

TM2T: Stochastic and Tokenized Modeling for the Reciprocal Generation of 3D Human Motions and Texts

Chuan Guo^[0000-0002-4539-0634], Xinxin Zuo, Sen Wang, and Li Cheng

University of Alberta
{cguo2, xzuo, sen9, lcheng5}@ualberta.ca

Abstract. Inspired by the strong ties between vision and language, the two intimate human sensing and communication modalities, our paper aims to explore the generation of 3D human full-body motions from texts, as well as its reciprocal task, shorthanded for text2motion and motion2text, respectively. To tackle the existing challenges, especially to enable the generation of multiple distinct motions from the same text, and to avoid the undesirable production of trivial motionless pose sequences, we propose the use of motion token, a discrete and compact motion representation. This provides one level playing ground when considering both motions and text signals, as the motion and text tokens, respectively. Moreover, our motion2text module is integrated into the inverse alignment process of our text2motion training pipeline, where a significant deviation of synthesized text from the input text would be penalized by a large training loss; empirically this is shown to effectively improve performance. Finally, the mappings in-between the two modalities of motions and texts are facilitated by adapting the neural model for machine translation (NMT) to our context. This autoregressive modeling of the distribution over discrete motion tokens further enables non-deterministic production of pose sequences, of variable lengths, from an input text. Our approach is flexible, could be used for both text2motion and motion2text tasks. Empirical evaluations on two benchmark datasets demonstrate the superior performance of our approach on both tasks over a variety of state-of-the-art methods. Project page: <https://ericguo5513.github.io/TM2T/>

Keywords: Motion captioning, text-to-motion generation

1 Introduction

The interplay of vision and language is important in our daily life and social functions. It has motivated considerable research progresses in related topics such as image or video captioning [52,47], and language grounded generation of images or videos [57,53,24]. On the other hand, when coming to human motion analysis, the connections between visual and textural aspects of human motions




Fig. 1. An illustration of our bidirectional TM2T approach that captures the interplay between text (left) and 3D human motion (right) through the text2motion and motion2text modules. Note the stochastic nature of our text2motion module allows the generation of different 3D motions from the same textual description.

are much less studied. Existing efforts primarily focus on unidirectional mapping of either motion captioning (motion2text) [13,41] or language grounded motion generation (text2motion) [2,11,25], with only two [35,54] exploring the integration of visual 3D motions and their textual descriptions. However, both studies tend to produce static pose sequences when motion lengths are longer than 3-4 seconds. Both requires as input the initial pose & target motion length. They are also deterministic methods. That is, each of them always generates the same motions from a given text script. The first phenomenon of lifeless motions could be largely attributed to the direct use of raw 3D poses as their motion representation, which is unnecessarily redundant and yet fails to capture the local contexts of the underlying motion dynamics. The second issue is rooted in their deterministic motion generation processes, that are in contrary to our daily experiences, where multiple distinct motion styles often exist for a character to perform under a same textual script. The conditioning on initial state and target length further imposes strict constraint toward being practically feasible.

The aim of this paper is to investigate the bi-directional ties between 3D human full-body motions and their language descriptions, as illustrated in Fig. 1. Given the asymmetric nature of the two underlying tasks, where text2motion is typically a much harder problem than the reciprocal task, motion2text, our primary focus is text2motion, with a secondary emphasis on motion2text. It is worth noting that in our approach, the module (also called motion2text for simplicity) developed for motion2text task, is also utilized as an integral part of our text2motion training process, referred to as inverse alignment in Fig. 2(c). Empirical evidences suggest the benefit of this strategy in improving our performance for the text2motion problem. To address the lifeless motion issue, we introduce motion token, a compact and semantically rich representation for 3D motions. This is achieved by adapting the deep vector quantization [43] in our context to learn a spatial-temporal codebook from the 3D pose sequences in the training set, with each entry in the codebook describing a particular kind of mo-

tion segments. 3D motions are then reconstructed by decoding the compositions of a list of codebook entries. This way, a 3D human motion is represented as a list of motion tokens (i.e. discrete indices to the codebook entries), each encoding its local spatial-temporal context. This discrete representation also facilitates the follow-up neural machine translators (NMTs) [44,4] to construct mappings between the stream of motion tokens from the motion side, and the stream of text tokens from the language side. Furthermore, our proposed approach is able to explicitly model the underlying distribution of 3D motions conditioned on texts, instead of regressing the mean motions as in previous works [2,11,25,54,35], thus allows non-deterministic text2motion generation.

Our main contributions can be summarized as follows: (i) a motion token representation that compactly encodes 3D human motions. Together with the other key ingredients, including NMT mappings in-between the motion-token and text-token sequences, the motion2text-based inverse alignment, as well as the distribution sampling for non-deterministic predictions, our approach is capable of generating 3D motions (i.e. pose sequences) that are distinct in their lengths and styles, visually pleasing, and importantly, semantically faithful to the same input script. Our approach is also flexible, in that it can be use for both text2motion and motion2text tasks. (ii) Extensive empirical evaluations over two motion-language benchmark datasets demonstrate the superior performance of our approach over a variety of state-of-the-art methods when examined on each of the two tasks.

2 Related Work

Image/Video Captioning and Motion2text. Vision grounded text generation has a long history with extended literature. Here we only focus on the closely related topic of image and video captioning. Early methods [22,21] commonly approach this problem by tagging parts of sentences such as nouns and verbs from visual contents, followed by filling in pre-defined sentence templates. With the advent of deep neural networks, the tools used for visual captioning have been significantly changed. Take [47] for example, it starts by extracting high-level image features from pre-trained GoogleNet, which are then fed into a LSTM decoder to produce captions. In [46], an RNN-based video captioning model is considered, that extracts individual frame features from pre-trained CNN, and translates them to sentences through sequence-to-sequence learning. Further extensions are made through e.g. incorporating attention mechanism for better vision-language alignment [52,49]. More recent methods consider the use of various deep learning apparatus such as GANs [17,30], deep reinforcement learning [10,36], and transformers [8,12].

In contrast, research efforts on captioning 3D human motions are considerably more limited. [41] learns the mapping from human motions to language relying on two statistical models: one associates motions with words; the other assembles words back to form sentences. Recurrent networks are utilized by [54,35] to address this task. In [54], motion and text features are extracted by

two autoencoders respectively; this is followed by generating texts and motions from each other through shared latent vectors. Sequence-to-sequence RNNs are adopted in [35] to translate motions to scripts. Recently, the work of [13] proposes SeqGAN that extends NMT model with a discriminator. Some common issues with existing motion2text results are typically short in length, often incomplete in content, and sometimes lack in details.

Human Motion Modeling and Text2motion. The importance of human motion modeling has been manifested through the extensive research efforts in recent years, where motions are produced based on various forms of inputs, such as partial pose sequences, control signals, action category, and text. Future motion prediction aims to generate short [51] and long [27,31] future pose sequences based on partial pose sequences. This has been traditionally modeled in one-to-one mapping fashion until recent works [3,56,28] that take account the stochastic nature of human motion dynamics. The efforts of [48,1,6,7] proceed to predict multi-person or scene-aware 3D motions. Meanwhile, [19,18,40] attempts to model human motions according to instant control signals such as velocity and directional readouts. In [19], feet contact information is fed into a phase function to produce blending weights of four expert MLP networks. The blended MLP network then predicts next pose state given current state and goal control signals. This is extended in [43,18] where the phase function is replaced by a learnable gating network. Action category based human motion generation also draws considerable interests by resorting to a diverse range of learning strategies, including GANs [50], VAEs [16,15], Transformers [33] and GCNs [55].

In terms of text based human motion modeling (text2motion), the sequence-to-sequence RNN models have been considered by [25,35]; in [2], a latent embedding space is proposed, which is shared by both text and pose sequences and is trained via curriculum learning. The work of [11] considers the topology of human skeleton, and proposes a hierarchical two-stream pose generator. Note existing techniques developed in text2motion are predominantly deterministic. This is in contrast to our proposed stochastic motion generation process.

Discrete Vector Quantization. [43] advocates the quantization of continuous features into discrete latent representation by training a variational autoencoder. This is followed up by several more recent efforts to improve the representation quality and reconstruction accuracy, including hierarchical feature representation [39], gumbel-softmax relaxation [38] and adversarial training [9]. In [32], hierarchical vector quantization is carried out in encoding and generating diverse image patches for inpainting; the work of [37] leverages quantized video frame representation to synthesize future frames. These prior arts inspire the motion token scheme considered in our approach.

3 Our Approach

In what follows, we first detail how discrete motion tokens are obtained from raw 3D motions via vector quantization in Sec. 3.1. Based on this new motion representation, autoregressive NMT networks are used for modeling the bi-modal

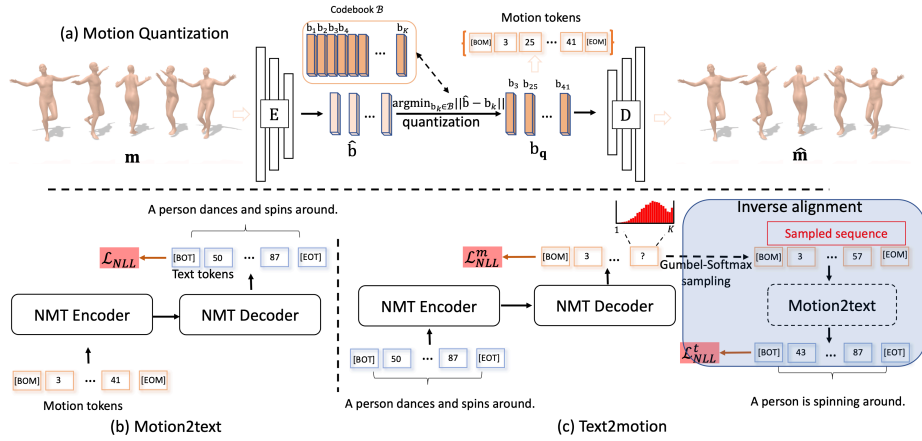


Fig. 2. Approach overview. (a) A 1D CNN based latent quantization model is firstly learned to reconstruct training motions. After training, a motion can be subsequently converted to a tuple of discrete motion tokens (i.e., codebook-indices). [BOM] and [EOM] are indicators of start and end added in a motion token sequence. (b-c) Mappings between motion and text tokens are modeled by autoregressive NMT networks and optimized by maximizing the log-likelihood of the targets (\mathcal{L}_{NLL} and \mathcal{L}_{NLL}^m). (c) While training text2motion, motion tokens sampled from the resulting discrete distributions are inversely mapped to the text space via the learned motion2text model. Loss \mathcal{L}_{NLL}^t penalizes the inverse alignment error. Finally, the 3D pose sequence is obtained by decoding motion tokens via the decoder D in (a).

mappings of motion2text (Sec. 3.2) and text2motion (Sec. 3.3), with inverse alignment elaborated in Sec. 3.3.

3.1 Motion Tokens

We pre-train a latent quantization model on 3D human motions as presented in Fig. 2 (a). Given the pose sequence $\mathbf{m} \in \mathbb{R}^{T \times D_p}$, where T denotes the number of poses and D_p pose dimension, a series of 1D convolutions are applied along the time (i.e. 1st) dimension that yields latent vectors $\hat{\mathbf{b}} \in \mathbb{R}^{t \times d}$ ($t < T$) with d being number of convolution kernels. This process could be written as $\hat{\mathbf{b}} = \mathbf{E}(\mathbf{m})$.

Then, $\hat{\mathbf{b}}$ is transformed to a collection of codebook entries $\mathbf{b}_q \in \mathbb{R}^{t \times d}$ through discrete quantization. Specifically, the learnable codebook $\mathcal{B} = \{\mathbf{b}\}_{k=1}^K \subset \mathbb{R}^d$ consists of K latent embedding vectors with each a d -dimensional vector. The process of quantization $Q(\cdot)$ is operated by replacing each row vector $\hat{\mathbf{b}}_i \in \mathbb{R}^d$ in $\hat{\mathbf{b}}$ with its nearest codebook entry \mathbf{b}_k in \mathcal{B} , defined as

$$\mathbf{b}_q = Q(\hat{\mathbf{b}}) := \left(\operatorname{argmin}_{\mathbf{b}_k \in \mathcal{B}} \|\hat{\mathbf{b}}_i - \mathbf{b}_k\| \right) \in \mathbb{R}^{t \times d}. \quad (1)$$

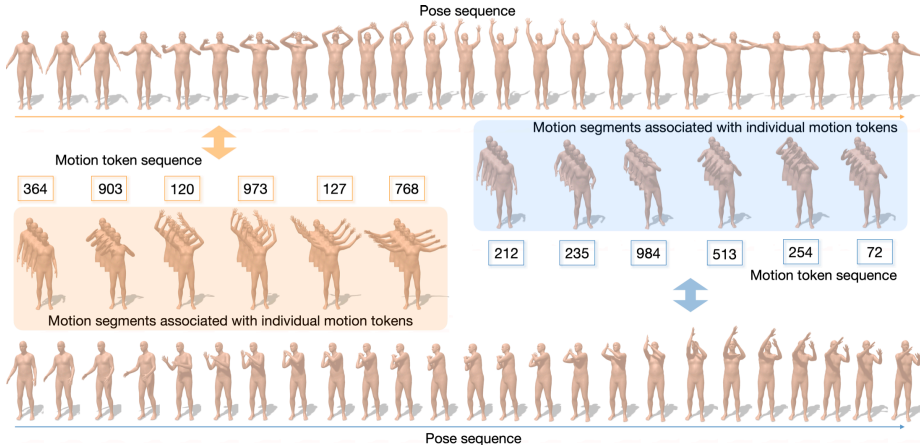


Fig. 3. Exemplar results of motion tokens (middle) and their corresponding pose sequences (top and bottom). Here two 24-frame pose sequence examples are presented; each is reconstructed from a motion token sequences of size 6. Each motion token is associated with a specific local spatial-temporal context, visualized in 4-frame motions.

A following de-convolutional decoder D projects \mathbf{b}_q back to the 3D motion space as a pose sequence, $\hat{\mathbf{m}}$. Now, the entire process can be formulated as

$$\hat{\mathbf{m}} = D(\mathbf{b}_q) = D(Q(E(\mathbf{m}))). \quad (2)$$

This is trained via a reconstruction loss combined with embedding commitment loss terms that encourage latent alignment and stabilize training process:

$$\mathcal{L}_{vq} = \|\hat{\mathbf{m}} - \mathbf{m}\|_1 + \|\text{sg}[E(\mathbf{m})] - \mathbf{b}_q\|_2^2 + \beta\|E(\mathbf{m}) - \text{sg}[\mathbf{b}_q]\|_2^2, \quad (3)$$

where $\text{sg}[\cdot]$ denotes the stop-gradient operation, and β a weighting factor. Straight-through gradient estimator [43] is employed to allow gradient backpropagation through the non-differentiable quantization operation in Eq.(1) that simply copies the gradients from the decoder D to the encoder E .

During inference, a pose sequence $\mathbf{m} \in \mathbb{R}^{T \times D_p}$ can be represented as a sequence of discrete codebook-indices $s \in \{1, \dots, |\mathcal{B}|\}^t$ (namely *motion tokens*) of quantized embedding vectors \mathbf{b}_q , where $s_i = k$ such that $(\mathbf{b}_q)_i = \mathbf{b}_k$. By mapping motion tokens back to their corresponding codebook entries $\mathbf{b}_q = (\mathbf{b}_{s_i})$, human poses are then readily recovered using decoder $\hat{\mathbf{m}} = D(\mathbf{b}_q)$. [BOM] and [EOM] are respectively added to the start and end of a motion token sequence as boundary indicators.

Motion Token Contexts. With vector quantization, each motion token is associated with a particular type of motion contexts, thus a 3D motion can be regarded as a meaningful composition of motion tokens. We decode each entry in the learned codebook \mathcal{B} using decoder D and get 4-frame motion segments

($t = \frac{T}{4}$ in our setting) that reflect the contexts associated with individual motion tokens. Fig. 3 presents two raw pose sequences and their motion token representations, as well as the associated motion segments. We can observe that, with global dependencies maintained in motion token sequences, each motion token successfully captures the spatial-temporal characteristics in local contexts.

3.2 Learning Motion2text

Given tokenized motion representation, we are able to efficiently build mapping from human motions to texts using NMT models such as Transformer [44]. Assume the target is a sequence of text tokens $x \in \{1, \dots, |\mathcal{V}|\}^N$, where \mathcal{V} is the word vocabulary and N number of words in the description. As described in Fig. 2 (b), source motion tokens are fed into Transformer encoder and then the decoder predicts the probability distribution of possible discrete text tokens at each step $p_\theta(x|s) = \prod_i p_\theta(x_i|x_{<i}, s)$. Thus the training goal is to maximize the log-likelihood of the target sequence,

$$\mathcal{L}_{NLL} = - \sum_{i=0}^{N-1} \log p_\theta(x_i|x_{<i}, s). \quad (4)$$

3.3 Learning Text2motion

Similarly, generating motions from language description can be modeled as autoregressive next-token predictions conditioned on textual inputs. Here we investigate two NMT models as our backbone: attentive GRU and Transformer, and examine our idea of *inverse alignment* on GRU-based model. Since Transformer is typically trained with full teacher force, optimizing the Transformer-based text2motion with inverse alignment is extremely complicated. In other words, every time when generating the density function of next motion token, we need to input the whole history to the Transformer decoder and feed forward. As a result, to sample a complete motion token sequence, the computational (or optimization) graph will be extremely high. Therefore, we specifically introduce the procedure of using GRU based model as an example.

As is shown in Fig. 2 (c), firstly, a bi-directional GRU (i.e., NMT Encoder) models the temporal dependencies in language $x \in \{1, \dots, |\mathcal{V}|\}^N$, and produces sentence feature vector $\mathbf{s} \in \mathbb{R}^{d_l}$ as well as word feature vectors $\mathbf{w} \in \mathbb{R}^{N \times d_l}$, with d_l denoting the dimensionality of hidden vectors. The NMT decoder, modeled as attention-based GRU, processes \mathbf{s} and \mathbf{w} and predicts the probability distribution over discrete motion tokens $\{1, \dots, |\mathcal{B}|\}$ autoregressively. In particular, GRU decoder is initialized by sentence vector \mathbf{s} , and then takes the attention vector \mathbf{w}_{att} together with motion token as input at each time step. The attention vector \mathbf{w}_{att}^t at time t is obtained via

$$\mathbf{Q} = \mathbf{h}_{t-1} \mathbf{W}^Q, \mathbf{K} = \mathbf{w} \mathbf{W}^K, \mathbf{V} = \mathbf{w} \mathbf{W}^V, \quad (5)$$

$$\mathbf{w}_{att}^t = \text{softmax} \left(\frac{\mathbf{Q} \mathbf{K}^T}{\sqrt{d_{att}}} \right) \mathbf{V}, \quad (6)$$

where $\mathbf{h}_{t-1} \in \mathbb{R}^{d_h}$ is previous hidden state in decoder, $\mathbf{W}^K, \mathbf{W}^V \in \mathbb{R}^{d_l \times d_{att}}$ and $\mathbf{W}^Q \in \mathbb{R}^{d_h \times d_{att}}$ are trainable weights with d_h and d_{att} denoting the dimension of hidden unit and attention vector respectively. During generation, motion tokens are sampled from predicted distribution $p_\phi(s_i | s_{<i}, x)$ recursively until the end token (i.e., [EOM]) comes with maximum probability.

Inverse Alignment. Here we re-utilize the motion2text model in Sec. 3.2 to further align the semantics between texts and generated motions. In detail, motion token sequence \hat{s} is sampled from the approximated distribution $p_\phi(s|x)$, which is taken as input to the learned motion2text model and mapped to language tokens x with probability $p_\theta(x|\hat{s})$. Note motion2text model is no longer updated here. However, sampling from discrete distribution is non-differentiable that does not allow the gradients back-propagating to the text2motion encoder and decoder. We instead resort to *Gumbel-Softmax* reparameterization trick [20] to approximate the discrete sampling process. As the temperature τ of Gumbel-Softmax approaches 0, the resulting Gumbel-Softmax distribution becomes identical to the discrete distribution $p_\phi(s_i | s_{<i}, x)$ and the sampled vectors become one-hot.

In summary, the final training objective turns to be

$$\mathcal{L} = - \left(\sum_{i=0}^{K-1} \log p_\phi(s_i | s_{<i}, x) + \sum_{i=0}^{N-1} \log p_\theta(x_i | x_{<i}, \hat{s}) \right). \quad (7)$$

3D pose sequences can finally be obtained by decoding sampled motion tokens \hat{s} using quantization decoder D as described in Sec. 3.1. With discrete motion tokens and autoregressive modeling, variable motion lengths are implicitly modeled by text2motion, that the NMT model particularly learns to predict the end token i.e. [EOM] with maximum probability as signal of termination. Moreover, our proposed approach is easy to train, and does not suffer from the known shortcomings in GAN and VAE such as "mode collapse".

4 Experiments

Extensive experiments are conducted to evaluate our learned motion2text (Sec. 4.3) and text2motion mapping models (Sec. 4.4).

4.1 Datasets

Two 3D human motion-language datasets are considered for evaluation:

- *HumanML3D* [14] is a large 3D human motion dataset that covers a broad range of human actions such as locomotion, sports, and dancing. It consists of 14,616 motions and 44,970 text descriptions. Each motion clip comes with at least 3 descriptions. Motions are re-scaled to 20 frames per second (FPS), resulting in duration ranges from 2 to 10 seconds.

- *KIT Motion-Language* [34] contains 3,911 3D human motion clips and 6,278 text descriptions. For each motion, the corresponding number of text descriptions ranges from one to four. Following [2,11], these pose sequences are all sub-sampled to 12.5 FPS.

Both datasets are split into training, testing and validation sets with ratio of 0.8:0.15:0.05, which are further augmented by mirroring motions and replacing corresponding words in their text descriptions (e.g., 'left'→'right').

4.2 Metrics

Besides traditional measurements, we also manage to evaluate the correspondences between motion and language using deep multimodal features. In particular, we train a simple framework that engages a motion feature extractor and a text feature extractor under contrastive assumption, that learn to produce geometrically closed feature vectors for matched text-motion pairs, and vice versa. Further details are relegated to supplementary file due to limited space.

R-Precision and Multimodal Distance are proposed to gauge how well a text and a motion are semantically aligned. Take the evaluation of motion2text mapping for an example. For each generated description, we take its corresponding motion as well as 31 randomly selected mismatched motions from the test set as a motion pool. With text and motion feature extractors available, Euclidean distances between the description feature and each motion feature in the pool are calculated and ranked. The ground truth entry falling into the top-k ($k=1,2,3$) candidates is regarded as a successful retrieval. Then we count the average accuracy at top-k places, known as *top-k R-precision*. Meanwhile, *multimodal distance* is computed as the average Euclidean distance between text feature of each generated description and motion feature of its corresponding motion in the test set. Computing R-precision and multimodal distance for text2motion mapping is analogically carried out except generated motions and ground truth description are accordingly used.

Overall, an extensive set of metrics including Bleu [29], Rouge [26], Cider [45], BertScore [58], R Precision and multimodal distance are adopted to quantitatively measure the performance of our motion2text mapping. For evaluation of non-deterministic text2motion mapping, we primarily follow [15] which uses Frechet Inception Distance (FID), diversity and multimodality, and our complementary metrics, R precision and multimodal distance. Details of metrics are deferred to be presented in supplementary file.

4.3 Evaluation of Motion-to-text Translation

We adopt RAEs [54] and SeqGAN [13] as our baseline methods; RAEs [54] learns a shared embedding space for language and human motions via two recurrent autoencoders, while SeqGAN [13] combines recurrent sequence-to-sequence model with a discriminator that judges whether a sentence is real or not. We further equip the vanilla RNN model in Seq2Seq [35] with late attention as another

Datasets	Methods	R Precision \uparrow			MM Dist. \downarrow	Bleu@1 \uparrow	Bleu@4 \uparrow	Rouge \uparrow	Cider \uparrow	BertScore \uparrow
		Top 1	Top 2	Top 3						
	Real Desc	0.523	0.725	0.828	2.901	-	-	-	-	-
	RAEs [54]	0.100	0.188	0.261	6.337	33.3	10.2	37.5	22.1	10.7
Human	Seq2Seq(Att)	0.436	0.611	0.706	3.447	51.8	17.9	46.4	58.4	29.1
ML3D	SeqGAN [13]	0.332	0.457	0.532	4.895	47.8	13.5	39.2	50.2	23.4
	Ours w/o MT	<u>0.483</u>	<u>0.678</u>	<u>0.783</u>	3.124	<u>59.5</u>	<u>21.2</u>	<u>47.8</u>	<u>68.3</u>	<u>34.9</u>
	Ours	0.516	0.720	0.823	2.935	61.7	22.3	49.2	72.5	37.8
	Real Desc	0.399	0.618	0.793	2.772	-	-	-	-	-
	RAEs [54]	0.034	0.063	0.106	9.364	30.6	0.10	25.7	8.00	0.40
KIT-	Seq2Seq(Att)	<u>0.293</u>	0.450	0.555	4.455	34.3	9.30	36.3	37.3	5.30
ML	SeqGAN [13]	0.109	0.345	0.425	6.283	3.12	5.20	32.4	29.5	2.20
	Ours w/o MT	0.284	<u>0.466</u>	<u>0.595</u>	3.979	<u>42.8</u>	<u>14.7</u>	<u>39.9</u>	<u>60.1</u>	<u>18.9</u>
	Ours	0.359	0.561	0.668	3.298	46.7	18.4	44.2	79.5	23.0

Table 1. Quantitative evaluation results for **motion-to-text** translation on HumanML3D and KIT-ML test sets. For each metric, the best score is highlighted in **bold**, with the second best highlighted using underscore.

strong baseline (termed as Seq2Seq(Att)). A variant of our method not using motion tokens (ours w/o MT) is also engaged to analyze the role of motion token. Note that grammatical tense and plural of words are neglected in our setting in order to ease the learning process. Descriptions are produced using beam search strategy with size of 2 throughout all experiments.

Quantitative Analysis. Table 1 presents the quantitative evaluation results of motion to language mapping on HumanML3D and KIT-ML test sets. The R precision and multimodal distance of **real** descriptions are provided for reference.

The high R precision of real descriptions also evidences the effectiveness of learned motion & text feature extractors and R precision metric. Overall, our method clearly outperforms all baseline methods over a large margin on all datasets and metrics. RAEs [50] suffers from limited capability on modeling long-term dependencies between 3D motion and language, thus resulting in low R precision and linguistic evaluation scores. This is mitigated by introducing attention mechanism in Seq2Seq(Att) or adversarial learning in SeqGAN, which effectively lifts the top-1 R precision up by more than 20% on HumanML3D and 10% on KIT-ML test sets. By utilizing motion token in our framework (ours), we can observe a obvious jump on both linguistic quality (i.e., Bleu, BertScore) and motion-retrieval precision (i.e., R precision) of generated language descriptions, which is surprisingly approaching the scores of real descriptions.

Qualitative Comparisons. Fig. 4 qualitatively compares the generated descriptions from different methods grounded on the same 3D human motions. RAEs [35] consistently produces descriptions with simple patterns like "is in a", resulting in meaningless linguistic combinations; descriptions from Seq2Seq(Att) and SeqGAN are relatively more complex which however are usually incom-

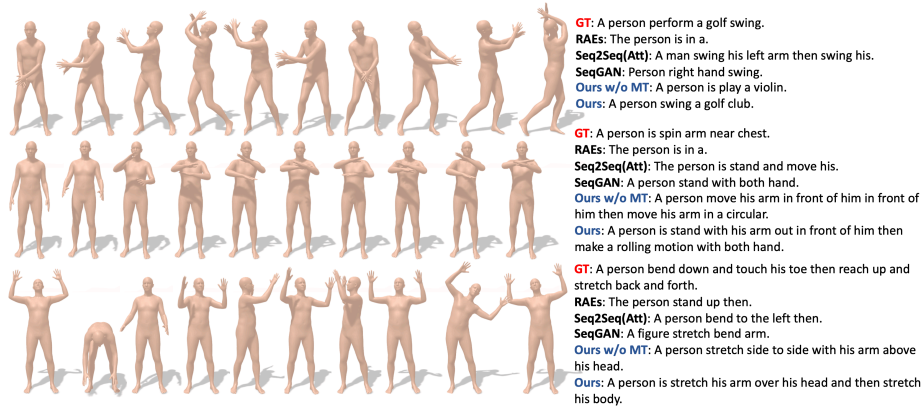


Fig. 4. Examples of motion-to-text translation results from different approaches. Grammatical tense and plural of words are not considered for simplifying learning process. More results are provided in supplementary files.

plete and lack of details. Our approach without motion tokens starts to generate long and complex descriptions. Nonetheless, these descriptions sometimes fail to capture the characteristics of the input 3D motions (e.g, "play a violin"). In contrast, our approach is able to provide fluent and descriptive sentences that accurately depict various aspects of 3D motions, such as body part ("both hand"), action category ("swing", "stretch"), spatial relations ("over head").

User Study. Beside the aforementioned objective evaluations, a crowd-sourced subjective assessment is also conducted on Amazon Mechanical Turk (AMT) involving hundreds of AMT users with *master* recognition. Particularly, descriptions are generated from 100 randomly selected 3D human motions using different methods. For each human motion, the corresponding generated and real descriptions are randomly reordered and shown to 3 AMT users, who are asked to rank their preference over these descriptions based on the accuracy and fluency.




Fig. 5. Statistics of human preference amongst the generated descriptions for given human motions. For each method, a color bar (from blue to red) indicated the the percentage of its preference level (from least to most preferred).

Datasets	Methods	R Precision \uparrow			FID \downarrow	MM Dist \downarrow	Diversity \rightarrow	MModality \uparrow
		Top 1	Top 2	Top 3				
Human ML3D	Real motions	0.511 \pm .003	0.703 \pm .003	0.797 \pm .002	0.002 \pm .000	2.974 \pm .008	9.503 \pm .065	-
	Seq2Seq[25]	0.180 \pm .002	0.300 \pm .002	0.396 \pm .002	11.75 \pm .035	5.529 \pm .007	6.223 \pm .061	-
	Language2Pose[2]	0.246 \pm .002	0.387 \pm .002	0.486 \pm .002	11.02 \pm .046	5.296 \pm .008	7.676 \pm .058	-
	Text2Gesture[5]	0.165 \pm .001	0.267 \pm .002	0.345 \pm .002	5.012 \pm .030	6.030 \pm .008	6.409 \pm .071	-
	Hier[11]	0.301 \pm .002	0.425 \pm .002	0.552 \pm .004	6.532 \pm .024	5.012 \pm .018	8.332 \pm .042	-
	MoCoGAN[42]	0.037 \pm .000	0.072 \pm .001	0.106 \pm .001	94.41 \pm .021	9.643 \pm .006	0.462 \pm .008	0.019 \pm .000
	Dance2Music[23]	0.033 \pm .000	0.065 \pm .001	0.097 \pm .001	66.98 \pm .016	8.116 \pm .006	0.725 \pm .011	0.043 \pm .001
	Ours baseline(T)	0.351 \pm .003	0.521 \pm .003	0.627 \pm .003	1.669 \pm .025	4.046 \pm .018	9.632 \pm .072	4.352 \pm .149
	Ours baseline	0.351 \pm .002	0.526 \pm .002	0.635 \pm .002	1.739 \pm .022	3.965 \pm .010	<u>8.651</u> \pm .083	<u>3.139</u> \pm .083
	Ours	0.424 \pm .003	0.618 \pm .003	0.729 \pm .002	1.501 \pm .017	3.467 \pm .011	8.589 \pm .076	2.424 \pm .093
KIT- ML	Real motions	0.424 \pm .005	0.649 \pm .006	0.779 \pm .006	0.031 \pm .004	2.788 \pm .012	11.08 \pm .097	-
	Seq2Seq[25]	0.103 \pm .003	0.178 \pm .005	0.241 \pm .006	24.86 \pm .348	7.960 \pm .031	6.744 \pm .106	-
	Language2Pose[2]	0.221 \pm .005	0.373 \pm .004	0.483 \pm .005	6.545 \pm .072	5.147 \pm .030	9.073 \pm .100	-
	Text2Gesture[5]	0.156 \pm .004	0.255 \pm .004	0.338 \pm .005	12.12 \pm .183	6.964 \pm .029	9.334 \pm .079	-
	Hier[11]	0.255 \pm .006	0.432 \pm .007	0.531 \pm .007	5.203 \pm .107	4.986 \pm .027	9.563 \pm .072	-
	MoCoGAN[42]	0.022 \pm .002	0.042 \pm .003	0.063 \pm .003	82.69 \pm .242	10.47 \pm .012	3.091 \pm .043	0.250 \pm .009
	Dance2Music[23]	0.031 \pm .002	0.058 \pm .002	0.086 \pm .003	115.4 \pm .240	10.40 \pm .016	0.241 \pm .004	0.062 \pm .002
	Ours baseline(T)	0.260 \pm .005	0.426 \pm .007	0.538 \pm .008	4.628 \pm .126	4.835 \pm .076	<u>12.16</u> \pm .120	<u>4.436</u> \pm .106
	Ours baseline	0.251 \pm .007	0.418 \pm .008	0.535 \pm .007	4.814 \pm .145	4.682 \pm .048	10.13 \pm .117	4.486 \pm .117
	Ours	0.280 \pm .005	0.463 \pm .006	0.587 \pm .005	3.599 \pm .153	4.591 \pm .026	9.473 \pm .117	3.292 \pm .081

Table 2. Quantitative evaluation results for **text-to-motion** mapping on HumanML3D and KIT-ML test sets. All baselines requires fixed motion lengths, and initial poses are further in demand for deterministic methods (first 4 baselines), which are all unnecessary in our approach. \pm indicates 95% confidence interval, and \rightarrow means the closer to the real motion the better. For each metric, the best score is highlighted in **bold**, while the second best is highlighted using underscore.

As shown in Figure.3, our method earns the most appreciation from users over all baselines. In detail, RAEs [54] is the least preferred method, from which 97% descriptions are ranked at the last place; Seq2Seq(Att) and SeqGAN [13] gain comparably more positive feedback from users; while our method without motion tokens comes to the second to the best. This objective study solidly substantiates the capability of our approach toward generating natural as well as motion-aligned language descriptions.

4.4 Evaluation of Text-to-motion Generation

Mapping language to 3D human motions in a non-deterministic fashion is relatively new. Here we compare our method to four state-of-the-art methods: Seq2Seq [25], Language2Pose [2], Text2Gesture [5] and Hier [11]. As with all existing methods, they are unfortunately deterministic methods. Therefore, two stochastic methods in other related fields are adopted here for more fair and in-depth evaluations: MoCoGAN [42] and Dance2Music [23]. MoCoGAN is widely used for conditioned video sequence synthesis, and Dance2Music learns to map sequential audio signals to 2D human dance motions. Proper changes are made to these methods for language-grounded 3D human motion generation. Ours baseline and ours baseline(T) ablates inverse alignment module during training Text2motion and map texts to motions using GRU and Transformer respec-

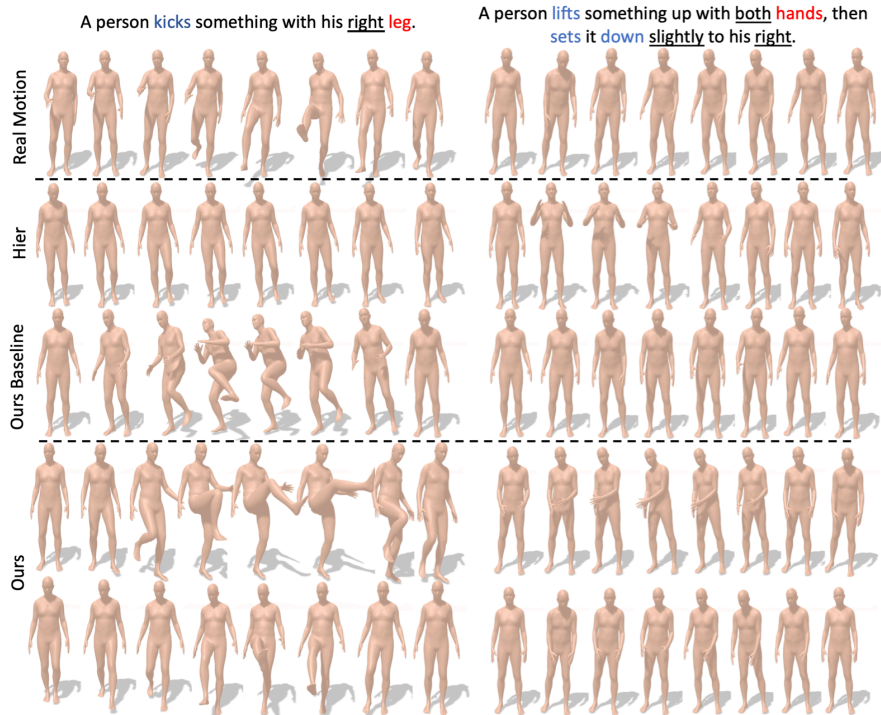


Fig. 6. Visual comparisons of generated motions from the same language descriptions. For each description, we show its corresponding real motion, one motion from Hier [11] (since it’s deterministic) and ours method without inverse alignment, as well as two motions from our method. Key frames of variable-length motion clips are shown. Refer to supplementary files for complete motions and more results.

tively. We repeat each experiment 20 times and report the mean value with 95% statistical confidence interval.

Quantitative Analysis. Table 2 shows the quantitative evaluation results of language grounded 3D human motion generation. We can observe that the motions from non-deterministic baselines, MoCoGAN [42] and Dance2Music [23], suffers from severely low quality and diversity, as reflected by their low R precision and multimodality score. Deterministic baselines such as Seq2Seq [25] and Text2Gesture [5] autoregressively regress human poses from textual input via vanilla sequence-to-sequence RNN and transformer respectively. However, such straightforward approaches find difficulty in maintaining textual semantics during generating human dynamics, which results in low motion-based text retrieval precision and high multimodal distance. Language2Pose [2] and Hier [11] propose to learn a co-embedding space between language and human motions, while Hier [11] go one step forward by incorporating the hierarchical topology of hu-

man skeleton. These have effectively boosted the performance on both datasets. Nevertheless, there still remain a significant gap between the synthetic results and real motions. Our framework of incorporating motion token and NMT model (*ours*, *ours baseline/baseline(T)*) in general achieve better performance, while the inverse alignment strategy greatly benefits this framework (*ours*) with the top-1 and top-3 precision increased by nearly 7% and 10% on HumanML3D.

Visual Comparisons. In Fig. 6, we visually compares the generated motions from our method (*ours*), our method not using inverse alignment (*ours baseline*), and the best performing state of the art, Hier [11]. The corresponding real motions are also provided for reference. Hier [11] could somewhat capture partial concept (e.g., "kick") in descriptions, while the produced motions are unfaithfully in low-mobility. Our method without inverse alignment is capable of generating natural and plausible human motions. It sometimes however still fail to present fine details (e.g., "right leg") from texts. On the contrary, our approach consistently produce visually appealing motions which precisely convey the language concepts in descriptions.

4.5 Limitations and Discussions

Although our proposed TM2T achieves superior performance on both tasks, some limitations and potential remedies can be taken into accounts in future studies. First, the approximation in motion quantization is unfortunately not lossless, which sometimes lead to blurriness and artifacts in local body (e.g., foot sliding). Second, dealing with long and complex descriptions for text2motion is somewhat beyond our capability. This could be possibly solved by using more advanced NMT models. Third, our motion2text model is trained independently with text2motion. Learning these two mapping functions jointly and reciprocally could be another interesting topic.

5 Conclusion

This paper presents TM2T, a general framework that works on the bi-modal mutual mappings between 3D human motions and texts, where motion2text is further reciprocally integrated as a part of text2motion learning through inverse alignment. A new motion representation, motion token, is proposed that compress 3D motions into short sequence of discrete variables. With motion token, neural machine translation networks efficiently build mappings in-between two modalities, that is able to produces accurate descriptions as well as sharp and diverse 3D human motions. Our proposed framework is shown to produce state-of-the-art results on two motion-language dataset in both tasks.

ACKNOWLEDGEMENTS

This work was partly supported by the NSERC Discovery, UAHJIC, and CFI-JELF grants. I also appreciate that the university of Alberta fund me with the Alberta Graduate Excellence Scholarship.

References

1. Adeli, V., Adeli, E., Reid, I., Niebles, J.C., Rezatofghi, H.: Socially and contextually aware human motion and pose forecasting. *IEEE Robotics and Automation Letters* **5**(4), 6033–6040 (2020) [4](#)
2. Ahuja, C., Morency, L.P.: Language2pose: Natural language grounded pose forecasting. In: 2019 International Conference on 3D Vision (3DV). pp. 719–728. *IEEE* (2019) [2](#), [3](#), [4](#), [9](#), [12](#), [13](#)
3. Aliakbarian, S., Saleh, F., Petersson, L., Gould, S., Salzmann, M.: Contextually plausible and diverse 3d human motion prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 11333–11342 (2021) [4](#)
4. Bahdanau, D., Cho, K., Bengio, Y.: Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473* (2014) [3](#)
5. Bhattacharya, U., Rewkowski, N., Banerjee, A., Guhan, P., Bera, A., Manocha, D.: Text2gestures: A transformer-based network for generating emotive body gestures for virtual agents. In: *IEEE Virtual Reality and 3D User Interfaces (VR)*. pp. 1–10. *IEEE* (2021) [12](#), [13](#)
6. Cao, Z., Gao, H., Mangalam, K., Cai, Q.Z., Vo, M., Malik, J.: Long-term human motion prediction with scene context. In: *European Conference on Computer Vision*. pp. 387–404. Springer (2020) [4](#)
7. Corona, E., Pumarola, A., Alenya, G., Moreno-Noguer, F.: Context-aware human motion prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6992–7001 (2020) [4](#)
8. Dubey, S., Olimov, F., Rafique, M.A., Kim, J., Jeon, M.: Label-attention transformer with geometrically coherent objects for image captioning. *arXiv preprint arXiv:2109.07799* (2021) [3](#)
9. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 12873–12883 (2021) [4](#)
10. Gao, J., Wang, S., Wang, S., Ma, S., Gao, W.: Self-critical n-step training for image captioning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6300–6308 (2019) [3](#)
11. Ghosh, A., Cheema, N., Oguz, C., Theobalt, C., Slusallek, P.: Synthesis of compositional animations from textual descriptions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1396–1406 (2021) [2](#), [3](#), [4](#), [9](#), [12](#), [13](#), [14](#)
12. Ging, S., Zolfaghari, M., Pirsiavash, H., Brox, T.: Coot: Cooperative hierarchical transformer for video-text representation learning. *Advances in neural information processing systems* **33**, 22605–22618 (2020) [3](#)
13. Goutsu, Y., Inamura, T.: Linguistic descriptions of human motion with generative adversarial seq2seq learning. In: 2021 IEEE International Conference on Robotics and Automation (ICRA). pp. 4281–4287. *IEEE* (2021) [2](#), [4](#), [9](#), [10](#), [12](#)
14. Guo, C., Zou, S., Zuo, X., Wang, S., Ji, W., Li, X., Cheng, L.: Generating diverse and natural 3d human motions from text. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5152–5161 (2022) [8](#)
15. Guo, C., Zuo, X., Wang, S., Liu, X., Zou, S., Gong, M., Cheng, L.: Action2video: Generating videos of human 3d actions. *International Journal of Computer Vision* pp. 1–31 (2022) [4](#), [9](#)

16. Guo, C., Zuo, X., Wang, S., Zou, S., Sun, Q., Deng, A., Gong, M., Cheng, L.: Action2motion: Conditioned generation of 3d human motions. In: Proceedings of the 28th ACM International Conference on Multimedia. pp. 2021–2029 (2020) [4](#)
17. Guo, L., Liu, J., Yao, P., Li, J., Lu, H.: Mscap: Multi-style image captioning with unpaired stylized text. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4204–4213 (2019) [3](#)
18. Holden, D., Kanoun, O., Perepichka, M., Popa, T.: Learned motion matching. *ACM Transactions on Graphics (TOG)* **39**(4), 53–1 (2020) [4](#)
19. Holden, D., Komura, T., Saito, J.: Phase-functioned neural networks for character control. *ACM Transactions on Graphics (TOG)* **36**(4), 1–13 (2017) [4](#)
20. Jang, E., Gu, S., Poole, B.: Categorical reparameterization with gumbel-softmax. arXiv preprint arXiv:1611.01144 (2016) [8](#)
21. Kojima, A., Tamura, T., Fukunaga, K.: Natural language description of human activities from video images based on concept hierarchy of actions. *International Journal of Computer Vision* **50**(2), 171–184 (2002) [3](#)
22. Kulkarni, G., Premraj, V., Ordonez, V., Dhar, S., Li, S., Choi, Y., Berg, A.C., Berg, T.L.: Babytalk: Understanding and generating simple image descriptions. *IEEE transactions on pattern analysis and machine intelligence* **35**(12), 2891–2903 (2013) [3](#)
23. Lee, H.Y., Yang, X., Liu, M.Y., Wang, T.C., Lu, Y.D., Yang, M.H., Kautz, J.: Dancing to music. *Advances in Neural Information Processing Systems* **32** (2019) [12](#), [13](#)
24. Li, Y., Min, M., Shen, D., Carlson, D., Carin, L.: Video generation from text. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018) [1](#)
25. Lin, A.S., Wu, L., Corona, R., Tai, K., Huang, Q., Mooney, R.J.: Generating animated videos of human activities from natural language descriptions. *Learning* **2018**, 1 (2018) [2](#), [3](#), [4](#), [12](#), [13](#)
26. Lin, C.Y.: Rouge: A package for automatic evaluation of summaries. In: Text summarization branches out. pp. 74–81 (2004) [9](#)
27. Liu, Z., Wu, S., Jin, S., Liu, Q., Lu, S., Zimmermann, R., Cheng, L.: Towards natural and accurate future motion prediction of humans and animals. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10004–10012 (2019) [4](#)
28. Mao, W., Liu, M., Salzmann, M.: Generating smooth pose sequences for diverse human motion prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13309–13318 (2021) [4](#)
29. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th annual meeting of the Association for Computational Linguistics. pp. 311–318 (2002) [9](#)
30. Park, J.S., Rohrbach, M., Darrell, T., Rohrbach, A.: Adversarial inference for multi-sentence video description. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6598–6608 (2019) [3](#)
31. Pavlo, D., Feichtenhofer, C., Auli, M., Grangier, D.: Modeling human motion with quaternion-based neural networks. *International Journal of Computer Vision* **128**(4), 855–872 (2020) [4](#)
32. Peng, J., Liu, D., Xu, S., Li, H.: Generating diverse structure for image inpainting with hierarchical vq-vae. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10775–10784 (2021) [4](#)
33. 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) [4](#)

34. Plappert, M., Mandery, C., Asfour, T.: The kit motion-language dataset. *Big data* **4**(4), 236–252 (2016) [9](#)
35. Plappert, M., Mandery, C., Asfour, T.: Learning a bidirectional mapping between human whole-body motion and natural language using deep recurrent neural networks. *Robotics and Autonomous Systems* **109**, 13–26 (2018) [2](#), [3](#), [4](#), [9](#), [10](#)
36. Qin, Y., Du, J., Zhang, Y., Lu, H.: Look back and predict forward in image captioning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8367–8375 (2019) [3](#)
37. Rakhimov, R., Volkhonskiy, D., Artemov, A., Zorin, D., Burnaev, E.: Latent video transformer. *arXiv preprint arXiv:2006.10704* (2020) [4](#)
38. Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., Sutskever, I.: Zero-shot text-to-image generation. In: *International Conference on Machine Learning*. pp. 8821–8831. PMLR (2021) [4](#)
39. Razavi, A., Van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems* **32** (2019) [4](#)
40. Starke, S., Zhang, H., Komura, T., Saito, J.: Neural state machine for character-scene interactions. *ACM Trans. Graph.* **38**(6), 209–1 (2019) [4](#)
41. Takano, W., Nakamura, Y.: Statistical mutual conversion between whole body motion primitives and linguistic sentences for human motions. *The International Journal of Robotics Research* **34**(10), 1314–1328 (2015) [2](#), [3](#)
42. Tulyakov, S., Liu, M.Y., Yang, X., Kautz, J.: Mocogan: Decomposing motion and content for video generation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1526–1535 (2018) [12](#), [13](#)
43. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. *Advances in neural information processing systems* **30** (2017) [2](#), [4](#), [6](#)
44. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017) [3](#), [7](#)
45. Vedantam, R., Lawrence Zitnick, C., Parikh, D.: Cider: Consensus-based image description evaluation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4566–4575 (2015) [9](#)
46. Venugopalan, S., Rohrbach, M., Donahue, J., Mooney, R., Darrell, T., Saenko, K.: Sequence to sequence-video to text. In: *Proceedings of the IEEE international conference on computer vision*. pp. 4534–4542 (2015) [3](#)
47. Vinyals, O., Toshev, A., Bengio, S., Erhan, D.: Show and tell: A neural image caption generator. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3156–3164 (2015) [1](#), [3](#)
48. Wang, J., Xu, H., Narasimhan, M., Wang, X.: Multi-person 3d motion prediction with multi-range transformers. *Advances in Neural Information Processing Systems* **34** (2021) [4](#)
49. Wang, L., Xiong, Y., Wang, Z., Qiao, Y., Lin, D., Tang, X., Gool, L.V.: Temporal segment networks: Towards good practices for deep action recognition. In: *European conference on computer vision*. pp. 20–36. Springer (2016) [3](#)
50. Wang, Z., Yu, P., Zhao, Y., Zhang, R., Zhou, Y., Yuan, J., Chen, C.: Learning diverse stochastic human-action generators by learning smooth latent transitions. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 34, pp. 12281–12288 (2020) [4](#), [10](#)
51. Xu, C., Govindarajan, L.N., Zhang, Y., Cheng, L.: Lie-x: Depth image based articulated object pose estimation, tracking, and action recognition on lie groups. *International Journal of Computer Vision* **123**(3), 454–478 (2017) [4](#)

52. Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., Bengio, Y.: Show, attend and tell: Neural image caption generation with visual attention. In: International conference on machine learning. pp. 2048–2057. PMLR (2015) [1](#), [3](#)
53. Xu, T., Zhang, P., Huang, Q., Zhang, H., Gan, Z., Huang, X., He, X.: Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1316–1324 (2018) [1](#)
54. Yamada, T., Matsunaga, H., Ogata, T.: Paired recurrent autoencoders for bidirectional translation between robot actions and linguistic descriptions. *IEEE Robotics and Automation Letters* **3**(4), 3441–3448 (2018) [2](#), [3](#), [9](#), [10](#), [12](#)
55. Yu, P., Zhao, Y., Li, C., Yuan, J., Chen, C.: Structure-aware human-action generation. In: European Conference on Computer Vision. pp. 18–34. Springer (2020) [4](#)
56. Yuan, Y., Kitani, K.: Dlow: Diversifying latent flows for diverse human motion prediction. In: European Conference on Computer Vision. pp. 346–364. Springer (2020) [4](#)
57. Zhang, H., Xu, T., Li, H., Zhang, S., Wang, X., Huang, X., Metaxas, D.N.: Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 5907–5915 (2017) [1](#)
58. Zhang, T., Kishore, V., Wu, F., Weinberger, K.Q., Artzi, Y.: Bertscore: Evaluating text generation with bert. arXiv preprint arXiv:1904.09675 (2019) [9](#)