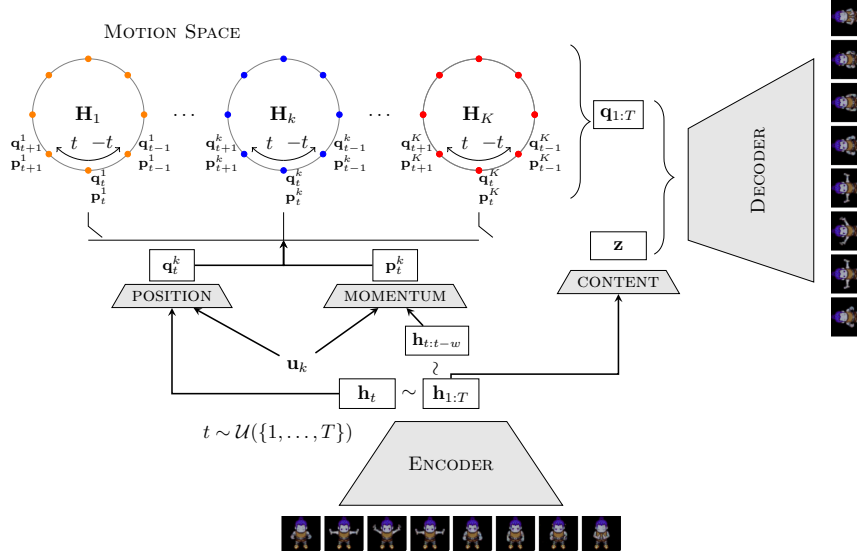


Figure 1: A framework of our model. We first map a sequence to the representation $\mathbf{h}_{1:T}$. Next, to unroll the dynamics of an action k , we map the encoded representation to the phase space. Specifically, we sample a starting index t and map \mathbf{h}_t to position coordinate \mathbf{q}_t^k . For momentum \mathbf{p}_t^k , we use temporal convolution with a kernel size w on $\mathbf{h}_{t:t-w}$. We then use the operator \mathbf{H}_k to trace out the forward and backward trajectory. We finally combine the position $\mathbf{q}_{1:T}$ and the content \mathbf{z} to generate the sequence.

also model the dynamics of high dimensional sequential data in latent space. However, the focus is only on sequence generation, and to our knowledge, it has not been used for disentanglement.

Group Transformations in Latent Space Models The early work of [31] proposed the algorithm to model the infinitesimal movement on data manifold using learnable Lie group operators. In [32], use the matrix exponents to learn the transport operators for modelling the manifold trajectory. Many other similar methods have investigated the use of geometric operators for learning the manifold representation from data [31, 32, 33, 34, 35]. The use of symmetries has recently got some attention for learning disentangled factor of variations. A disentanglement is generally identified as learning representations with independent latent factors. The main goal is that each latent factor should control a distinct data factor, and a single latent variable should control no two data factors [36, 37, 38]. In [7], proposed a symmetry-based definition of disentanglement. The goal was to decompose latent space into subspaces, and on each subspace, learn a unique group transformation such that the subspace is unchanged by the action of other groups. In [39], build such a model using interaction with the environment. Some other similar approaches recently proposed to learn group transformations in latent space [40, 41, 42]. However, the application is restricted to toy problems and to our knowledge has not been investigated on high dimensional videos.

3 Method

We propose a deep latent variable model for sets of sequential image data. Each set of sequences depicts the temporal evolution associated with one of a number of *actions*. In this context an *action* is simply a label associated with a particular sequence set, but where it is understood the sequences within a set may have very different content, but same correspondence of dynamic. E.g. in sprites data (discussed later) the actions are ‘walking’, ‘spell cast’ and ‘slash’. The sequences within a set are different individuals performing the relevant action.

Let $\mathbf{x}_{1:T}^i$ denote the i th image sequence, with \mathbf{x}_t^i the t th image in the sequence. Let \mathbf{u}^i be an indicator vector denoting the action associated with the i th sequence; i.e. $u_k^i = 1$ iff sequence i follows action k and $u_k^i = 0$ otherwise. These sequences and corresponding actions are collected into a dataset $\{(\mathbf{x}_{1:T}^i, \mathbf{u}^i)\}_{i=1}^N$ of size N , where, for the sake of simplicity, we assume they all are of same length T .

In this paper, we use a latent space to aid modelling of each sequence and decompose that latent space into two parts, which we call a *content* space (denoted by \mathbf{Z}) and a *motion* space (denoted

by \mathbf{S}). As the data comprises sequences of various actions that take different dynamical form, we further decompose the latent motion space $\mathbf{S} = \mathbf{S}^1 \oplus \mathbf{S}^2 \oplus \dots \oplus \mathbf{S}^K$, with one part for each action. In modelling a sequence corresponding to action k , only the subspace \mathbf{S}^k will be allowed to change across the length of that sequence. Each motion subspace is further decomposed into generalised *position* and *momentum* parts: $\mathbf{S}^k = (\mathbf{p}^k, \mathbf{q}^k)$. Only the position part of this latent space is used generatively to create an individual image. The momentum part *only* affects the dynamics.

Critical to this work, this decomposition makes it straightforward to define learnable Hamiltonian mechanics to model the dynamic process. This Hamiltonian model provides many advantages; it prevents the neural network from leaking constant *content* information via the motion representation, and it ensures the possibility of preserving key conservation quantities that must be implicit there are many motion dynamics. This is discussed further in Section 3.1.1. The full framework of our model is illustrated in Figure 1. In the next section, the generative model will be introduced, followed by the variational formalism for inference and the learning.

3.1 Generative Model

For completeness we first present the full probabilistic model in (3-6) before describing each component. The dynamic is defined in terms of an *action* \mathbf{u} ; the model is conditioned on this action vector. First, in (3), we sample the *content* variable \mathbf{z} from a prior $p(\mathbf{z})$. The content variable will describe the constant appearance characteristics expressed through the whole sequence. Next we sample a starting position from prior $p(\mathbf{q}_1^k)$, and a momentum from a prior $p(\mathbf{p}_1^k)$. The full state-space representation is given by $\mathbf{s}_1^k = (\mathbf{p}_1^k, \mathbf{q}_1^k)$. We then use our dynamical model in (8) to trace out the forward trajectory in the phase space. Finally, we combine the position trajectory with the content representation and use a decoder neural network to get the emission distribution of sequence in the data space. In summary,

GIVEN: k denoting action label for a sequence, (1)

$$\mathbf{z} \sim p(\mathbf{z}) \quad (2)$$

$$\mathbf{q}_1^k \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d), \quad \mathbf{p}_1^k \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d), \quad \mathbf{s}_1^k = [\mathbf{p}_1^k, \mathbf{q}_1^k] \quad (3)$$

$$\mathbf{s}_t^k = f(\mathbf{s}_{t-1}^k; \omega_k, t), \quad \forall t > 1 \quad (4)$$

$$\mathbf{s}_t^{k'} = \mathbf{0}, \quad \forall t, k' \neq k \quad (5)$$

$$\mathbf{q}_t = [\mathbf{q}_t^1, \dots, \mathbf{q}_t^K], \quad \mathbf{x}_t \sim \mathcal{N}(\mathbf{x}_t | \phi(\mathbf{z}, \mathbf{q}_t), \alpha^2 \mathbf{I}_m), \quad \forall t \quad (6)$$

where d is the dimensionality of k^{th} subspace, m is the dimensionality of data space, f is a dynamical model (8) and ω_k are its parameters corresponding to the k^{th} subspace. The emission distribution is a spherical Gaussian, with a parameterised mean $\phi(\cdot, \cdot)$, and a spherical covariance $\alpha^2 \mathbf{I}_m$. In our work, we choose $\alpha = 1$. We have provided probabilistic graph of our GM in Appendix 1.

3.1.1 Dynamical Model

In image sequences, we can view each frames of a sequence as a point in some representation space; the temporal dynamics trace a path connecting the frames forming a 1-submanifold of the image manifold. Most dynamical models either try to capture this structure deterministically [43] or probabilistically [44, 45, 27] via linear or non-linear state-space models. In either case, small errors in dynamical steps can accumulate and result in a significant deviation from the manifold when unrolling long-term trajectories at inference time [12, 13]. Interestingly, Hamiltonian systems associate conserved quantities with the motions and, we argue, alleviate some of these issues by constraining the dynamics to be reversible and volume-preserving. In our work, without significant loss of generality, we propose a linear Hamiltonian system in the latent layer, relying on the nonlinear neural network mapping to data space to handle all nonlinear aspects. This also enhances the interpretability of the dynamics.

Definition 1. A matrix $\mathbf{H} \in \mathbb{R}^{2d \times 2d}$ is an Hamiltonian matrix if $\mathbf{H}^T \mathbf{J} \mathbf{H} = \mathbf{J}$, where \mathbf{J} is a skew-symmetric matrix $\mathbf{J} = \begin{pmatrix} 0 & \mathbf{I}_d \\ -\mathbf{I}_d & 0 \end{pmatrix}$.

Consider a coordinate vector $\mathbf{s} \in \mathbb{R}^{2d}$ in the phase space \mathbf{S} that evolves under Hamiltonian energy \mathbf{E} ,

$$\mathbf{E} = \frac{1}{2} \mathbf{s}^T \mathbf{M}(t) \mathbf{s} \quad (7)$$

where $\mathbf{M}(t)$ is a symmetric matrix and \mathbf{s} is a coordinate in phase space at a time t . In Hamiltonian mechanics, the coordinates are specified in terms of position \mathbf{q} and momentum \mathbf{p} variables as $\mathbf{s} = (\mathbf{q}, \mathbf{p})$. Then the equation of motion is given by, $\frac{d\mathbf{s}(t)}{dt} = \mathbf{J}\mathbf{M}(t)\mathbf{s}$. Let $\mathbf{H}(t) = \mathbf{J}\mathbf{M}(t)$, we can rewrite the equation of motion as, $\frac{d\mathbf{s}(t)}{dt} = \mathbf{H}(t)\mathbf{s}$.

For a time-invariant Hamiltonian we can obtain the solution of the system in the closed-form using matrix exponential $\mathbf{s}(t) = e^{t\mathbf{H}}\mathbf{s}(0)$. The matrix exponent has a connection to Lie algebras, and for small t we can interpret $e^{t\mathbf{H}}$ as an infinitesimal transformation of $\mathbf{s}(0)$ under the group action. We discuss this further in Appendix A .2. For a detailed introduction to the topic we refer to [46].

In this work, we consider K Hamiltonians of the form $\mathbf{H}_1, \dots, \mathbf{H}_K$, each potentially acting on separate parts of the state space $\mathbf{S}^1, \mathbf{S}^2, \dots, \mathbf{S}^K$. To unroll the trajectory of a sequence i with associated motion k , we use the group action defined by the matrix exponent of the operator \mathbf{H}_k on a starting phase space representation $\mathbf{s}_1^k \in \mathbf{S}^k$ given by,

$$\mathbf{s}_t^k = f(\mathbf{s}_{t-1}^k; \omega_k, t) = e^{t\mathbf{H}_k} \mathbf{s}_1^k, \quad \forall t > 1 \quad (8)$$

$$\mathbf{s}_t^{k'} = \mathbf{0}, \quad \forall t, k' \neq k. \quad (9)$$

The backward dynamics can simply be obtained by negating time $-t$. We assume all time steps are equally spaced. The above formulation provides an explicit disentanglement of the motion space. It further allows us to parallelise the computation of matrix exponential by leveraging the block diagonal form of \mathbf{H} . To make use of the structure of Hamiltonians, we consider the group of real Hamiltonian matrices that form a Symplectic Lie group under multiplication $Sp(2d)$ with $2d^2 + d$ independent elements. We also look at the symplectic orthogonal group $SpO(2d)$ that further restricts Hamiltonians to skew-symmetric form with $(d^2 - d)/2$ independent elements. We briefly introduce it in Appendix A .2. For a more comprehensive overview, we refer readers to [47].

3.2 Inference

In order to learn the model parameters, we need to infer the distribution over latent variables; we follow a variational formalism that provides following evidence lower bound (ELBO),

$$\max_q \mathbb{E}_{q(\mathbf{z}, \mathbf{s}_{1:T} | \mathbf{x}_{1:T}, \mathbf{u})} \log \left[\frac{p(\mathbf{x}_{1:T}, \mathbf{z}, \mathbf{s}_{1:T} | \mathbf{u})}{q(\mathbf{z}, \mathbf{s}_{1:T} | \mathbf{x}_{1:T}, \mathbf{u})} \right] \quad (10)$$

where $q(\cdot | \cdot)$ is the approximate posterior distribution and $\mathbf{s}_t = [\mathbf{q}_t, \mathbf{p}_t]$. It remains to define the approximate posterior we use. Since the Hamiltonian dynamics are reversible in time, at inference time, we randomly sample a choice of frame t and use forward and backward dynamics to trace the trajectory of states after and before that frame. For a sequence $\mathbf{x}_{1:T}$, we use the process in (11)-(14) to draw samples from a variational distribution $q(\mathbf{z}, \mathbf{s}_{1:T} | \mathbf{x}_{1:T}, \mathbf{u})$. Simply, we sample the content variable \mathbf{z} conditioned on the observed data and independently sample the motion states $\mathbf{s}_t^k = [\mathbf{q}_t^k, \mathbf{p}_t^k]$ for the reference frame t conditioned on the observed data and the relevant action k . Motion states corresponding to other actions are set to zero. The motion states for all the other frames are then created from the forward and backward application of the Hamiltonian motion for the relevant action. In equations, this is,

$$\mathbf{z} \sim q(\mathbf{z} | \mathbf{x}_{1:T}), \quad t \sim \mathcal{U}(\{1, \dots, T\}) \quad (11)$$

$$\mathbf{q}_t^k \sim q(\mathbf{q}_t^k | \mathbf{x}_t, \mathbf{u}), \quad \mathbf{p}_t^k \sim q(\mathbf{p}_t^k | \mathbf{x}_{t-w}, \mathbf{u}), \quad \mathbf{s}_t^k = [\mathbf{q}_t^k, \mathbf{p}_t^k] \quad (12)$$

$$\mathbf{s}_{t+1}^k = f(\mathbf{s}_t^k; \omega_k, t), \quad \mathbf{s}_{t-1}^k = f(\mathbf{s}_t^k; \omega_k, -t), \quad \forall t \quad (13)$$

$$\mathbf{s}_t^{k'} = \mathbf{0}, \quad \forall t, k' \neq k, \quad (14)$$

where t is a starting index, $q(\mathbf{q}_t^k | \mathbf{x}_t, \mathbf{u})$ is the posterior distributions of k^{th} position subspace conditioned on the frame \mathbf{x}_t , $q(\mathbf{p}_t^k | \mathbf{x}_{t-w}, \mathbf{u})$ is the posterior distributions of k^{th} momentum subspace conditioned on w previous frames and $q(\mathbf{z}^i | \mathbf{x}_{1:T}^i)$ is the posterior distribution of the content space conditioned on the entire sequence. We parameterise the factorised posterior as a spherical Gaussian distribution,

$$q_{\theta_1 \circ \theta_2}(\mathbf{z} | \mathbf{x}_{1:T}) = \mathcal{N}(\mathbf{z} | \boldsymbol{\mu}_{\mathbf{z}}, \boldsymbol{\Sigma}_{\mathbf{z}}^2), \quad q_{\gamma_k \circ \theta_2}(\mathbf{q}_t^k | \mathbf{x}_t, \mathbf{u}) = \mathcal{N}(\mathbf{q}_t^k | \boldsymbol{\mu}_{\mathbf{q}_t^k}^k, \boldsymbol{\Sigma}_{\mathbf{q}_t^k}^{2k} \mathbf{I}), \quad (15)$$

$$q_{\delta_k \circ \theta_2}(\mathbf{p}_t^k | \mathbf{x}_{t-w}, \mathbf{u}) = \mathcal{N}(\mathbf{p}_t^k | \boldsymbol{\mu}_{\mathbf{p}_t^k}^k, \boldsymbol{\Sigma}_{\mathbf{p}_t^k}^{2k} \mathbf{I}) \quad (16)$$

where $\{\theta_1, \theta_2, \gamma_k, \delta_k\}$ are parameters of encoder neural network. The parameters θ_1 are specific to the content network, γ_k of the network mapping to the k^{th} position subspace, δ_k of the network mapping to the k^{th} momentum subspace and θ_2 are shared parameters of content and motion network. We use reparametrisation trick [48] to sample from latent distribution $\mathbf{z} = \boldsymbol{\mu} + \boldsymbol{\sigma} \odot \boldsymbol{\epsilon}$ where $\boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I})$.

3.3 Learning Objective

The final learning problem reduces to the optimisation of the following objective function,

$$\begin{aligned} \max_{\theta_1, \theta_2, \phi, \omega_k, \gamma_k, \delta_k} & -KL[q(\mathbf{q}_t^k | \mathbf{x}_t, \mathbf{u}) || p(\mathbf{q}_t^k)] - KL[q(\mathbf{p}_{t:t-w}^k | \mathbf{x}_t, \mathbf{u}) || p(\mathbf{p}_t^k)] \\ & - KL[q(\mathbf{z} | \mathbf{x}_{1:T}, \mathbf{u}) || p(\mathbf{z})] + \mathbb{E}_{q(\mathbf{q}_t^k | \mathbf{x}_t, \mathbf{u})} \left[\sum_{t'} \log p(\mathbf{x}_t | \mathbf{q}_{t'}, \mathbf{z}) \right]. \end{aligned} \quad (17)$$

We have provided the derivation in the Appendix A .1.

4 Experiments

We conduct experiments on the following video datasets, i) **Sprites** a sequence of animated character performing different actions as per sprites sheets.¹ It comprises three actions: ‘walking’, ‘spell cast’ and ‘slashing’ from three viewing angles: ‘left’, ‘right’ and ‘straight’. The sequences are of length 8 with each frame as an RGB image of size $64 \times 64 \times 3$. The appearance of each character has four attributes: colour of skin, hairstyle, tops and trousers/pants. Each attribute can take six values resulting in 1296 unique characters. We used 1000 characters for training and the rest for evaluation. ii) **MUG** [49] is a dataset of 52 individual performing six facial expressions: anger, disgust, fear, happiness, sadness and surprise. The dataset is made available by signing the license agreement available. The dataset consists sequences of variable length ranging from 50 to 160 frames. For training purpose, we downsample the sequences by a factor of two and use random subsequence of length 8, crop face region and resize it to 64×64 . The training and evaluation splits are based on [9].

4.1 Results and Discussion

We perform the qualitative and quantitative analysis of all the models. We investigate two choices of structure for the Hamiltonian matrices. We refer to the symplectic group structure as **H** and the symplectic orthogonal group as skew-**H**. To compute the matrix exponential, we use fast Taylor approximation [50] that provides a stable solution under various matrix norms.

Quantitative Evaluation We evaluate our model on sequence generation as well as the disentanglement of representations. As a first step we evaluate the generation of two operators **H** and skew-**H** using the per-frame structural similarity index measure (SSIM) peak signal to noise ratio (PSNR) and mean squared error (MSE). We generate a longer sequence of length 16 (twice the length used for training purpose) conditioned on the starting frame. The sprites consists of periodic sequences of length 8 where the start and end frames are identical, in this case we duplicate the sequence to get a ground truth of length 16. For MUG, we draw sequence of length 16 from the evaluation set. The SSIM scores are between -1 and 1 , with a more significant score indicating more similarity between the ground truth and generated sequence. Likewise, higher PSNR and lower MSE implies better generation. Table 1 describes the performance of our model on generation task under different scores. The scores show our model can generate high-quality sequences from an input image. For rest of the paper, we consider only **H** as it consistently works better at sequence generation. We compare its performance with the state-of-the-art baselines for sequence disentanglement, namely DSAE [27], MoCoGAN[9] and S3VAE [29].

To evaluate the disentanglement, we use a classifier pre-trained on the task of action prediction, to evaluate the generated sequences. To begin with, we draw a starting position and momentum from a prior distribution and use a dynamical model to unroll the trajectory in the phase space. Next, we sample the content variable \mathbf{z} from real sequences and combine it with position variables to generate the sequence. We report the performance of the classifier in predicting the action from these generated

¹<https://github.com/jrconway3/Universal-LPC-spritesheet>

Table 1: Evaluation for sequence generation. For SSIM and PSNR higher is better; for MSE lower is better. We can see both choices of operator can generate sequences close to the ground truth.

Model	Dataset	SSIM \uparrow	PSNR \uparrow	MSE \downarrow
H	Sprites	0.982 \pm 0.005	36.76 \pm 1.096	0.0005 \pm 0.0002
	MUG	0.797 \pm 0.003	24.49 \pm 0.099	0.0040 \pm 0.0001
Skew- H	Sprites	0.950 \pm 0.021	33.88 \pm 2.03	0.0026 \pm 0.0012
	MUG	0.791 \pm 0.003	24.25 \pm 0.094	0.0044 \pm 0.0001

Table 2: Quantitative evaluation of disentanglement and diversity of generated samples

Method	Data	Accuracy \uparrow	$H(y \mathbf{x})\downarrow$	$H(y)\uparrow$	Sprites (Attr.)	Accuracy \uparrow
Ours H	MUG	0.929	0.108	1.778	Skin Color	0.925
DSVAE [27]		0.543	0.374	1.657	Shirt	0.948
MoCoGAN [9]		0.631	0.183	1.721	Pant	0.968
S3VAE [29]		0.705	0.135	1.760	Hair	0.992
Ours H	Sprites	0.994	0.011	2.009	Identity (MUG)	0.998
DSVAE [27]		0.907	0.072	2.192		
MoCoGAN [9]		0.928	0.090	2.192		
S3VAE [29]		0.994	0.041	2.197		

(a) Results of classifier on MUG and sprites data. The high score of accuracy and Inter-Entropy $H(y)$ while low scores of Intra-Entropy $H(y|\mathbf{x})$ are expected from a better model. Our model performs best across all three scores on MUG. On sprites we are comparable to S3VAE. This is due to simplicity of classes in sprites.

(b) Performance of our model in terms of accuracy of individual attributes in sprites and identity of actors in MUG dataset. This shows our model can preserve content when the motion representation is changed.

sequences. This score gives us a measure of how well a model can keep the motion intact with the modified content variable. We also report intra-Entropy $H(y|\mathbf{x})$ and inter-Entropy $H(y)$ score using the same classifier to evaluate the diversity of generated sequences. $H(y|\mathbf{x})$ gives a measure of closeness of generated sequences to the real sequences and $H(y)$ gives a measure of diversity of generated sequences [51]. The results are reported in part (a) of Table 2. Our model outperforms the other approaches on the MUG dataset and is comparable with S3VAE on sprites. One reason for the improvement is that we explicitly associate every action with a unique subspace. This split allows our model to achieve separability of the dynamics and hence avoids the possibility of any mixing or ambiguity of action in the motion space. The results on sprites are comparable; we attribute this to the simplicity of classes in sprites that results in high performance across all models.

We further evaluate our model in preserving the identity of sequences. The identity for sprites is described in terms of four different attributes, and for MUG [49] it refers to the identity of the person. We pre-train a classifier on the task of identity prediction and use it for evaluating the generated sequences. This evaluation gives us a measure of the model’s ability to keep the identity when the motion is changed. For sprites, we report the accuracy of individual attributes. Results are outlined in Table 2. We can see on MUG our model can preserve the identity with high accuracy. We can make a similar observation for different attributes of sprites sequences. Thus, the good performance indicates that the content is preserved when traversing the motion subspace, and the motion space is invariant when changing the content variables. This is also reflected in the qualitative results.

Qualitative Evaluation For the qualitative analysis, we report results with the Hamiltonian **H**. We observed the skew-**H** resulted in similar performance; we omit it here due to limited space. To begin with, Figure 2 gives an example of original, reconstruction and generated sequences. We generate a sequence by applying the motion operator on the latent encoding of the first time step. On the left are the results for sprites and on the right of MUG video sequences. To evaluate disentanglement, we use the model for the task of motion swapping. We start by encoding two sequences $\mathbf{x}_{1:T}^1$ and $\mathbf{x}_{1:T}^2$ to their latent representations $(\mathbf{z}^1, \mathbf{q}_{1:T}^1)$ and $(\mathbf{z}^2, \mathbf{q}_{1:T}^2)$, next we swap the motion variables $(\mathbf{z}^1, \mathbf{q}_{1:T}^2)$ and $(\mathbf{z}^2, \mathbf{q}_{1:T}^1)$ between the two representation spaces, and then pass the resulting representations through the decoder to generate the sequences $\mathbf{x}_{1:T}^{1\rightarrow 2}$ and $\mathbf{x}_{1:T}^{2\rightarrow 1}$. Figure 3 shows the result of this. On the left, the pair of consecutive rows are of original sequences and on the right of the sequences generated by swapping the motion representations. We can see swapping the motion part does not affect the identity of the sequences.




Figure 2: In each patch, the first row is the original sequence, the second row is its reconstruction, and third row is a sequence generated by an action of the operator on the phase space representation of the starting time step. On left are results on sprites and on right on MUG. The reconstruction shows our model can learn good representations and the generation shows the dynamical operator can generate realistic motions from a starting frame.




Figure 3: Left: original sequence pairs; right: reconstructions after swapping the motion variables in row pairs. The content space is disentangled from the motion space.

To further investigate the generation quality of different motion operators, we use our model for an image to sequence generation. We first encode the image to its latent space representation. Next, we obtain its representation in the different motion spaces and use the respective operators to unroll the trajectories in phase space for each different motion, which are then combined with content and transformed to the image space using a decoder network. The left side of Figure 4 shows examples of decoding different motion from the same input image. For sprites, the actions are in order ‘walk’, ‘spell card’, ‘slash’ and for MUG they are ordered ‘anger’, ‘disgust’, ‘fear’, ‘happiness’, ‘sadness’ and ‘surprise’. We can see the visual dynamics associated with all the operators are well separated from one another. More results are presented in Appendix A .3.3, where it is apparent that longer-term sequences maintain the consistency associated with the content-motion pair.

5 Conclusions and Future Work

We introduced a DGM that uses Hamiltonian operators in the latent space to model the manifold of various motions. We decompose the latent space into content and motion subspace, and for each action, associate a unique partition in the motion subspace. The dynamics in each partition are unrolled using a Hamiltonian operator. Our formulation, associates conserved quantities with the




Figure 4: On left (a) and (c) are examples of image to sequence generated by an action of Hamiltonian operator on the phase space representation of the starting frame. We can observe all dynamics are well separated that demonstrates the disentanglement of various actions. On right (b) and (c) are examples of unconditional generation.

dynamics; we empirically show it provides a helpful notion of disentanglement for image sequence data. We demonstrate the performance on several tasks such as image-to-sequence, motion swapping and unconditional generation. The main advantage of our model is we can generate long term trajectories and traverse the motion manifolds of different actions in the latent space. We look forward to future applications to other sequential data types, including music audio and speech.

One potential limitation associated with the presented approach is that it is rigid in its form since we explicitly restrict every motion to a unique subspace that requires prior knowledge of the action label of a sequence. In future work we will address this by allowing each motion in the data to be composed from primitive motion components via a prior mixing distribution. Another limitation of our model is that it is less able to deal with irregularly sampled sequences, changes in tempo or reversals. In future work, this will be addressed by allowing a more flexible prior on the spacing between time steps. Nevertheless, all such approaches depend critically on the decomposition demonstrated here, and we argue that this work has demonstrated the significant promise of content and Hamiltonian motion decomposition of image sequences.

Acknowledgments and Disclosure of Funding

The authors would like to thank Joseph Mellor and William Toner for the helpful discussion during the project and Elliot J. Crowley for the valuable feedback on the paper. This research was funded in part by an unconditional gift from Huawei Noah’s Ark Lab, London.

References

- [1] David Hogg. Model-based vision: A program to see a walking person. *Image and Vision Computing*, 1(1):5–20, 1983.
- [2] Jarmo Hurri and Aapo Hyvärinen. Temporal and spatiotemporal coherence in simple-cell responses: a generative model of natural image sequences. *Network: Computation in Neural Systems*, 14(3):527–551, 2003. PMID: 12938770.
- [3] Daniel Cremers and Alan Yuille. A generative model based approach to motion segmentation. In *Joint Pattern Recognition Symposium*, pages 313–320. Springer, 2003.

- [4] AJ Storkey and CKI Williams. Image modeling with position-encoding dynamic trees. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(7):859–871, 2003.
- [5] Anitha Kannan, Nebojsa Jojic, and Brendan J Frey. Generative model for layers of appearance and deformation. In *AISTATS*, volume 2005, 2005.
- [6] Emmy Noether. Invariant variation problems, gott, 1918.
- [7] Irina Higgins, David Amos, David Pfau, Sebastien Racaniere, Loic Matthey, Danilo Rezende, and Alexander Lerchner. Towards a definition of disentangled representations. *arXiv preprint arXiv:1812.02230*, 2018.
- [8] Peter Toth, Danilo Jimenez Rezende, Andrew Jaegle, Sébastien Racanière, Aleksandar Botev, and Irina Higgins. Hamiltonian generative networks. *arXiv preprint arXiv:1909.13789*, 2019.
- [9] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. MoCoGAN: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018.
- [10] Jinsung Yoon, Daniel Jarrett, and Mihaela van der Schaar. Time-series generative adversarial networks. 2019.
- [11] Yatin Dandi, Aniket Das, Soumye Singhal, Vinay Namboodiri, and Piyush Rai. Jointly trained image and video generation using residual vectors. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3028–3042, 2020.
- [12] Maximilian Karl, Maximilian Soelch, Justin Bayer, and Patrick Van der Smagt. Deep variational Bayes filters: Unsupervised learning of state space models from raw data. *arXiv preprint arXiv:1605.06432*, 2016.
- [13] Marco Fraccaro, Simon Kamronn, Ulrich Paquet, and Ole Winther. A disentangled recognition and nonlinear dynamics model for unsupervised learning. *arXiv preprint arXiv:1710.05741*, 2017.
- [14] Cagatay Yildiz, Markus Heinonen, and Harri Lähdesmäki. ODE2VAE: Deep generative second order ODEs with Bayesian neural networks. 2019.
- [15] Alex Bird and Christopher KI Williams. Customizing sequence generation with multi-task dynamical systems. *arXiv preprint arXiv:1910.05026*, 2019.
- [16] Sam Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. *arXiv preprint arXiv:1906.01563*, 2019.
- [17] Roberto Bonadesan and Austen Lamacraft. Learning symmetries of classical integrable systems. *arXiv preprint arXiv:1906.04645*, 2019.
- [18] Alvaro Sanchez-Gonzalez, Victor Bapst, Kyle Cranmer, and Peter Battaglia. Hamiltonian graph networks with ODE integrators. *arXiv preprint arXiv:1909.12790*, 2019.
- [19] Alvaro Sanchez-Gonzalez, Jonathan Godwin, Tobias Pfaff, Rex Ying, Jure Leskovec, and Peter Battaglia. Learning to simulate complex physics with graph networks. In *International Conference on Machine Learning*, pages 8459–8468. PMLR, 2020.
- [20] Miles Cranmer, Alvaro Sanchez-Gonzalez, Peter Battaglia, Rui Xu, Kyle Cranmer, David Spergel, and Shirley Ho. Discovering symbolic models from deep learning with inductive biases. *arXiv preprint arXiv:2006.11287*, 2020.
- [21] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary differential equations. *arXiv preprint arXiv:1806.07366*, 2018.
- [22] Sam Roweis and Zoubin Ghahramani. A unifying review of linear Gaussian models. *Neural computation*, 11(2):305–345, 1999.
- [23] Robert J Elliott and Vikram Krishnamurthy. New finite-dimensional filters for parameter estimation of discrete-time linear Gaussian models. *IEEE Transactions on Automatic Control*, 44(5):938–951, 1999.
- [24] Thad Starner and Alex Pentland. Real-time American sign language recognition from video using hidden Markov models. In *Motion-based recognition*, pages 227–243. Springer, 1997.
- [25] Vladimir Pavlovic, James M Rehg, and John MacCormick. Learning switching linear models of human motion. In *NIPS*, volume 2, page 4, 2000.
- [26] Rudolph Emil Kalman. A new approach to linear filtering and prediction problems. 1960.
- [27] Li Yingzhen and Stephan Mandt. Disentangled sequential autoencoder. In *International Conference on Machine Learning*, pages 5670–5679. PMLR, 2018.
- [28] Đorđe Miladinović, Waleed Gondal, Bernhard Schölkopf, Joachim M Buhmann, and Stefan Bauer. Disentangled state space models: Unsupervised learning of dynamics across heterogeneous environments. 2019.

- [29] Yizhe Zhu, Martin Renqiang Min, Asim Kadav, and Hans Peter Graf. S3VAE: Self-supervised sequential VAE for representation disentanglement and data generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6538–6547, 2020.
- [30] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [31] Rajesh PN Rao and Daniel L Ruderman. Learning Lie groups for invariant visual perception. *Advances in neural information processing systems*, pages 810–816, 1999.
- [32] Benjamin J Culpepper and Bruno A Olshausen. Learning transport operators for image manifolds. In *NIPS*, pages 423–431, 2009.
- [33] Roland Memisevic. On multi-view feature learning. *arXiv preprint arXiv:1206.4609*, 2012.
- [34] Jascha Sohl-Dickstein, Ching Ming Wang, and Bruno A Olshausen. An unsupervised algorithm for learning Lie group transformations. *arXiv preprint arXiv:1001.1027*, 2010.
- [35] Taco Cohen and Max Welling. Learning the irreducible representations of commutative Lie groups. In *International Conference on Machine Learning*, pages 1755–1763. PMLR, 2014.
- [36] Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.
- [37] Brenden M Lake, Tomer D Ullman, Joshua B Tenenbaum, and Samuel J Gershman. Building machines that learn and think like people. *Behavioral and brain sciences*, 40, 2017.
- [38] Cian Eastwood and Christopher KI Williams. A framework for the quantitative evaluation of disentangled representations. In *International Conference on Learning Representations*, 2018.
- [39] Hugo Caselles-Dupré, Michael Garcia-Ortiz, and David Filliat. Symmetry-based disentangled representation learning requires interaction with environments. *arXiv preprint arXiv:1904.00243*, 2019.
- [40] Marissa Connor and Christopher Rozell. Representing closed transformation paths in encoded network latent space. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 3666–3675, 2020.
- [41] Robin Quessard, Thomas Barrett, and William Clements. Learning disentangled representations and group structure of dynamical environments. *Advances in Neural Information Processing Systems*, 33, 2020.
- [42] Emilien Dupont, Miguel Bautista Martin, Alex Colburn, Aditya Sankar, Josh Susskind, and Qi Shan. Equivariant neural rendering. In *International Conference on Machine Learning*, pages 2761–2770. PMLR, 2020.
- [43] Nitish Srivastava, Elman Mansimov, and Ruslan Salakhudinov. Unsupervised learning of video representations using LSTMs. In *International conference on machine learning*, pages 843–852. PMLR, 2015.
- [44] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. *arXiv preprint arXiv:1506.02216*, 2015.
- [45] Wei-Ning Hsu, Yu Zhang, and James Glass. Unsupervised learning of disentangled and interpretable representations from sequential data. *arXiv preprint arXiv:1709.07902*, 2017.
- [46] Claude Chevalley. *Theory of Lie Groups (PMS-8), Volume 8*. Princeton University Press, 2016.
- [47] Robert W Easton. Introduction to Hamiltonian dynamical systems and the N-body problem (KR Meyer and GR Hall). *SIAM Review*, 35(4):659–659, 1993.
- [48] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [49] Niki Aifanti, Christos Papachristou, and Anastasios Delopoulos. The MUG facial expression database. In *11th International Workshop on Image Analysis for Multimedia Interactive Services WIAMIS 10*, pages 1–4. IEEE, 2010.
- [50] Philipp Bader, Sergio Blanes, and Fernando Casas. Computing the matrix exponential with an optimized Taylor polynomial approximation. *Mathematics*, 7(12):1174, 2019.
- [51] Jiawei He, Andreas Lehrmann, Joseph Marino, Greg Mori, and Leonid Sigal. Probabilistic video generation using holistic attribute control. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 452–467, 2018.