

# Generative Locomotion Model of Snake Robot with Hierarchical Networks for Topological Representation

Yoonkyu Hwang and Masato Ishikawa

*Graduate School of Engineering, Osaka University, Yamadaoka, Suita, Osaka, Japan*

**Keywords:** Locomotion, Generative Model, Topological Representation, Sequential Variational Inference, Hierarchical Networks, User-oriented Interface.

**Abstract:** We propose novel generative locomotion models for snake robots. Locomotion researches have been relied on human experts with rich domain-knowledge and experience. Although recent data-driven approaches can achieve explicit controllers to make robots move, results often do not show enough interpretability with respect to user-oriented interface. The proposed model focuses on interpretable locomotion generation to help intuitive locomotion planning by end-users. First, we introduce the topological shaping for time-series training data. This allows us to bound the data to specific region, which leads to training/inference simplification, intuitive visualization, and finally high generalization property for the proposed framework. Second, the dedicated hierarchical networks were designed to propagate complex contexts in the snake robot locomotion. The result shows that our generative locomotion models can be utilized as user-oriented interface for interpretable locomotion design.

## 1 INTRODUCTION

Robot locomotion helps us to improve our intuition about how the living things act, as well as general control principle for animal-oriented robots including humanoids, dog robots, and snake robots. However, studies on robot locomotion have heavily relied on human experts with multidisciplinary knowledge/inspiration across system control theory, biology, and neuroscience. The main difficulty is to search and distinguish the structured input-output patterns. Furthermore, the multi-modality between body central pattern and joint pattern makes problem more severe.

In this paper, the locomotion problem is defined as the probabilistic generative process of an autonomous agent having latent representations for locomotion. The generative models are core machine learning framework to learn data distribution (Goodfellow et al., 2014) usually with low-dimensional latent representation (Kingma and Welling, 2013). The learned latent representation is utilized to generate the desired novel observations which were not seen before. The generative process is actively studied also in practical engineering problems. For example, the inverse design problems such as (Sanchez-Lengeling and Aspuru-Guzik, 2018) can be efficiently solved by

searching complex, multi-modal patterns realizing the desired goal.

Interpretability is an crucial factor as well as the estimation error for all machine learning techniques including generative models. The result having low interpretability is not helpful for practical generation work. Popular approach in machine-learning framework is to encode information to latent variables, and improve information connection between the desired effects and the specific latent variable. Most direct way is to reveal the related metrics such as mutual information by direct estimation (Zhang et al., 2018b; Hjelm et al., 2018; Poole et al., 2018) or some techniques like annealing (Dupont, 2018). Utilizing discrete distribution (Maddison et al., 2016) in latent variables also showed the interpretable results for the target application having the discrete characteristics. Another useful way to improve high-level interpretability is based on hierarchical architectures. The hierarchy helps information to be divided usually by overall meaning and residual information, which is called as summary vector (Veličković et al., 2018), context- nuisance vector (Tomczak and Welling, 2018), and sequence-segment level attributes (Hsu et al., 2017; Li and Mandt, 2018). Our generative locomotion model is based on the hierarchical configuration.

Modeling sequential data is a core part because a

great number of applications including robot locomotion are operated and measured in time-series. Convolutional neural networks have received huge attention for time-series application such as motion recognition (Ha and Choi, 2016) and classification (Cui et al., 2016). The convolution operation also can be utilized in causal settings (van den Oord et al., 2016). Recurrent neural networks are considered as an ideal way to estimate time-series data. However, for the sequences having too long length or complex correlation suffers from low modeling performance. Stochastic recurrent networks introduce latent variable and temporal generational process. With the tractable variational inference (Chung et al., 2015), the latent variable at each sequence-step improves modeling capacity, and can be used by high-level information indicator.

On the other hand, the time-series can be changed by other form via proper transformation such as Fourier transform, which leads to dimensional reduction of raw data. In this paper, we transform our time-series locomotion data to topological shape via Fourier transform. Then, our aim is turned to learn shape characteristics. To learn the representation of topological shape (Li et al., 2017) in the transformed locomotion data, We adopt the mentioned stochastic recurrent neural networks with variational inference. By the topological shape and stochastic recurrent neural networks, we can model complex shape pattern with meaningful representation.

Locomotion and motion generation has been intensively studied. Reinforcement Learning (Schulman et al., 2015) is widely used to make the robots to move. In (Zhang et al., 2018a), the mode-adaptive locomotion is searched. (Grochow et al., 2004) realize inverse kinematics with style change for the desired body pattern. Motion is generated for character in (Holden et al., 2016), and for heterogeneous agents in (Wampler et al., 2014). However, there remains problems about how to shape locomotion data to explore. Furthermore, consideration on user-oriented interface is limited.

This paper is organized as follows. In Section 2, the variational inference for sequential data is briefly introduced. In Section 3, the proposed generative locomotion model is derived with the problem definition for snake robot. In Section 4, we show experiments results including training fitness and locomotion generation. Finally, the conclusion is given in Section 5.

## 2 PRELIMINARIES

In this section, we briefly introduce the generative model with sequential variational inference. Our goal is to identify below data distribution  $p_\theta(x)$  with underlying valuable latent representation  $z$  and model parameter  $\theta$ .

$$p_\theta(x) = \int p_\theta(x, z) z. \quad (1)$$

However, the distribution is usually intractable because of the invisible latent variables. One of the tractable alternatives can be obtained by variational inference. It is called as the Evidence Lower Bound (ELBO)  $\mathcal{L}(\theta, \phi)$  with arbitrary alternative distribution  $q_\phi(x)$  having model parameter  $\phi$ .

$$\begin{aligned} \log p_\theta(x) &= \log \int q_\phi(z|x) \frac{p_\theta(x, z)}{q_\phi(z|x)} = \log \mathbb{E}_{q_\phi(z|x)} \\ &\geq \mathbb{E}_{q_\phi(z|x)} \log \frac{p_\theta(x, z)}{q_\phi(z|x)} = \mathcal{L}(\theta, \phi). \end{aligned} \quad (2)$$

Although the ELBO has a few different representation ways (Hoffman and Johnson, ), we focus on the below form which is decomposed as reconstruction likelihood and Kullback-Leibler divergence  $D_{KL}$ .

$$\mathbb{E}_{q_\phi(z|x)} [\log p_\theta(x|z)] - D_{KL}(q_\phi(z|x) \| p_\theta(z)). \quad (3)$$

The variational method can be also expanded for the sequential data  $x_{1:K}$  with sequence length  $K$  and corresponding sequential latent variable  $z_{1:K}$ .

$$p_\theta(x_{\leq K}, z_{\leq K}) = \prod_{k=1}^K p_\theta(x_k | z_{\leq k}, x_{<k}) p_\theta(z_k | x_{<k}, z_{<k}). \quad (4)$$

Using recurrent neural networks with hidden states  $h_k$ , the information is propagated autoregressively.

$$h_k = f(h_{k-1}, x_k, z_k). \quad (5)$$

Then, the ELBO for sequential data is defined as in (Chung et al., 2015).

$$\begin{aligned} &\mathbb{E}_{q_\phi(z_{\leq K} | x_{\leq K})} \left[ \sum_{k=1}^K \log p_\theta(x_k | z_{\leq k}, x_{<k}) - \right. \\ &\left. D_{KL}(q_\phi(z_k | x_{\leq k}, z_{<k}) \| p_\theta(z_k | x_{<k}, z_{<k})) \right] \end{aligned} \quad (6)$$

, where the prior is conditioned on the hidden state  $h_k$ . Our generative locomotion model utilizes the ELBO in the next section.

### 3 PROPOSED GENERATIVE LOCOMOTION MODEL

The locomotion of snake robot is formulated as generative models. Accordingly, the inference model and trainable ELBO is derived. We start from the definition of system states and consider the representation way of the measured data to be learned.

#### 3.1 States Definition of Target Robot

The articulated snake robot in our setting has 8-actuated joints. Two neighboring joints are perpendicularly connected as drawn in Figure 1. The robot has 8 motor angle position  $x^{(j)}(1:8)$  as joint states and 6 states of central body frame  $x^{(B)}(1:6)$ , which means position in vertical, position in forward, position in sideways, roll, pitch and yaw angle.

#### 3.2 Data Representation

How to shape data is the critical part for the efficient learning, plausible generation, and interpretability. The robot states are measured in usually time-series. However, one can utilize spectral transform if the data is expected to have spectral information. With the fact that locomotion is usually driven by spectral joint pattern, we can promote spectral characteristics of joint responses by injecting the group of excitation signal having sparse spectral components without loss of generality in both respect of time and frequency-domain. The detailed explanation about how we can achieve the excitation is given in the next section with experiments settings.

##### 3.2.1 Topological Arrangement of Joint States

After Fourier transform for the measured time-series of each joint angle state, the same length of spectral vectors, that are complex values, are achieved. Filtering out frequencies having too small power relative to input power, we obtain more compact spectral vector. Specifically, the  $L$  length spectral vector is obtained for  $T$  length time-series for each joint angle, where usually  $L \ll T$  with the filtering. The resulting spectral vectors are defined as below.

$$X_k^{(j)}(1:L) = [c_{k1}, c_{k2}, \dots, c_{kl}, \dots, c_{kL}] \quad (7)$$

, where  $c_{kl} = [r_{kl}, i_{kl}]$ .

The terms of  $X_k^{(j)}(1)$ ,  $c_{kl}$ ,  $r_{kl}$  and  $i_{kl}$  means the spectral vector for  $k$  th joint angle, the spectral components for  $l$  th frequency, and values corresponding

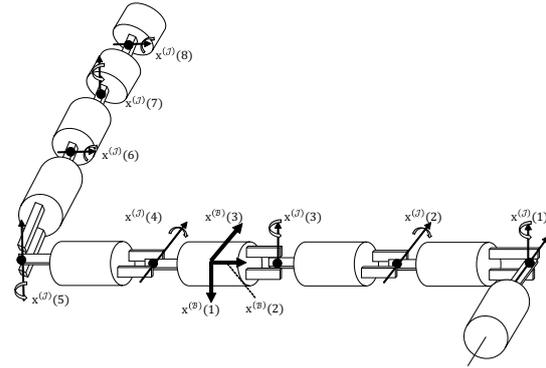


Figure 1: The schematic of 3d snake robot system.

real and imaginary axis, respectively. Then, the set of  $X_k^{(j)}(1)$  is re-arranged as topological shape. The spectral component  $c_{kl}$  is mapped as the point in coordinates  $(r_{kl}, i_{kl}, k)$ , which is illustrated in Figure 2. This approach turns our locomotion problem to one to extract and generate meaningful pattern in point cloud group that is topologically connected. Finally, we can view locomotion data as the pair of the topological joint shape  $X_{1:8}^{(j)}(1:L)$  and jointly distributed body pattern in time series  $x^{(B)}$ .

#### 3.3 Overall Architecture

The overall architecture of the proposed generative locomotion model is illustrated in Figure 3. It is composed of three parts (at top right box in Figure 3) including the encoder of body states, the decoder of body states, and sequential variational autoencoder regarding topological joint shape. Again, the last part is composed by hierarchically connected two stochastic recurrent neural networks, which is written as ① in Figure 3. We call central recurrent network that is illustrated by large squares as root network and small squares as child network in ①. The  $X_k^{(j)}(1:L)$  is

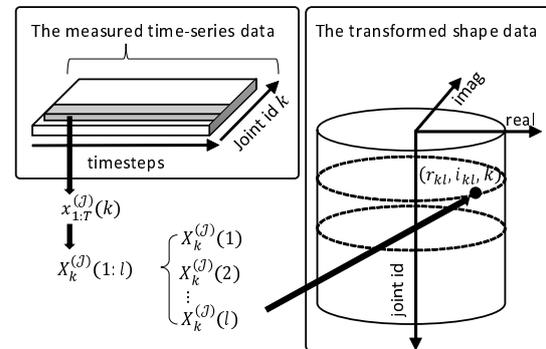


Figure 2: Topological arrangement for time-series of joint states.

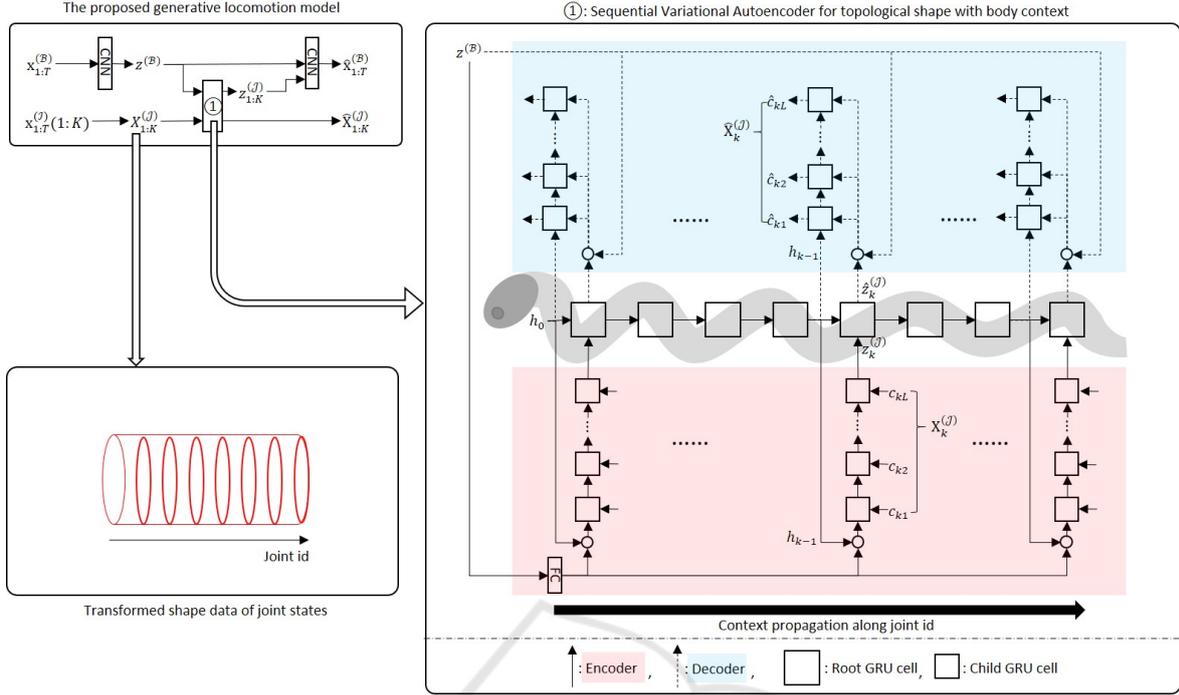


Figure 3: The overall architecture of the proposed generative locomotion model.

the topological joint shape for  $x_{1:T}^{(j)}(k)$ . The  $z^{(B)}$  and  $z_k^{(j)}$  is the encoded latent variables for body states and topological joint shape, respectively. Then, the hidden states  $h_k$  of root network at  $k$  th step is updated as following recurrence function.

$$h_k = f(h_{k-1}, X_k^{(j)}, z_k^{(j)}, z^{(B)}). \quad (8)$$

The child networks are propagated with three context vector such as hidden states of root networks, body latent variable, and joint shape latent variables. For the encoder and decoder of body states, the convolutional neural networks are used. However, because the topological joint shape has small number of spectral components and joint, the conventional powerful convolutional neural networks-based approach is difficult to be applied for small number of pixels. Also, deeply stacked recurrent networks without the stochastic generation at each sequence step is hard to train practically regardless of hierarchical configuration. On the other hand, the proposed architecture propagates latent context information autoregressively and hierarchically.

### 3.4 Learning Objective

Our goal is to identify the locomotion data distribution with below decomposition.

$$p_{\theta} \left( X_{1:K}^{(j)}, x_{1:T}^{(B)}, z_{1:K}^{(j)}, z^{(B)} \right). \quad (9)$$

The generative model can be derived for the proposed structure in Figure 3 with similar way in the explanation of section 2.

$$p_{\theta} \left( x_{1:T}^{(B)} | z_{1:K}^{(j)}, z^{(B)} \right) p_{\theta} \left( z^{(B)} \right) \times \prod_{k=1}^K p_{\theta} \left( X_k^{(j)} | z_{\leq k}^{(j)}, z^{(B)} \right) p_{\theta} \left( z_k^{(j)} | z_{< k}^{(j)}, z^{(B)} \right). \quad (10)$$

By the proposed hierarchical architecture, the prior distribution for body latent variable and joint latent variables is defined as below.

$$z^{(B)} \sim \mathcal{N}(0, I). \quad (11)$$

$$z_k^{(j)} \sim \mathcal{N} \left( \mu_k^{(j_{prior})}, \sigma_k^{(j_{prior})} \right) \\ \left[ \mu_k^{(j_{prior})}, \sigma_k^{(j_{prior})} \right] = f^{(prior)} \left( h_{k-1}, z^{(B)} \right). \quad (12)$$

Also, for the proposed network, the inference model is written by

$$q_{\phi} \left( z^{(B)} | x_{\leq T}^{(B)} \right) \prod_{k=1}^K q_{\phi} \left( z_k^{(j)} | X_{\leq K}^{(B)}, z_{< k}^{(j)}, z^{(B)} \right). \quad (13)$$

Before the derivation of final training objectives, it should be noted that the posterior collapse often occurs in the variational representation learning frameworks (Dai and Wipf, 2019; Alemi et al., 2017). The problem states the posterior distribution  $q_\theta(z|x)$  tends to  $q_\theta(z)$ , which leads to uninformative representation learning. A few solutions are proposed such as mutual information regularization, KL divergence annealing, and the utilization of Maximum Mean Discrepancy (MMD). The MMD can be alternative divergence measure to KL divergence in original formulation of variational inference as in (Tolstikhin et al., 2017; Zhao et al., 2019). The MMD  $D_{MM}$  is defined with Gaussian kernel.

$$D_{MM}(p(z) \| q(z)) = \mathbb{E}_{p(z), p(z')} [\kappa(z, z')] + \mathbb{E}_{q(z), q(z')} [\kappa(z, z')] - 2\mathbb{E}_{p(z), q(z')} [\kappa(z, z')]. \quad (14)$$

, where Gaussian kernel  $\kappa(z, z') = e^{-\frac{\|z-z'\|^2}{2\sigma^2}}$ . Finally, the learning objective for our generative locomotion model is derived by

$$\mathbb{E}_{q_\phi(\mathbf{z}|\mathbf{x})} [\mathcal{L}_{REC}^J + \alpha \mathcal{L}_{REC}^B - \beta \mathcal{L}_{MMD}^J] - \gamma \mathcal{L}_{MMD}^B, \quad (15)$$

where:

$$q_\phi(\mathbf{z}|\mathbf{x}) = q_\phi\left(z^{(B)}, z_{<k}^{(j)} \mid x_{\leq T}^{(B)}, X_{\leq k}^{(j)}\right),$$

$$\mathcal{L}_{REC}^J = \sum_{k=1}^K \log p_\theta\left(X_k^{(j)} \mid z_{\leq k}^{(j)}, X_{<k}^{(j)}, z^{(B)}\right),$$

$$\mathcal{L}_{REC}^B = \log p_\theta\left(x_{\leq T}^{(B)} \mid z_{\leq K}^{(j)}, z^{(B)}\right),$$

$$\mathcal{L}_{MMD}^J = D_{MM}\left(q_\phi\left(z_k^{(j)} \mid X_{\leq k}^{(j)}, z_{<k}^{(j)}, z^{(B)}\right) \parallel p_\theta\left(z_k \mid X_{<k}^{(j)}, z_{<k}^{(j)}, z^{(B)}\right)\right),$$

$$\mathcal{L}_{MMD}^B = D_{MM}\left(q_\phi\left(z^{(B)} \mid x_{\leq T}^{(B)}\right) \parallel p_\theta\left(z^{(B)}\right)\right).$$

The terms of  $\mathcal{L}_{REC}^J$ ,  $\mathcal{L}_{REC}^B$ ,  $\mathcal{L}_{MMD}^J$ , and  $\mathcal{L}_{MMD}^B$  respectively represent reconstruction loss of joint states, reconstruction loss of body states, regularization term on joint states distribution, and regularization term on body states distribution. At the training phase, the total loss is weighted by hyperparameters of  $\alpha$ ,  $\beta$ , and  $\gamma$ .

## 4 EXPERIMENTS

The results are shown after describing the experiments settings, the detailed network specification, and the training settings.

### 4.1 Acquisition of Training Samples

The experiments were performed on virtual robot simulator V-REP. For the acquisition of training samples, the Riemann-equivalent excitation signal (Schoukens et al., 2016) is injected as motor reference angles for each joint. The excitation signal is known for good alternative to Gaussian random noise excitation, chirp, or the sum of multi-sines, because the excitation shows the controlled amplitude distribution both in time-domain and frequency-domain with small number of spectral components (Pintelon and Schoukens, 2012). Specifically, below signal is injected.

$$u_k(t) = \sum_{l=1}^L U \sin(2\pi f_l t + \phi_l) \quad (16)$$

, where  $u_k(t)$ ,  $U$ ,  $f_l$ ,  $L$ , and  $\phi_l$  means reference input of  $k$ th joint, amplitude constant, frequency, and phase that is distributed randomly in  $[0, 2\pi]$ , respectively. In our experiments, the number of frequencies ( $L$ ) is 8, the frequencies  $f_l$  are selected in logarithmically equal interval from 0.05 Hz to 3.0 Hz. The excitation signal  $u_{k=1,2,\dots,8}(t) \in [-2.2, 2.2]$  rad is obtained in 20 Hz sampling rate with 2000 timesteps and 4096 batch samples. After downsampling with 1/4 factor, the topological joint shape is obtained by following section 3. The topological shape and corresponding body states are illustrated in Figure 4.

### 4.2 Network Specification

The 3 layers of convolutional neural networks are stacked for the encoder and decoder of body states. The body encoder has separable convolution neural network layer for its first layer and the filter size is increased two times with 1/2 average pooling layer for each stacking. Global average pooling is done at final layer of body encoder. Inversely, the decoder of body states is upsampled two times with the decreased filter size for each stacking. For the hierarchically connected two stochastic recurrent neural networks, Gated Recurrent Unit is used for parameter efficiency. All child network shares parameters between each other, and the root network has independent parameters. The body and joint latent distribution is isotropic Gaussian. For the body states encoder/decoder, the detailed architecture is illustrated in Appendix.

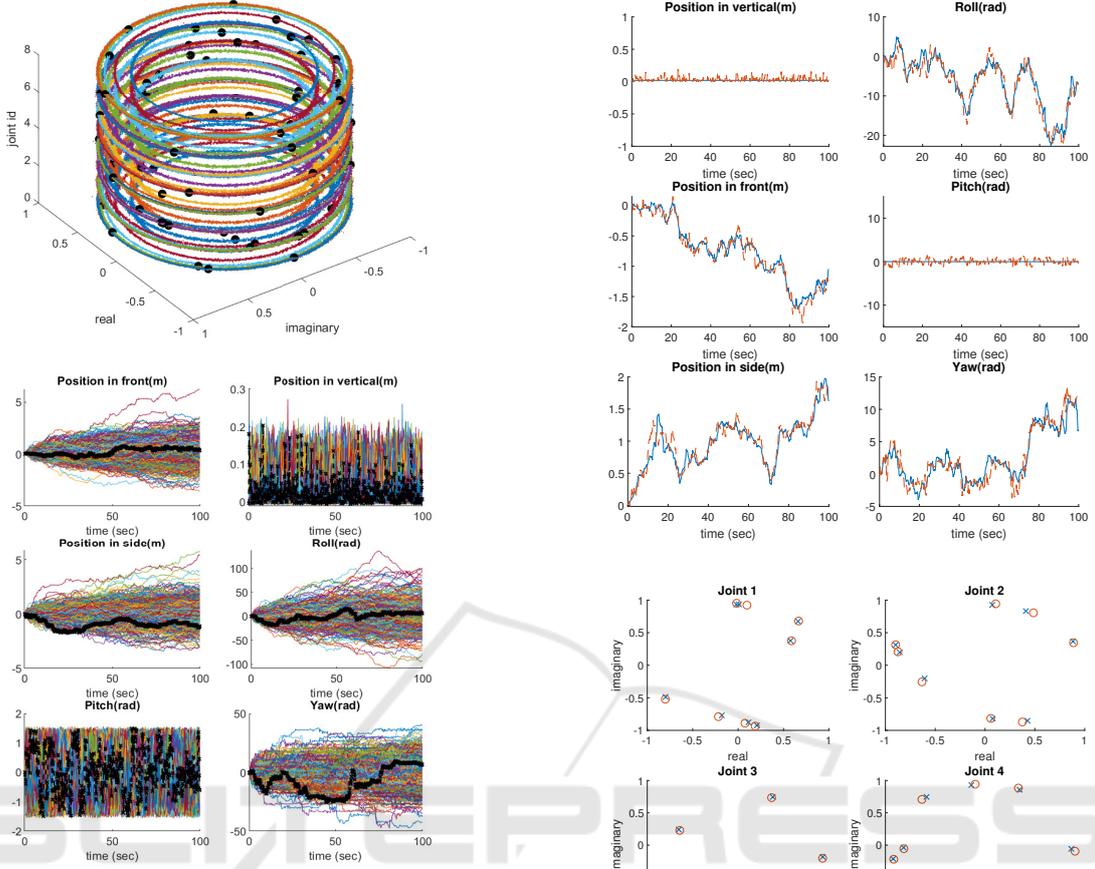


Figure 4: The topological joint shape of the measured training samples. In the top figure, the point group of black circles corresponds one example sample in all batch experiments. In the bottom figure, the highlighted black line is the paired body states corresponding to the black circles in top figure.

### 4.3 Results

#### 4.3.1 Training Results

The training was done by the Adam optimizer with cyclic learning rate for faster convergence (Smith, 2017). The learning rate is cyclically repeated from 0.0002 to 0.004 in 800 epoch period with exponentially decaying profile. The training results in Figure 5 show good fitting both in body states and topological joint shape.

#### 4.3.2 Locomotion Generation

We perform generation process for two different settings. First, the body latent variable  $z^{(B)}$  is randomly sampled for the prior distribution with zero the joint latent variables  $z_{1:K}^{(J)}$ . The joint latent variable  $z_{1:K}^{(J)}$  is principally sampled by dynamic condition. However,

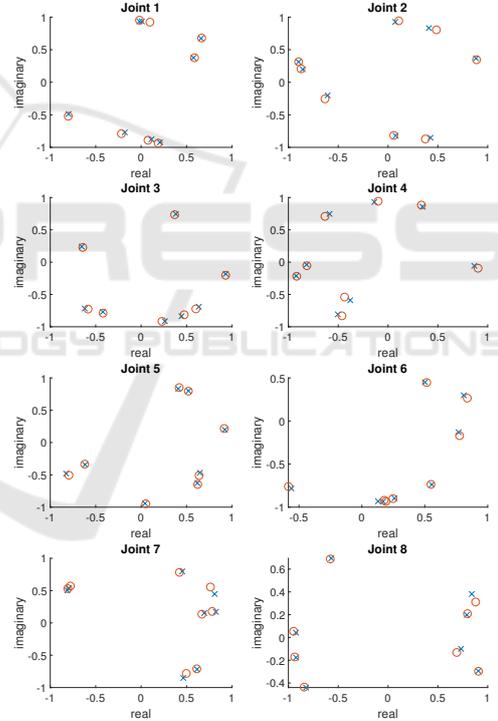


Figure 5: The training result. For the topological joint shape (bottom figure), the circle means experiments data, and the cross does the estimated one. For the body states (top figure), the line means experiments data, and the dashed line means the estimated one.

we set them as zero to identify the disentanglement ability of body latent variables. Second, the joint latent variables  $z_{1:K}^{(J)}$  is sampled by the prior distribution with the body latent variable  $z^{(B)}$  fixed. The generated samples are shown in Figure 6 for the first setting, and in Figure 7 for the second setting. In the Figure 6, The

generation result of first settings shows that the body latent variable  $z^{(B)}$  efficiently spans across sequence level that is represented as variation in vertical direction. It means that we can design locomotion based on the intuitive body latent variables. The result of second setting in Figure 7 indicates that two body pattern have inter-sequence variation, not the intra-sequence variation. The same result also can be found samples of spectral joint shape in Figure 8. Thus, we can generate locomotion stably for the desired intuitive body pattern with corresponding topological joint shape. The results from two different settings, the proposed generative locomotion model has user-oriented interface to help locomotion design.

### 5 CONCLUSIONS

We introduced the generative locomotion model. Compared with the conventional works that were relied on human expert or specific domain-knowledge, the proposed approach provides user-oriented high level intuition/interface. In our model, one can generate complex locomotion by selecting intuitive body patterns. The advantages come from two methods. Topological shaping allows locomotion data to be bounded in shape, which leads to generalize datasets. The designed hierarchical networks efficiently estimated complex contexts via latent information propagation. We expect that our generative model approach on locomotion proposes alternative ways for conventional locomotion studies, and the method can be adopted for other robot locomotion researches.

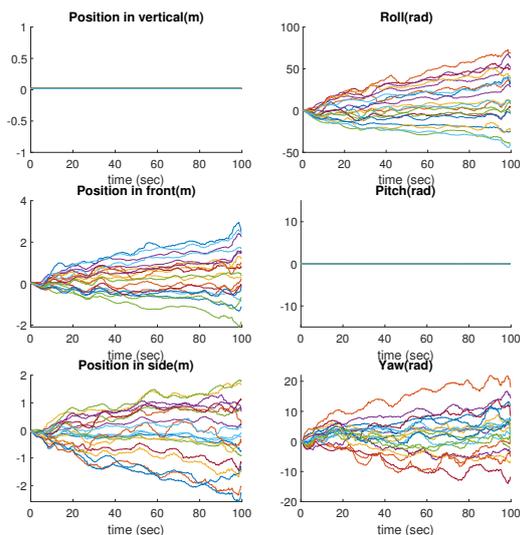


Figure 6: The body latent variable is randomly sampled with zero joint latent variables.

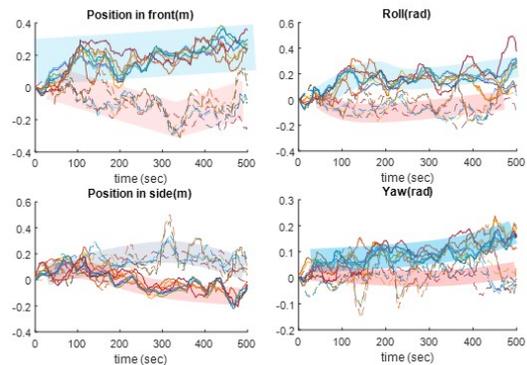


Figure 7: Two differnt latent variables of body state were sampled. For each body latent variable, 8 different joint latent variables were sampled. The shaded region represents trajectories for same body latent variable. In the shaded region, 8 different joint latent sample trajectories were drawn.

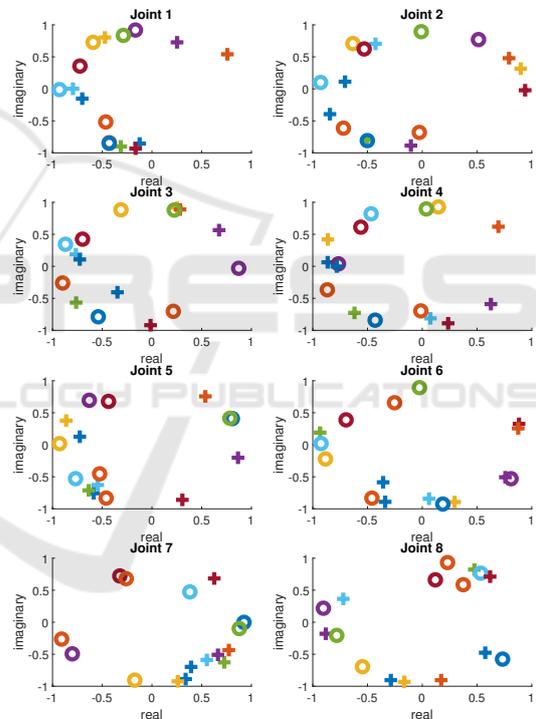


Figure 8: For two different body latent variables in Figure 7, independet 8 joint samples were drawn. The symbols 'o' and 'x' respectively represents two differnt body latent variables. For the symbol 'o' or 'x', independent joint latent samples were illustrated by differnt colors.

### ACKNOWLEDGEMENTS

This work was partially supported by CREST, JST, and Komatsu MIRAI construction institute, Osaka University.

## REFERENCES

- Alemi, A. A., Poole, B., Fischer, I., Dillon, J. V., Saurous, R. A., and Murphy, K. (2017). Fixing a broken ELBO. *arXiv preprint arXiv:1711.00464*.
- Chung, J., Kastner, K., Dinh, L., Goel, K., Courville, A., and Bengio, Y. (2015). A Recurrent Latent Variable Model for Sequential Data. pages 1–9.
- Cui, Z., Chen, W., and Chen, Y. (2016). Multi-Scale Convolutional Neural Networks for Time Series Classification.
- Dai, B. and Wipf, D. (2019). Diagnosing and enhancing vae models. *arXiv preprint arXiv:1903.05789*.
- Dupont, E. (2018). Learning disentangled joint continuous and discrete representations. In *Advances in Neural Information Processing Systems*, pages 710–720.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680.
- Grochow, K., Martin, S. L., Hertzmann, A., and Popović, Z. (2004). Style-based inverse kinematics. In *ACM transactions on graphics (TOG)*, volume 23, pages 522–531. ACM.
- Ha, S. and Choi, S. (2016). Convolutional neural networks for human activity recognition using multiple accelerometer and gyroscope sensors. In *2016 International Joint Conference on Neural Networks (IJCNN)*, pages 381–388. IEEE.
- Hjelm, R. D., Fedorov, A., Lavoie-Marchildon, S., Grewal, K., Bachman, P., Trischler, A., and Bengio, Y. (2018). Learning deep representations by mutual information estimation and maximization. *arXiv preprint arXiv:1808.06670*.
- Hoffman, M. D. and Johnson, M. J. ELBO Surgery: Yet another way to carve up the evidence lower bound.
- Holden, D., Saito, J., and Komura, T. (2016). A deep learning framework for character motion synthesis and editing. *ACM Transactions on Graphics (TOG)*, 35(4):138.
- Hsu, W.-N., Zhang, Y., and Glass, J. (2017). Unsupervised learning of disentangled and interpretable representations from sequential data. In *Advances in neural information processing systems*, pages 1878–1889.
- Kingma, D. P. and Welling, M. (2013). Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Li, J., Xu, K., Chaudhuri, S., Yumer, E., Zhang, H., and Guibas, L. (2017). Grass: Generative recursive autoencoders for shape structures. *ACM Transactions on Graphics (TOG)*, 36(4):52.
- Li, Y. and Mandt, S. (2018). Disentangled sequential autoencoder. *arXiv preprint arXiv:1803.02991*.
- Maddison, C. J., Mnih, A., and Teh, Y. W. (2016). The concrete distribution: A continuous relaxation of discrete random variables. *arXiv preprint arXiv:1611.00712*.
- Pintelon, R. and Schoukens, J. (2012). *System identification: a frequency domain approach*. John Wiley & Sons.
- Poole, B., Ozair, S., van den Oord, A., Alemi, A. A., and Tucker, G. (2018). On variational lower bounds of mutual information. In *NeurIPS Workshop on Bayesian Deep Learning*.
- Sanchez-Lengeling, B. and Aspuru-Guzik, A. (2018). Inverse molecular design using machine learning: Generative models for matter engineering. *Science*, 361(6400):360–365.
- Schoukens, J., Vaes, M., and Pintelon, R. (2016). Linear system identification in a nonlinear setting: Nonparametric analysis of the nonlinear distortions and their impact on the best linear approximation. *IEEE Control Systems Magazine*, 36(3):38–69.
- Schulman, J., Moritz, P., Levine, S., Jordan, M., and Abbeel, P. (2015). High-Dimensional Continuous Control Using Generalized Advantage Estimation. pages 1–14.
- Smith, L. N. (2017). Cyclical learning rates for training neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 464–472. IEEE.
- Tolstikhin, I., Bousquet, O., Gelly, S., and Schoelkopf, B. (2017). Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558*.
- Tomczak, J. M. and Welling, M. (2018). VAE with a vamp-prior. *International Conference on Artificial Intelligence and Statistics, AISTATS 2018*, pages 1214–1223.
- van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., and Kavukcuoglu, K. (2016). Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*.
- Veličković, P., Fedus, W., Hamilton, W. L., Liò, P., Bengio, Y., and Hjelm, R. D. (2018). Deep graph infomax. *arXiv preprint arXiv:1809.10341*.
- Wampler, K., Popović, Z., and Popović, J. (2014). Generalizing locomotion style to new animals with inverse optimal regression. *ACM Transactions on Graphics (TOG)*, 33(4):49.
- Zhang, H., Starke, S., Komura, T., and Saito, J. (2018a). Mode-adaptive neural networks for quadruped motion control. *ACM Transactions on Graphics (TOG)*, 37(4):145.
- Zhang, Y., Galley, M., Gao, J., Gan, Z., Li, X., Brockett, C., and Dolan, B. (2018b). Generating informative and diverse conversational responses via adversarial information maximization. In *Advances in Neural Information Processing Systems*, pages 1810–1820.
- Zhao, S., Song, J., and Ermon, S. (2019). Infovae: Balancing learning and inference in variational autoencoders. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 5885–5892.

