



PirouNet: Creating Dance Through Artist-Centric Deep Learning

Mathilde Papillon¹ , Mariel Pettee² , and Nina Miolane³ 

¹ Department of Physics, University of California Santa Barbara,
Santa Barbara, CA, USA
papillon@ucsb.edu

² Lawrence Berkeley National Lab, Berkeley, CA, USA

³ Department of Electrical and Computer Engineering,
University of California Santa Barbara, Santa Barbara, CA, USA

Abstract. Using Artificial Intelligence (AI) to create dance choreography is still at an early stage. Methods that conditionally generate dance sequences remain limited in their ability to follow choreographer-specific creative direction, often relying on external prompts or supervised learning. In the same vein, fully annotated dance datasets are rare and labor intensive. To fill this gap and help leverage deep learning as a meaningful tool for choreographers, we propose “PirouNet”, a semi-supervised conditional recurrent variational autoencoder together with a dance labeling web application. PirouNet allows dance professionals to annotate data with their own subjective creative labels and subsequently generate new bouts of choreography based on their aesthetic criteria. Thanks to the proposed semi-supervised approach, PirouNet only requires a small portion of the dataset to be labeled, typically on the order of 1%. We demonstrate PirouNet’s capabilities as it generates original choreography based on the “Laban Time Effort”, an established dance notion describing a given intention for a movement’s time dynamics. We extensively evaluate PirouNet’s dance creations through a series of qualitative and quantitative metrics, validating its applicability as a tool for choreographers.

Keywords: Deep learning · Neural network · Semi-supervised · Generative · Recurrent · LSTM · VAE · Pose · Laban

1 Introduction

Recent years have witnessed the development of deep learning methods that can generate original dance sequences. Yet, these methods have not been widely adopted by communities of dancers and choreographers. A possible reason for this lack of adoption is the fact that existing AI-powered dance generation tools cannot create dance based on artistic criteria, or “user-inputs” that are meaningful to choreographers, such as elements of Laban Movement Analysis [8, 17].

Laban Movement Analysis (LMA) [8,17] stands as one of the most established Western methods for describing and understanding human motion [33], often used by dance professionals to transmit choreography, teach technique [13,56], and rigorously assess performance [2]. LMA’s corresponding notation, Labanotation (LN), enables detailed and objective recording of the quantitative and qualitative characteristics of movement [22]. As a result, LMA has become a cross-disciplinary staple of technologically-driven movement analysis and has been extensively assessed for reliability and effectiveness [5,26,35]. LMA now provides a valuable tool for many movement-centric fields beyond dance, such as physical rehabilitation [8], theater [58], sports [19], and psychology [40,44,51].

Laban Axes and Laban Efforts. LMA is divided into four axes which independently contribute to its portrayal of movement [12]. “Body” outlines actions of body parts and relationships between those parts, “Shape” addresses the body’s more general evolution of shape, “Space” describes the geometrical and directional emphasis of movement, and “Effort” describes the energy content and inner intention of a movement. In the context of this paper, we use the Effort axis as a leading example of aesthetic criterion, depicted in Fig. 1, for AI-powered dance creation. We describe choreographic intention as the intended Effort of a movement. Effort is a multi-dimensional axis of LMA, described by four independent Laban Efforts. It truly sets LMA apart, as it describes an intention of movement, rather than the resulting execution [15,29].




Fig. 1. LMA’s Effort axis consists of a four-dimensional space spanned by four orthogonal efforts called Laban Efforts. For every movement, there exists a 4-vector of total effort. For example, a movement that is described as bound, light, sustained, and direct has a total effort comprised of Low Flow Effort, Low Weight Effort, Low Time Effort, and High Space Effort.

PirouNet: Creating Dance from User-Defined Labels. Given the lack of user-tailored creative dance AIs, we propose “PirouNet”, a semi-supervised generative recurrent deep learning model that creates new dance sequences from choreographers’ aesthetic inputs. We illustrate the use of PirouNet with categorical intensities of the Laban Time Effort, a staple of LMA’s descriptive richness. Future users can choose different Laban-inspired labels, or define any subjective

label they wish to use for creation, regardless of their knowledge of LMA. The artist then categorically labels a very small portion of an input dance dataset, typically of the order of 1%. We provide a label augmentation tool that smartly increases this portion to about 25%, in our case. Once trained, PirouNet creates new dances prompted by the artist’s desired style.

Contributions. Our contributions are three-fold.

1. We introduce the first deep generative model for dance based on LMA’s Effort axis.
2. We provide a novel dance labeling web application with a label augmentation toolkit, which distinguishes itself from currently available customizable annotation databases [14] as it may be used on the user’s own motion dataset. Figure 2 shows an annotated screenshot of the app.
3. We present a semi-supervised approach that, coupled with our web app, (i) allows the use of our model for dance generation based on *any* class of the choreographer’s choosing and (ii) limits manual annotation labor to 1% of the dance dataset.

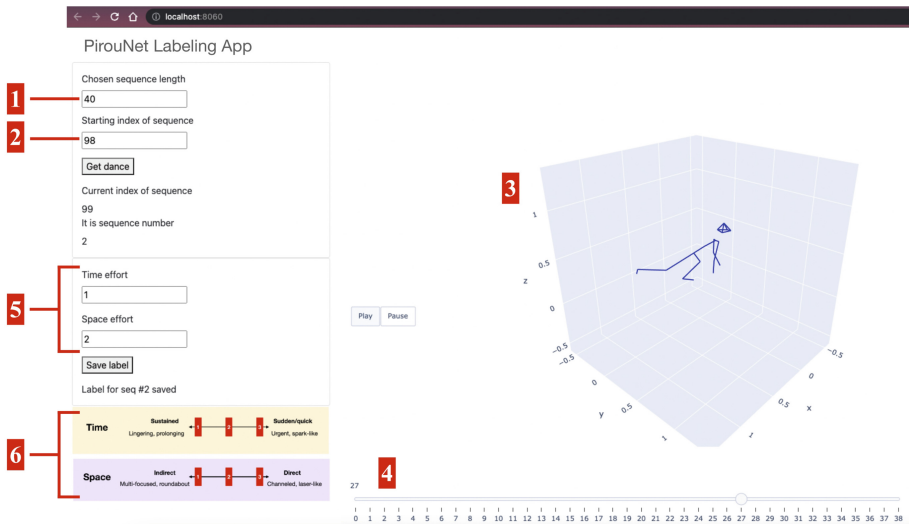


Fig. 2. Screen capture of web labeling app. User enters amount of poses per dance sequence in (1), and the index of the first sequence’s starting pose in (2). Upon clicking “Get Dance”, an animation of the fully connected skeleton appears in (3). The user can zoom and rotate the animation directly, and click into specific frames with (4). User enters Laban Effort (or any chosen label) in (5) and clicks “Save label” to record the inputs to a CSV file. A spot is reserved in (6) for an infographic with instructions to ensure consistent labeling.

Outline. The paper is organized as follows. Section 2 describes related work coupling LMA, deep learning and automatic generation of human motion and dance. Section 3 introduces the architecture and semi-supervised training of the proposed model, PirouNet. Section 4 presents experimental results on dance creation from Laban Efforts, validated with a series of qualitative and quantitative metrics that demonstrate PirouNet’s potential for AI-powered and choreographer-controlled dance generation.

2 Related Works

We review related works on LMA and machine learning, as well as on the automatic creation of human motion and dance.

Predicting LMA from Videos or Pose Sequences. A first class of learning methods focuses exclusively on predicting LMA from input dance videos. Researchers have successfully leveraged motion sensors [48] and pose detection systems [6, 23] to automatically infer LMA labels from dance sequences, which subsequently allows assessing and improving movement techniques [47]. Recent advances in the field have extended the ability of these systems to generate the analyzed movement’s corresponding LN using feed-forward time-series neural networks [53] and bi-directional RNNs [61].

Predicting High-Level Dance Information from LMA. Another class of methods uses existing LMA annotations (e.g. produced from the first class of methods) as inputs to perform downstream tasks, demonstrating the richness of LMA notations. For example, [2] can predict the similarity of two dancing motions from LMA. Supervised deep learning using LMA as inputs has also successfully extended computational movement analysis to emotion recognition [59], often relying on the complex temporal learning processes offered by Long Short Term Memory networks (LSTMs) [20, 54] and Convolutional Neural Networks (CNNs) [50].

Creating Dance from LMA: Non-deep AI and Algorithmic Approaches. LMA has played a key role in the emergence of a myriad of tools for generating dance directly from written scores [57] or vice versa [11]. This approach’s machine learning component mostly relies on associative memory [10] or look-up tables [60], with an associated retrieval criterion. The subsequent adaptation of the dance motion is either implemented through hard-coded dynamic constraints or through a combination of such constraints and direct user-inputs. While such methods testify of the interest in generating dance from LMA, they are rather restrictive in the creative context of dance generation. Carlson et al. [9] make use of a genetic algorithm to generate static dance poses from Laban Efforts, meant to catalyze choreographic creation.

Creating Dance from Non-LMA Inputs: Deep Learning Approaches. Outside of the LMA community, researchers have introduced deep learning tools into the choreographic creative process. Most of these tools generate sequences of movement using variants of Recurrent Neural Networks (RNNs) [43] to represent the time-component of the dance dynamical system, and (variational) autoencoders [55] to transform abstract latent variables into dance movements. A first class of algorithms generates choreography from scratch: the generated dance is completely new movement [3,62], without using any input from the choreographer. Other instances of innovative techniques achieving this include self-organizing maps [34] and autoencoders for reacting to live movement [4]. Recent research has also introduced variational autoencoders [42] that encode sequences towards and generate sequences from a lower dimensional latent space. This enables generation of variations on given sequences as well as sampling new sequences. Lastly, other methods prompt sequences with music [1,30–32], which a choreographer can use to influence dance outputs. In all of these cases, the choreographer’s creative control over the movement output is either limited to the choice of training data or that of an external prompt. To the best of our knowledge, there exists no model able to create dance sequences from choreographer-specific aesthetic labels or Laban Efforts.

Creating Movement from Non-LMA Inputs: Conditional Deep Learning Approaches. Recent methods, called label-conditioned movement generations, create human motion based on a discrete set of labels. To our knowledge, only two label-conditioned motion generation systems exist in the literature: Action2Motion [18] and Actor [41]. Neither are adaptable to LMA-based motion generation nor to choreographer-guided dance generation, as they require a large, fully labeled dataset. Both Action2Motion and Actor train their methods from motion capture (MoCap) datasets, including the NTU dataset (containing over 100k movements) [45] or the UESTC dataset (containing over 25k movements) [24] labeled with human actions such as walking or throwing. Such labeled MoCap datasets are imbalanced in categories or noisy (inspiring [18] to produce their own action dataset) and do not consider dance moves nor LMA labels. Certain annotated dance databases exist [14,25,52], although they are limited to their producers’ specific genres, styles, and creative processes. We postulate that the lack of creative, yet precise, dance generation algorithms come from a lack of large dance datasets with labels meaningful to choreography as an artistic practice. Coming up with such a dataset for fully supervised models is especially difficult due to the time-consuming nature of such annotations.

3 Methods

This section presents the proposed conditional dance generative model and the core components of PirouNet’s architecture: a dance encoding and generative model that uses (i) a variational autoencoder (VAE) [55] inspired from [42], coupled with (ii) motion dynamics through a LSTM [20] network. In order to

minimize the need for manual labeling, PirouNet also leverages (iii) a semi-supervised learning approach [28]. PirouNet’s source code implementing this model is available at github.com/bioshape-lab/pirounet. As PirouNet accepts both labeled and unlabeled input data for training, its VAE adopts two different forms, showcased in Fig. 3. The VAE makes use of a linear classifier for the case of unlabeled data.

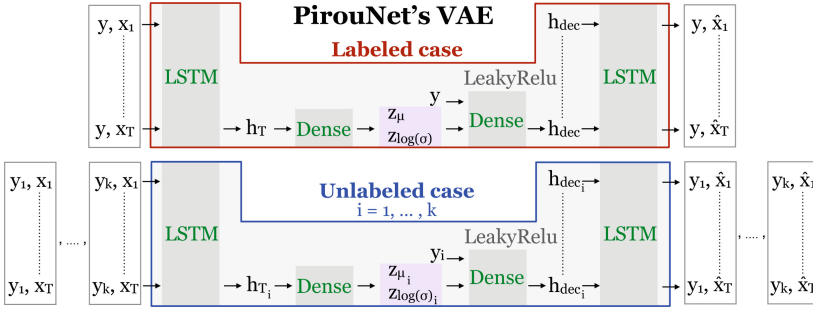


Fig. 3. PirouNet’s supervised and unsupervised training. A dance pose at time t is represented by an input x_t , and a dance sequence is represented by an array of such poses, $x = [x_1, \dots, x_T]$. The last hidden variable h_T in the LSTM output is encoded into a continuous latent variable z . The decoder’s LSTM takes in a hidden variable h_{dec} repeated T times. The Laban Effort y associated to a pose x_t is a categorical latent variable that represents the input label for that particular sequence. There are $k = 3$ distinct possible classes for the label y . Importantly, the VAE in the unlabeled case reconstructs every dance sequence k times, using a different label value each time.

3.1 Conditional Dance Generative Model

We propose the following conditional dance generative model:

$$p_{\pi}(y) = \text{Cat}(y \mid \pi); \quad p(z) = \mathcal{N}(z \mid 0, I); \quad p_{\theta}(x \mid y, z; \theta) = \mathcal{N}(f_{\theta}(y, z), \sigma^2), \quad (1)$$

where z is a continuous latent variable representing the dynamics - i.e. which movement is performed - and y is a categorical variable denoting the choreographer specific or LMA-based label - i.e. how the movement is performed. If no label is available, y is treated as a latent variable. The observed variable x represents a dance sequence. The prior on y is modeled as a multinomial distribution (chosen as a uniform categorical distribution in practice) $\text{Cat}(y \mid \pi)$ with parameter π , while the prior on z is represented by $\mathcal{N}(z \mid 0, I)$, the multivariate standard Gaussian following the VAEs’ classical prior. The latent variables y, z are assumed to be marginally independent. The forward generative model $p_{\theta}(x \mid y, z)$ is a Gaussian distribution parameterized by a mean $f_{\theta}(y, z)$ defined as a non-linear transformation of the latent variables y, z , and a variance σ^2 .

In practice, the non-linear transformation $f_\theta(y, z)$ is implemented as a neural network with weights θ called the decoder f_θ .

The decoder further implements the dance dynamics through a series of temporal updates on the hidden variable h^t :

$$h^{(t+1)} = \text{Dyn}_\theta(h^{(t)}, z, y), \quad t = 1, \dots, T - 1 \quad (2)$$

producing T poses that are concatenated to produce a dance sequence. In practice, PirouNet uses LSTM layers to represent the dynamic operator Dyn_θ [20].

3.2 Semi-supervised Conditional Variational Autoencoder

Based on the conditional dance generative model of Eq. (1), PirouNet solves the inverse problem of inferring the latent variables of movement and intention (z, y) from a dance sequence (x) . PirouNet is built with a conditional variational autoencoder (VAE) architecture [46, 55], modified following the semi-supervised approach proposed in [28] in order to include the (only partially observed) labels of movement intentions, as shown in Fig. 3 and described below.

Encoder: Posterior of the Movement Latent Variable z (Labeled and Unlabeled Cases.) Inference in (conditional) VAEs involves computing an approximate posterior q_ϕ for the continuous latent variable z , which writes:

$$\begin{cases} q_\phi(z|x, y) & \text{when a label } y \text{ is available,} \\ q_\phi(z|x) = \sum_{y=1}^k q_\phi(z, y|x) = \sum_{y=1}^k q_\phi(z|y, x) \cdot q_\phi(y|x) & \text{otherwise.} \end{cases} \quad (3)$$

Both labeled and unlabeled cases in Eq. 3 require the computation of the distribution $q_\phi(z|y, x)$, which is modeled as a multi-variate diagonal Gaussian distribution, as is standard in VAEs [55]. The role of the encoder g_ϕ is to output $q_\phi(z|y, x)$. In the unlabeled case, the encoder is used k times (see Fig. 3): once for each possible value of the label y within the sum defining the unlabeled case in Eq. 3. Specifically, the encoder g_ϕ takes a sequence x and a label y as inputs and outputs the mean $\mu_\phi(x, y)$ and variance $\sigma_\phi^2(x, y)$ of the amortized approximate posterior $q_\phi(z|x, y)$, written as:

$$q_\phi(z|x, y) = \mathcal{N}(g_\phi(x, y)) = \mathcal{N}(\mu_\phi(x, y), \sigma_\phi^2(x, y)). \quad (4)$$

Importantly, the encoder extracts this posterior from the dynamics of the input sequence x , which is represented as:

$$h^{(t+1)} = \text{Dyn}_\phi(h^{(t)}, x, y), \quad t = 1, \dots, T - 1, \quad (5)$$

where the last hidden variable h_T is fed to fully-connected layers to give $\mu_\phi(x, y)$ and $\sigma_\phi^2(x, y)$. In practice, just as for the decoder, this is implemented with LSTM layers [20].

Classifier: Posterior of the Intention Latent Variable y (Unlabeled Case.) The unlabeled case requires the additional computation of the posterior $q_\phi(y|x)$ of

the categorical latent variable y in Eq.(3). This is modeled as a multinomial distribution, as is standard in multi-class classification. Compared to a traditional VAE, we use an additional encoder-classifier (called classifier, for simplicity) whose role is to output $q_\phi(y|x)$. Specifically, the classifier takes a sequence x as input and outputs the k probabilities of the k possible values for the label y in the probability vector $\pi_\phi(x)$, i.e.:

$$q_\phi(y|x) = \text{Cat}(y | \pi_\phi(x)). \quad (6)$$

In practice, the classifier is implemented with fully-connected layers (Fig. 3).

Reparameterization Trick. The VAE reparameterization trick is used to generate a sample z from the approximate posterior $q_\phi(z|y, x)$, i.e. an element of the VAE’s latent space that corresponds to a low-dimensional representation of the dance sequence x . We emphasize that z represents a dynamic *sequence*, and not a static dance pose—which is crucial for the encoding of Laban Efforts that intrinsically characterize dynamics. Specifically, z should be understood as a parameter of dynamics (e.g. describing the movement of a jump or spin), which is fed to the generative model’s LSTM layers.

Decoder/Generator: Conditional Dance Generative Model. The VAE’s decoder corresponds to the implementation of the conditional dance generative model from Eq. (1). The decoder takes the encoded dance sequence z (sampled with the reparameterization trick) and an assigned label y as inputs. During unsupervised training, the decoder is used k times (see Fig. 3): once for each possible value of the label y . The decoder then outputs a reconstructed dance sequence $\hat{x} = f_\theta(z, y)$. To generate original dance choreography, a new random latent variable z is sampled from an approximation of the marginal distribution $q_\phi(z|y)$ - corresponding to the body motion - where the label y is chosen by the user.

3.3 Semi-supervised Training and Loss Function

The training uses the semi-supervised framework and loss function prescribed in [28] which considers the labeled and unlabeled cases separately.

Labeled Case. In the labeled case, our objective is to maximize the log-likelihood $\log p_\theta(x, y)$ to learn the parameters θ of the conditional dance model. This is traditionally performed via the maximization of its lower bound $-\mathcal{L}(x, y)$:

$$\begin{aligned} \log p_\theta(x, y) &\geq \mathbb{E}_{q_\phi(z|x, y)}[\log p_\theta(x | z, y) + \log p_\pi(y)] - \text{KL}(p(z) || q_\phi(z | x, y)) \\ &= -\mathcal{L}(x, y), \end{aligned}$$

where KL denotes the Kullback-Leibler (KL) divergence. In our case, the prior $p_\pi(y)$ is uniform over all labels, and the distribution $p_\theta(x | z, y)$ is modeled as a Gaussian. Thus, the quantity $\mathcal{L}(x, y)$ resembles a regularized L_2 reconstruction loss of a VAE with a modified KL divergence. The L_2 reconstruction loss is averaged over each body joint in a pose, and over each pose per sequence.

Unlabeled Case. In the unlabeled case, y is missing and treated as another latent variable, in addition to z , over which we perform posterior inference. In this case, the objective is to maximize the marginalized log-likelihood $\log p_\theta(x)$, via the maximization of its lower bound $-\mathcal{U}(x)$ written as:

$$\begin{aligned} \log p_\theta(x) &\geq \mathbb{E}_{q_\phi(y,z|x)} [\log p_\theta(x | y, z) + \log p_\theta(y) + \log p(z) - \log q_\phi(y, z | x)] \\ &= \sum_{y=1}^k q_\phi(y | x) (-\mathcal{L}(x, y)) + \mathcal{H}(q_\phi(y | x)) = -\mathcal{U}(x), \end{aligned}$$

where $-\mathcal{H}$ denotes a loss term participating to the entropy:

$$H(q_\phi(y | x)) = -\sum_{y=1}^k q_\phi(y | x) \log q_\phi(y | x) = -\sum_{y=1}^k q_\phi(y | x) \mathcal{H}(q_\phi(y | x)). \quad (7)$$

The loss $-\mathcal{H}(q_\phi(y | x))$ and the regularized reconstruction loss, $\mathcal{L}(x, y)$, are now weighted by the confidence associated to each label y , $q_\phi(y | x)$.

Loss Function. The training uses the loss function computed from a lower bound that encompasses the labeled and unlabeled cases:

$$\text{loss} = \sum_{(x,y) \sim \tilde{p}_l} \mathcal{L}(x, y) + \sum_{x \sim \tilde{p}_u} \mathcal{U}(x) + \alpha \cdot \mathbb{E}_{\tilde{p}_l(x,y)} [-\log q_\phi(y | x)], \quad (8)$$

where $\tilde{p}_l(x, y)$ and $\tilde{p}_u(x)$ are the empirical distributions over the labeled and unlabeled subsets, and the third term is a classification loss weighted by the hyperparameter α which controls the relative weight between dance generation and classification objectives [28].

4 Results

We qualitatively and quantitatively evaluate PirouNet for semi-supervised classification and generation of choreography. During experiments, we find that these two goals are best achieved with different hyperparameters and levels of supervision. Therefore, we present a model primarily trained for the purpose of creating novel and conditional dance (PirouNet_{dance}), as well as a model that achieves better classification accuracy (PirouNet_{watch}).

Like most deep learning tools for dance, we leverage motion data in keypoint format. This format, often included in large movement datasets [36, 45, 52], represents the body as a cloud of 3D points, each representing a unique joint. Available and high-performing pose-estimation software [16, 38, 49] makes keypoint format accessible to smaller, homemade datasets as well. We use half of Pettee’s keypoint dataset [42], featuring a trained contemporary dancer in solo improvisation. We choose this particular style of dance for its predisposition to LMA classification, a tool developed for Western styles of dance. Future work is needed to validate PirouNet for other forms of dance and annotations. In this spirit, we propose a Dash [21] app, pictured in Fig. 2, for easy manual labeling of any keypoint dataset.

4.1 Datasets

The dataset [42] is comprised of 36,396 poses extracted from six uninterrupted dances captured at a rate of 35 frames per second. This amounts to about 20 min of real-time movement of an experienced dancer. Each pose features 53 joints captured in three dimensions, normalized such that the dance fits within inside a unit box. The dancer’s barycenter is fixed to one point on the (x, y) plane. From the pose data, we extract 36,356 sliding sequences of 40 continuous poses, and manually label 350 of these sequences (0.96% of the dataset) which do not share any of the same poses. This takes an experienced dancer (the principal author) about 3 h, identifying if the movement’s Laban Time Effort is Low, Medium, or High. We apply two automated techniques to augment this labeling to 9,167 labeled sequences (representing 25.2% of our unlabeled dataset) in total, split between 45% Low, 34% Medium, and 21% High Efforts. (i) We automatically label all sequences between sequences that share a same Effort. For example, if two back-to-back sequences are deemed to have Low Time Effort, all sequences that are a combination of the poses in these two sequences are also labeled with Low Time effort. (ii) We extend every label to all sequences starting within 6 frames (0.17 s) before or after its respective sequence.

4.2 Training

All experiments are built using the PyTorch library [39] and run on a server with two Nvidia A30 GPUs and two CPUs, each with 16 cores. We train using an Adam optimizer with standard hyperparameters [27]. We present results for the PirouNet architecture resulting from a hyperparameter search using Wandb [7] on batch size, learning rate, number of LSTM and dense layers, as well as hidden variable sizes. PirouNet uses 5 LSTM layers with 100 nodes in both the encoder and the decoder. The classifier features 2 ReLU-activated [37] linear layers with 100 nodes. The latent space is 256-dimensional, which is approximately 25 times smaller than the 6360-dimensional initial space. We train for 500 epochs with a learning rate of $3e^{-4}$ and a batch size of 80 sequences. We select different epochs for PirouNet_{dance} and PirouNet_{watch} to minimize validation loss. For unsupervised training, we use 35,538 40-pose sequences, with the remaining sequences being reserved for testing. For supervised training, PirouNet_{dance} and PirouNet_{watch} are trained on 79% (16.6% of entire training set) and 92% (18.8% of entire training set) of the labeled sequences, respectively. We reserve 5% of the labeled sequences for validation, and 3% for testing.

4.3 Semi-supervised Classification of Laban Efforts

While the primary purpose of this work is to generate Effort-specific dance, we examine classification performance to investigate training: we evaluate the accuracy of the classifier in attributing the correct categorical Time Effort to a dance. Because evaluating Laban efforts is subjective, the ground-truth labels provided by the human labeler may be considered “noisy”. To account for this

in our evaluation of PirouNet_{watch}, the human labeler also evaluates their own self-accuracy by relabeling the validation and entire datasets (Fig. 4(b, c)).

When classifying the validation set (Fig. 4a), PirouNet_{watch} succeeds with 50.1% accuracy, which represents 75% of the labeler’s self-accuracy. In the case of test data, PirouNet_{watch} is 72% as accurate as the labeler. This is a satisfactory performance, considering the subjective nature of the task, the limited access to labeled data, and the observation that classifying the validation set was challenging even for the human labeler.




Fig. 4. Confusion matrices for classification of validation dance sequences’ Time effort as being “Low”, “Medium”, or “High”. a. PirouNet_{watch} classifies the validation dataset. b. The labeler reclassifies a comparable validation dataset. c. The labeler reclassifies a comparable entire dataset (train, validation, and test).

We hypothesize that results could be further improved by increasing the training dataset size, as well as optimizing the trade-off between training the VAE and the classifier. This would include examining the impact of the loss’s control term α during later training. A hyperparameter search limited to early training shows $\alpha = 0.1 \cdot \frac{n_{\text{unlabeled}}}{n_{\text{labeled}}}$ is optimal, where $n_{\text{unlabeled}}$ and n_{labeled} are the unlabeled and labeled dataset sizes. Another potential solution would be switching to a regression setting. Using continuous scalar variables as labels would better reflect the continuous nature of the Effort spectrum and be better suited to deal with noise between Low and High Efforts. Figure 4b shows an example of how the discontinuous categories can result in strong self-disagreement for humans, which affects the performance of the classifier on those same sequences (Fig 4a). While selected for classification, PirouNet_{watch}’s generation performance is also satisfactory.

It is worth noting that the ground truth and relabeled sets of data are not exactly identical in labeling procedure: while the dance sequences are the same, a sequence that was manually classified in one dataset may have been classified via label augmentation in the other. We do not expect this to significantly impact training, as the augmentation procedure was designed to agree with manual labeling. However, this signifies that a misclassification of a small number of sequences by the labeler can propagate through augmentation, and result in 100% misclassification for a given label, as is the case in Fig. 4b. We note that unlike PirouNet_{watch}, the labeler almost never re-predicts a sequence to be more than one categorical bin away from the ground truth. This highlights the use of

regression classification as a means of preserving information about the labels’ meanings. We leave the improvements discussed above for future iterations.

4.4 Reconstructions of Choreographies

Qualitative Assessment. Artifacts obtained on validation and test datasets, two of which are depicted in Fig. 5, show that PirouNet_{dance} successfully encodes and decodes a given sequence with very little variation. Reconstructions capture complex movements that include rotations (“pirouettes”), changes in height, and changes in velocity. The Laban Time Effort is preserved, as indicated by the identical time evolution of the reconstructions.




Fig. 5. Reconstructions of dance sequences. a. Sequence extracted from test data featuring a fast pirouette. b. Sequence extracted from validation data featuring a multi-axial rotation on the floor with gradual leg and arm extension.

Quantitative Assessment. To assess PirouNet_{dance}’s ability to reconstruct movement, we use the root mean squared error representing the average joint distance (AJD) between original and reconstructed sequences, pose per pose. For N sequences with $T = 40$ poses of $J = 53$ joints each, the AJD writes as:

$$\bar{D} = \frac{1}{N \cdot T \cdot J} \sum_{n=1}^N \sum_{t=1}^T \sum_{j=1}^J \sqrt{|x_j^{t,n} - \hat{x}_j^{t,n}|^2 + |y_j^{t,n} - \hat{y}_j^{t,n}|^2 + |z_j^{t,n} - \hat{z}_j^{t,n}|^2}. \quad (9)$$

For the validation dataset ($N = 455$), the average joint difference $\bar{D}_{\text{validation}}$ is $2.8e^{-2}$. On the test dataset ($N = 249$), the average difference is $\bar{D}_{\text{test}} = 1.8e^{-2}$, which, remarkably, matches \bar{D}_{train} ($N = 7,161$). Since the joint coordinates are scaled to fit inside a unit box, $\bar{D}_{\text{validation}}$ and \bar{D}_{test} represent less than 2.8% and 1.8% of the dancer’s height, 4.8 cm and 3.1 cm. While trained for classification purposes, PirouNet_{watch} has a satisfactory reconstruction error, with $\bar{D}_{\text{validation}} = 3.6e^{-2}$ and $\bar{D}_{\text{test}} = 2.0e^{-2}$.

4.5 Creation of Choreographies Based on Laban Efforts

We propose conditionally generating dance by sampling from an approximation of the marginal conditional distribution $q_\phi(z|y)$ of the latent variable z , given an input Laban effort y . Effectively, this means sampling dance from neighborhoods in the latent space featuring a high density of previously encoded same-Effort

sequences (manually or automatically labeled). Future work will focus on disentangling the latent space in order to enable conditional generation independently from previously encoded data. Outside of the conditional framework, PirouNet can generate dance in the general style of the training data by sampling random latent variables.

Qualitative Assessment. For each of the three labels (Low/Medium/High Effort), we conditionally generate 75 sequences from PirouNet_{dance}'s marginal distribution. The assessment of this 225-sequence benchmark set is two-fold.

First, we examine the overall “danceability” of the generated material. We determine danceable material as movement that (i) would be physically realizable by the training set’s dancer in at least one physical environment and (ii) is largely continuous in space and time. Within this broad category, we find three sub-categories, outlined in Fig. 6. Most of the danceable sequences belong to at least two of these sub-categories, which shows that PirouNet_{dance} has learned to produce danceable materials.

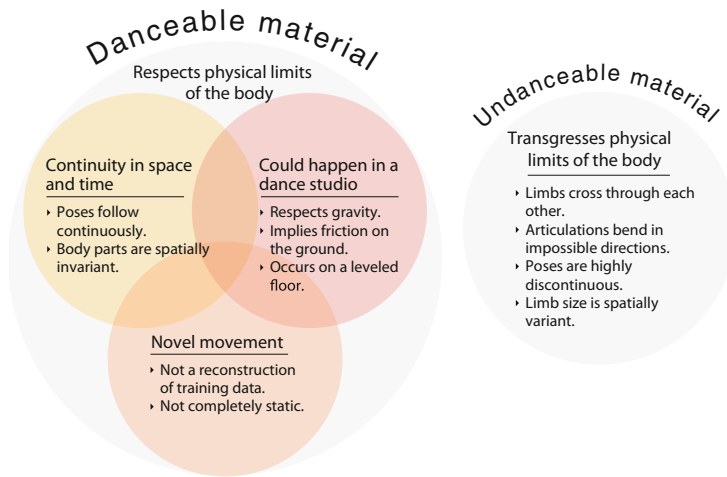


Fig. 6. Infographic of danceable versus non-danceable material, as generated by PirouNet_{dance}.

Second, we examine PirouNet_{dance}'s display of Time Effort in the danceable material. Figure 7 depicts PirouNet_{dance} originals grouped by intended Effort. Looking at the video format of such sequences, we remark that instances of discontinuity between poses, appearing as quick jumps, arise in a sequence's first or last five frames. The (danceable) movement is otherwise realistic, and often totally reproducible by an experienced dancer. We provide video examples at github.com/bioshape-lab/pirounet.

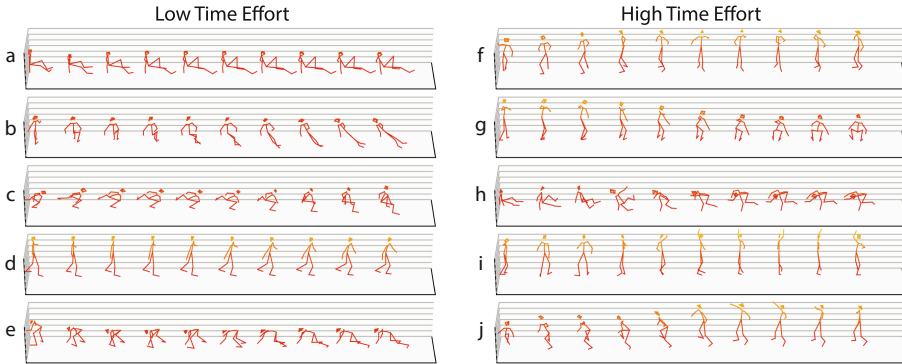


Fig. 7. Low/High Time Effort dance sequences created by PirouNet_{dance}. The labeler describes the experience of Time Effort as relating to the presence (or absence) of impulse, sense of urgency, and explosive momentum. a. Slow and continuous leg extension from V-sit. b. Continuous transition from crouched to plank position. c. Smooth weight transfer through deep pli . d. Forward standing weight transfer through small pli . e. Transition from standing to plank position through double rond-de-jambe. f. Demi-pirouette lead by elbow, ending in quick change in direction. g. Spontaneous drop to a crouched position. h. Leg swing leads a sharp torso rotation. i. Pirouette with arm thrown up. j. Explosive extension of arms and torso from pli .

Quantitative Assessment. We evaluate PirouNet_{dance}'s Effort-centric generative model using the 225-sequence benchmark set. The labeler identifies 96.0% of the dataset as danceable, which is 11.5% better than a benchmark set randomly sampled from the latent space. Of the danceable portion, 98.1% of the conditionally sampled set respects the constraints of the physical environment, compared to 72.2% of the randomly sampled set. Only 0.46% of the former shows some non-physical discontinuity, such as a noticeable change in limb size, versus 8.8% for the latter. Therefore, the marginal conditional method proves to be advantageous for the purposes of creating realizable and meaningful dance for the user, outside of providing Effort-specific movement. To determine PirouNet_{dance}'s accuracy in generating appropriate Efforts, the labeler blindly labels the danceable portion of the benchmark set without knowledge of PirouNet_{dance}'s intended Effort. On average, the human labeler correctly classifies 63% of PirouNet_{dance}'s sequences for a given Effort (Fig. 8). This represents 83% of the labeler's own self-agreement on the validation and training data (Fig. 4c).




Fig. 8. Confusion matrix on the danceable portion of the benchmark set (225 sequences). Generation is the least accurate for High Effort, possibly due to the smaller portion of High Effort labels in the training data. Notably, the human classifier correctly recognizes the majority of the PirouNet’s intended Efforts.

5 Conclusions

This paper presents PirouNet, the first Laban Effort based deep generative tool for dance, made possible via a semi-supervised conditional dynamic variational autoencoder. While illustrated on Laban Efforts, our tool can be readily used with any choreographer-specific categorical labels. Our keypoint labeling web-app enables anyone to classify any keypoint data-set with any annotations, empowering artists to shape their own AI-tool. We demonstrate PirouNet’s ability to classify the Laban Time Effort and generate new choreography based on this Effort as comparable to that of an experienced dancer. Conditional generation from the marginal distribution $q_\phi(z|y)$ offers largely accurate, diverse, and realizable choreography. Randomly sampling from the latent space at large also produces realistic and novel dance, providing an exciting choreographic tool outside of categorical generation. Future work will focus on improving classification accuracy and disentanglement of the latent space, as well as examining PirouNet’s implementation and functionality for actual choreographic practices. We hope this artist-centric, adaptable tool will act as a launching pad for engaging with old repertoire and inspiring new choreography from a completely new creative standpoint. While demonstrated on dance, the proposed method can be extended to other forms of art creation, inspiring AI-based tools tailored to the style and intuition of their artist.

References

1. Alemi, O., Françoise, J., Pasquier, P.: GrooveNet: real-time music-driven dance movement generation using artificial neural networks. *Networks* **8**(17), 26 (2017)
2. Aristidou, A., Stavrakis, E., Charalambous, P., Chrysanthou, Y., Himona, S.L.: Folk dance evaluation using Laban movement analysis. *J. Comput. Cult. Heritage* **8**(4), 1–19 (2015). <https://doi.org/10.1145/2755566>
3. Berman, A., James, V.: Kinetic imaginations: exploring the possibilities of combining AI and dance. In: *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI 2015*, pp. 2431–2437. AAAI Press (2015)

4. Berman, A., James, V.: Learning as performance: autoencoding and generating dance movements in real time. In: Liapis, A., Romero Cardalda, J.J., Ekárt, A. (eds.) *EvoMUSART 2018*. LNCS, vol. 10783, pp. 256–266. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-77583-8_17
5. Bernardet, U., Fdili Alaoui, S., Studd, K., Bradley, K., Pasquier, P., Schiphorst, T.: Assessing the reliability of the Laban Movement Analysis system. *PLoS One* **14**(6), e0218179 (2019)
6. Bernstein, R., Shafir, T., Tsachor, R., Studd, K., Schuster, A.: Laban movement analysis using kinect. *Int. J. Comput. Inform. Eng.* **9**(6), 1567–1571 (2015)
7. Biewald, L.: Experiment tracking with weights and biases (2020). <https://www.wandb.com/>
8. Brooks, L.M.: Harmony in space: a perspective on the work of Rudolf Laband. *J. Aesthetic Educ.* **27**(2), 29–41 (1993)
9. Carlson, K., Schiphorst, T., Pasquier, P.: Scuddle: generating movement catalysts for computer-aided choreography. In: *ICCC*, pp. 123–128 (2011)
10. Choensawat, W., Hachimura, K.: Generating stylized dance motion from labanotation by using an autonomous dance avatar. In: *GRAPP/IVAPP* (2012)
11. Choensawat, W., Nakamura, M., Hachimura, K.: GenLaban: a tool for generating labanotation from motion capture data. *Multimed. Appl.* **74**(23), 10823–10846 (2015)
12. Davies, E.: *Beyond Dance: Laban’s Legacy of Movement Analysis*. Routledge (2007)
13. Davis, J.: Laban movement analysis: a key to individualizing children’s dance. *J. Phys. Educ. Recreat. Dance* **66**(2), 31–33 (1995)
14. El Raheb, K., Ioannidis, Y.: Annotating the captured dance: reflections on the role of tool-creation. *Int. J. Perform. Arts Digit. Media* **17**(1), 118–137 (2021)
15. Ewan, V., Sagovsky, K.: *Laban’s Efforts in Action: A Movement Handbook for Actors with Online Video Resources*. Bloomsbury Publishing (2018)
16. Fang, H.S., Xie, S., Tai, Y.W., Lu, C.: RMPE: regional multi-person pose estimation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2334–2343 (2017)
17. Groff, E.: Laban movement analysis: charting the ineffable domain of human movement. *J. Phys. Educ. Recreat. Dance* **66**(2), 27–30 (1995)
18. Guo, C., et al.: Action2Motion: conditioned generation of 3D human motions. *CoRR abs/2007.15240* (2020)
19. Hamburg, J.: Coaching athletes using Laban movement analysis. *J. Phys. Educ. Recreat. Dance* **66**(2), 34–37 (1995). <https://doi.org/10.1080/07303084.1995.10607040>
20. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
21. Hossain, S., Calloway, C., Lippa, D., Niederhut, D., Shupe, D.: Visualization of bioinformatics data with Dash bio. In: *Proceedings of the 18th Python in Science Conference*, pp. 126–133 (2019)
22. Hutchinson, A.: *Labanotation: The System of Analyzing and Recording Movement*, No. 27, Taylor & Francis (1977)
23. Jang, M., Kim, D., Kim, Y., Kim, J.: Automated dance motion evaluation using dynamic time warping and Laban movement analysis. In: *2017 IEEE International Conference on Consumer Electronics (ICCE)*, pp. 141–142. IEEE (2017)
24. Ji, Y., Xu, F., Yang, Y., Shen, F., Shen, H.T., Zheng, W.S.: A large-scale RGB-D database for arbitrary-view human action recognition. In: *Proceedings of the 26th ACM International Conference on Multimedia*, pp. 1510–1518 (2018)

25. Kim, D., Kim, D.H., Kwak, K.C.: Classification of K-Pop dance movements based on skeleton information obtained by a Kinect sensor. *Sensors* **17**(6), 1261 (2017)
26. Kim, H.J., Neff, M., Lee, S.H.: The perceptual consistency and association of the LMA effort elements. *ACM Trans. Appl. Percept.* **19**(1) (2022). <https://doi.org/10.1145/3473041>
27. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. In: 3rd International Conference on Learning Representations, ICLR 2015 (2015)
28. Kingma, D.P., Mohamed, S., Jimenez Rezende, D., Welling, M.: Semi-supervised learning with deep generative models. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., Weinberger, K. (eds.) *Advances in Neural Information Processing Systems*, vol. 27. Curran Associates, Inc. (2014)
29. Knight, H., Simmons, R.: Expressive motion with x, y and theta: Laban effort features for mobile robots. In: *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pp. 267–273. IEEE (2014)
30. Lee, H.Y., et al.: Dancing to music. *Adv. Neural Inf. Process. Syst.* **32** (2019)
31. Li, B., Zhao, Y., Zhelun, S., Sheng, L.: DanceFormer: music conditioned 3D dance generation with parametric motion transformer. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, pp. 1272–1279 (2022)
32. Li, R., Yang, S., Ross, D.A., Kanazawa, A.: AI choreographer: music conditioned 3D dance generation with AIST+++. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13401–13412 (2021)
33. Maletic, V.: *Body-Space-Expression: The Development of Rudolf Laban's Movement and Dance Concepts*, vol. 75. Walter de Gruyter (2011)
34. McCormick, J., Vincs, K., Nahavandi, S., Creighton, D.: Learning to dance with a human. *ISEA International* (2013). <http://hdl.handle.net/2123/9638>
35. Mentis, H.M., Johansson, C.: Seeing movement qualities. In: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2013*. Association for Computing Machinery, New York (2013). <https://doi.org/10.1145/2470654.2466462>
36. Müller, M., Röder, T., Clausen, M., Eberhardt, B., Krüger, B., Weber, A.: *Documentation MoCap database HDM05* (2007)
37. Nair, V., Hinton, G.E.: Rectified linear units improve restricted Boltzmann machines. In: *ICML*, pp. 807–814 (2010)
38. Newell, A., Yang, K., Deng, J.: Stacked hourglass networks for human pose estimation. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016*. LNCS, vol. 9912, pp. 483–499. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46484-8_29
39. Paszke, A., et al.: PyTorch: an imperative style, high-performance deep learning library. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 32, pp. 8024–8035. Curran Associates, Inc. (2019)
40. Payne, H.: The psycho-neurology of embodiment with examples from authentic movement and Laban movement analysis. *Am. J. Dance Ther.* **39**(2), 163–178 (2017)
41. Petrovich, M., Black, M.J., Varol, G.: Action-conditioned 3D human motion synthesis with transformer VAE. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10985–10995 (2021)
42. Pettee, M., Shimmin, C., Duhaime, D., Vidrin, I.: Beyond imitation: generative and variational choreography via machine learning. In: *Proceedings of the Tenth International Conference on Computational Creativity*, pp. 196–203 (2019)

43. Rumelhart, D.E., Hinton, G.E., Williams, R.J.: Learning internal representations by error propagation. Technical report, California Univ San Diego La Jolla Inst for Cognitive Science (1985)
44. Santos, L., Dias, J.: Laban movement analysis towards behavior patterns. In: Camarinha-Matos, L.M., Pereira, P., Ribeiro, L. (eds.) DoCEIS 2010. IAICT, vol. 314, pp. 187–194. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-11628-5_20
45. Shahroudy, A., Liu, J., Ng, T.T., Wang, G.: NTU RGB+ D: a large scale dataset for 3D human activity analysis. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1010–1019 (2016)
46. Sohn, K., Yan, X., Lee, H.: Learning structured output representation using deep conditional generative models. In: Proceedings of the 28th International Conference on Neural Information Processing Systems, NIPS 2015, vol. 2, p. 3483–3491. MIT Press, Cambridge (2015)
47. Sun, G., et al.: A virtual reality dance self-learning framework using Laban movement analysis. *J. Eng. Sci. Technol. Rev.* **10**, 25–32 (2017). <https://doi.org/10.25103/jestr.105.03>
48. Tan, F., Woo, G., Tsang, H.H.: CGLER: laban effort framework analysis with conducting gestures using neural networks. In: 2020 IEEE Symposium Series on Computational Intelligence (SSCI), pp. 1452–1459. IEEE (2020)
49. Toshev, A., Szegedy, C.: DeepPose: human pose estimation via deep neural networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1653–1660 (2014)
50. Truong, A., Boujut, H., Zaharia, T.: Laban descriptors for gesture recognition and emotional analysis. *Vis. Comput.* **32**(1), 83–98 (2016)
51. Tsachor, R.P., Shafir, T.: A somatic movement approach to fostering emotional resiliency through Laban movement analysis. *Front. Hum. Neurosci.* **11**, 410 (2017)
52. Tsuchida, S., Fukayama, S., Hamasaki, M., Goto, M.: AIST dance video database: multi-genre, multi-dancer, and multi-camera database for dance information processing. In: ISMIR, no. 5, p. 6 (2019)
53. Wang, J., Miao, Z., Xie, N., Xu, W., Li, A.: Labanotation generation from motion capture data for protection of folk dance. *IEEE Access* (2020). <https://doi.org/10.1109/ACCESS.2020.3014157>
54. Wang, S., Li, J., Cao, T., Wang, H., Tu, P., Li, Y.: Dance emotion recognition based on Laban motion analysis using convolutional neural network and long short-term memory. *IEEE Access* **8**, 124928–124938 (2020)
55. Welling, M., Kingma, D.P.: Auto-encoding variational bayes. In: ICLR (2014)
56. Whittier, C.: Laban movement analysis approach to classical ballet pedagogy. *J. Dance Educ.* **6**(4), 124–132 (2006)
57. Wilke, L., Calvert, T., Ryman, R., Fox, I.: From dance notation to human animation: the LabanDancer project. *J. Vis. Comput. Animat.* **16**, 201–211 (2005). <https://doi.org/10.1002/cav.90>
58. Young, J., Wood, L.L.: Laban: a guide figure between dance/movement therapy and drama therapy. *Arts Psychother.* **57**, 11–19 (2018)
59. Zacharatos, H., Gatzoulis, C., Chrysanthou, Y., Aristidou, A.: Emotion recognition for exergames using Laban movement analysis. In: Proceedings of Motion on Games, pp. 61–66. Association for Computing Machinery (2013)
60. Zhang, S., Li, Q., Yu, T., Shen, X.J., Geng, W., Wang, P.: Implementation of a notation-based motion choreography system. In: Zha, H., Pan, Z., Thwaites, H., Addison, A.C., Forte, M. (eds.) VSMM 2006. LNCS, vol. 4270, pp. 495–503. Springer, Heidelberg (2006). https://doi.org/10.1007/11890881_54

61. Zhang, X., Miao, Z., Yang, X., Zhang, Q.: An efficient method for automatic generation of labanotation based on bi-directional LSTM. *J. Phys. Conf. Ser.* **1229**(1) (2019). <https://doi.org/10.1088/1742-6596/1229/1/012031>
62. Zhou, Y., Li, Z., Xiao, S., He, C., Huang, Z., Li, H.: Auto-conditioned recurrent networks for extended complex human motion synthesis. In: *International Conference on Learning Representations* (2018)