

DeFeeNet: Consecutive 3D Human Motion Prediction with Deviation Feedback

Xiaoning Sun¹, Huaijiang Sun^{1✉}, Bin Li², Dong Wei¹, Weiqing Li¹, Jianfeng Lu¹
¹Nanjing University of Science and Technology, China
²Tianjin AiForward Science and Technology Co., Ltd., China
 {sunxiaoning, sunhuaijiang}@njust.edu.cn, libin@aiforward.com

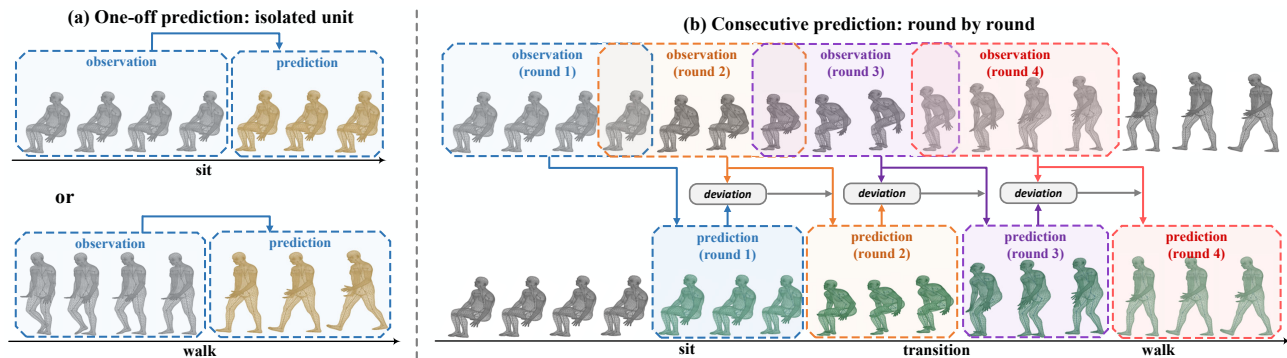


Figure 1. (a) Current human motion prediction task: a one-off “observe then predict” process within the isolated unit. (b) Consecutive human motion prediction with deviation feedback: covering multiple “observe then predict” rounds, where the prediction deviation (i.e., mistake) in previous round could constrain the model to make better prediction in the following round.

Abstract

Let us rethink the real-world scenarios that require human motion prediction techniques, such as human-robot collaboration. Current works simplify the task of predicting human motions into a one-off process of forecasting a short future sequence (usually no longer than 1 second) based on a historical observed one. However, such simplification may fail to meet practical needs due to the neglect of the fact that motion prediction in real applications is not an isolated “observe then predict” unit, but a consecutive process composed of many rounds of such unit, semi-overlapped along the entire sequence. As time goes on, the predicted part of previous round has its corresponding ground truth observable in the new round, but their deviation in-between is neither exploited nor able to be captured by existing isolated learning fashion. In this paper, we propose DeFeeNet, a simple yet effective network that can be added on existing one-off prediction models to realize deviation perception and feedback when applied to consecutive motion prediction task. At each prediction round, the deviation generated by previous unit is first encoded by our DeFeeNet, and then incorporated into the existing predictor to enable a deviation-aware prediction manner, which, for the first time,

allows for information transmit across adjacent prediction units. We design two versions of DeFeeNet as MLP-based and GRU-based, respectively. On Human3.6M and more complicated BABEL, experimental results indicate that our proposed network improves consecutive human motion prediction performance regardless of the basic model.

1. Introduction

In the age of intelligence, humans have to share a common space with robots, machines or autonomous systems for a sustained period of time, such as in human-motion collaboration [33], motion tracking [10] and autonomous driving [38] scenarios. To satisfy human needs while keeping their safety, deterministic human motion prediction, which is aimed at forecasting future human pose sequence given historical observations, has become a research hotspot and already formed a relatively complete implement procedure. As human poses may become less predictable over time, current works abstract the practical needs into a simplified task of learning to “observe a few frames and then predict the following ones”, with the prediction length mostly set to ≤ 1 second [4, 7–9, 16, 20, 21, 24, 27, 28, 34, 40].

Essentially, such simplified prediction task can be re-

garded as a one-off acquired target limited within the short, isolated “observe then predict” unit. Nevertheless, this unit is actually not applicable in reality that requires consecutive observation and prediction on humans during the long period of human-robot/machine coexistence. Though, intuitively, sliding such unit round by round along the time may roughly satisfy the need for consecutive prediction, one neglected fact is that each round of prediction unit is arranged in a *semi-overlapped* structure (see Figure 1). As time goes on, what was predicted *before* has its corresponding ground truth observable *now*, but their deviation in-between (i.e., mistake) remains unexplored.

Our key insight lies in that robots/machines should be able to detect and learn from the mistakes they have made. Due to the inherent continuity and consistency of human motion data, such information would be very powerful to improve future prediction accuracy. The semi-overlapped, multi-round unit structure offers a natural way to transmit deviation feedback across adjacent prediction units, which, however, are unable to be realized by current one-off unit prediction strategy.

Based on this situation, in this paper, we propose DeFeeNet, a simple yet effective network which can be added on existing one-off human motion prediction models to implement **Deviation** perception and **Feedback** when applied to consecutive prediction. By mining the element “deviation” that neglected in previous works, we introduce inductive bias, where DeFeeNet can learn to derive certain “patterns” from past deviations and then constrain the model to make better predictions. To be specific, our DeFeeNet is constructed in two versions: MLP-based version containing one temporal-mixing MLP and one spatial-mixing MLP; GRU-based version containing one-layer GRU with fully-connected layers only for dimension alignment. At each prediction round, DeFeeNet serves as an additional branch inserted into the existing predictor, which encodes the deviation of the previous unit into latent representation and then transmits it into the main pipeline. In summary, our contribution are as follows:

- We mine the element “deviation” that neglected in existing human motion prediction works, extending current within-unit research horizon to cross-unit, which, for the first time, allows for information transmit across adjacent units.
- We propose DeFeeNet, a simple yet effective network that can be added on existing motion prediction models to enable consecutive prediction (i) with flexible round number and (ii) in a deviation-aware manner. It can learn to derive certain patterns from past mistakes and constrain the model to make better predictions.
- Our DeFeeNet is agnostic to its basic models, and capable of yielding stably improved performance on Hu-

man3.6M [15], as well as a more recent and challenging dataset BABEL [32] that contains samples with different actions and their transitions.

2. Related Work

2.1. Human Motion Prediction

Model Design Paradigm. Two mainstreams of model design in human motion prediction are based on sequential networks [7, 9, 16, 28, 30] and feed-forward networks [1, 5, 8, 20, 21, 27, 33, 34, 40]. For the former one, RNNs [7, 16], LSTMs [9, 28] and GRUs [28, 30] are employed to extract temporal features of motion sequences, while for the latter, CNNs [20], GNNs [21], GCNs [8, 27, 33–35, 40] and transformers [1, 5] learn spatial dependencies or spatio-temporal information to better depict human skeletal structure and joint spatio-temporal connections. All these current designs simplify human motion prediction as a short, isolated “observe then predict” *unit*, with predicted length mostly ≤ 1 second, and with no consideration of its context, which, therefore, could not satisfy practical needs like human-robot collaboration scenario that requires consecutive observation and prediction on human behaviors.

Sample Selection Protocol. Currently in this field, samples are drawn in a uniform way of cutting small pieces from long motion sequences with known action categories. Each sample piece (i.e., the *unit*) is composed of the first several frames as observation and the others as prediction target, where poses in one piece all belong to the same action category. Such short-length yet single-action sampling convention, therefore, determines that current trained predictors are unable to handle real-world prediction task that covers multiple categories of action and action-changing periods.

Motivated by the above, we reformulate human motion prediction task from the consecutive perspective, and argue that previous prediction deviation should be detected and learned by robots/machines to help the following prediction. We accordingly present a resampling method by drawing multiple yet flexible rounds of *semi-overlapped* unit as *one* sample with a relatively long horizon (Section 3.4 for details), which enables (i) our deviation-aware consecutive prediction as well as (ii) the first attempt to consider action-changing in deterministic human motion prediction. Note that in this paper we only refer to *deterministic* motion prediction that aims to predict only *one* future sequence. Stochastic prediction (predicting multiple possible sequences) [2, 25, 26, 37, 39] or motion synthesis (generating sequences without observations) [11, 31] are not within the scope of our discussion.

2.2. Deviation-Aware Prediction

Other seemingly similar ideas on deviation-aware prediction include Local-Behavior-Aware (LBA) framework in

trajectory prediction [41] and Residual Correction framework in node regression [17], node classification [14] and traffic prediction [18, 22]. For the former, LBA [41] refers to involving a collection of all historical trajectories at the object’s current location to help narrow down the search space of future trajectories. Though the “historical trajectories of other objects” notion somewhat resembles our “previous prediction round”, it is entirely unrelated to the use of deviation. For the latter, residual is defined as the difference between GT and the prediction just like our deviation, but all these works first generate an “original” prediction, which is then *directly* added with the estimated residual, so that the original prediction could be corrected to some extent (note that the *estimated* residual is calculated based on residuals of training vertices [14, 17] or based on historical residuals [18, 22]). In other words, their residual correction serves as a post-processing step, and the residual expression is completely independent of the original pipeline. In contrast, we adopt the “encoder-separated while decoder-unified” strategy, which allows for separate coding on previous deviation while joint decoding on both information of current observation and previous deviation. Therefore, our prediction is naturally improved rather than corrected by post-processing, and that is why we regard it as deviation-aware prediction but not prediction correction. As human motion data is more *granular* yet delicate that cannot be corrected by simply adding a residual/deviation vector, our joint decoding strategy could leverage the decoding power of the original predictor, enabling deviation-aware prediction only with a lightweight network branch inserted.

3. Proposed Method

In most cases, 3D human structure is represented by pure pose parameters based on human joint coordinates or SMPL model [23] which depicts 3D human mesh with both pose and shape. According to task demand, we follow [26, 31] to discard the shape term and only predict the pose term. Currently, human motion prediction [4, 27, 34] aims to predict a future sequence $\hat{\mathbf{Y}} = [\hat{\mathbf{y}}_1, \hat{\mathbf{y}}_2, \dots, \hat{\mathbf{y}}_T]$ with T frames based on the observed N -frame $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]$ sequence, where each frame (i.e., pose) is represented by $\hat{\mathbf{y}}_i \in \mathbb{R}^K$ and $\mathbf{x}_i \in \mathbb{R}^K$, respectively. We define such process of predicting T frames with the N -frame observation as *unit*. Research within this isolated unit, however, totally neglects the importance of “prediction mistake” element to consecutive motion prediction, and is unable to realize deviation perception and feedback to improve prediction accuracy.

3.1. Reformulation from Consecutive Perspective

We introduce our reformulated human motion prediction task. Consecutive prediction involves multiple rounds of unit with $N + T$ frames semi-overlapped along the entire sequence. Suppose a sequence $\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_L]$ with

length L long enough, we define the first round as $\mathbf{R}_1 = \mathbf{s}_{1:N+T}$, and every following round is T frames further, i.e., the r -th round $\mathbf{R}_r = \mathbf{s}_{1+(r-1)T:N+T+(r-1)T}$, where $r \geq 1$, and in this paper we set $N \geq T$. Therefore, if we define all the motion rounds $\mathbf{R} = [\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_r, \dots]$ and rewrite $\mathbf{R}_r = [\mathbf{X}_r, \mathbf{Y}_r] = [\mathbf{x}_{r,1}, \dots, \mathbf{x}_{r,N}, \mathbf{y}_{r,1}, \dots, \mathbf{y}_{r,T}]$, it can be deduced that $\mathbf{x}_{r,N-T:N} = \mathbf{y}_{r-1,1:T}$, and the deviation generated in round \mathbf{R}_{r-1} is expressed as $\mathbf{D}_{r-1} = f_d(\mathbf{x}_{r,N-T:N}, \hat{\mathbf{y}}_{r-1,1:T})$, to be used to improve the prediction accuracy of round \mathbf{R}_r . When $r = 1$, \mathbf{D}_0 is empty as no deviation can be used. Basically, we regard the task of consecutive motion prediction as implement \mathbf{R} round by round.

3.2. Deviation-Aware Prediction

To express deviation $\mathbf{D}_{r-1} = f_d(\mathbf{x}_{r,N-T:N}, \hat{\mathbf{y}}_{r-1,1:T})$ mentioned in Section 3.1 in detail, we draw inspiration from [4, 12] that design their prediction loss function based on joint displacement of two adjacent poses (i.e., velocity) rather than joint position. Instead of directly subtracting between $\mathbf{x}_{r,N-T:N}$ and $\hat{\mathbf{y}}_{r-1,1:T}$, we first calculate their velocities as $v(\mathbf{x}_{r,N-T:N}), v(\hat{\mathbf{y}}_{r-1,1:T}) \in \mathbb{R}^{K \times (T-1)}$ and then obtain:

$$\begin{aligned} \mathbf{D}_{r-1} &= f_d(\mathbf{x}_{r,N-T:N}, \hat{\mathbf{y}}_{r-1,1:T}) \\ &= v(\mathbf{x}_{r,N-T:N}) - v(\hat{\mathbf{y}}_{r-1,1:T}) \in \mathbb{R}^{K \times (T-1)}. \end{aligned} \quad (1)$$

In this way, compared to existing works that focus on learning a prediction function $\mathbf{Y} = \varphi_o(\mathbf{X})$ between the observation \mathbf{X} and the target \mathbf{Y} in the isolated unit, our deviation-aware consecutive prediction aims to find a function for every prediction round with deviation feedback incorporated:

$$\mathbf{Y}_r = \varphi(\mathbf{X}_r, \mathbf{D}_{r-1}), \quad (2)$$

where $\varphi : \mathbb{R}^{K \times N} \times \mathbb{R}^{K \times (T-1)} \mapsto \mathbb{R}^{K \times T}$. Note that $\mathbf{D}_0 = \mathbf{0}$ when $r = 1$.

In Figure 2, we provide an overview of how DeFeeNet is embedded into the existing prediction framework and helps implement deviation-aware motion prediction round by round during a relatively long period. At round r , we feed the current round of observation $\mathbf{X}_r = \mathbf{x}_{r,1:N}$ into the main pipeline (i.e., the existing predictor). Meanwhile, our DeFeeNet could detect the prediction deviation that *just happened* as the prediction of round $r - 1$ (i.e., $\hat{\mathbf{Y}}_{r-1} = \hat{\mathbf{y}}_{r-1,1:T}$) has its ground truth observable now. Therefore, with this deviation \mathbf{D}_{r-1} coded and incorporated, the current prediction $\hat{\mathbf{Y}}_r = \hat{\mathbf{y}}_{r,1:T}$ could be naturally improved.

3.3. Network Architecture

We build two versions of DeFeeNet based on MLP and GRU (see Figure 3), both of which are simple yet effective, as we do not encourage DeFeeNet to increase the computation burden too much on original prediction models.




Figure 2. (a) Existing baselines that simplify human motion prediction as a one-off process of predicting a T -pose sequence based on N -pose observation (i.e., the isolated prediction unit). (b) DeFeeNet allows for human motion prediction *round by round* with a deviation-aware manner. Round r is marked in mint green while $r - 1$ in orange. During the prediction of round r , our DeFeeNet detects the deviation between current observation and previous round of prediction, to realize deviation feedback across adjacent units.

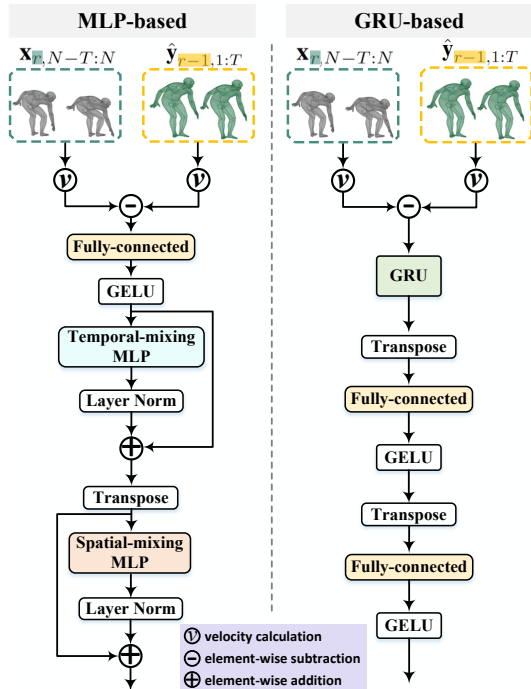


Figure 3. Architectures of MLP-based and GRU-based DeFeeNet.

Inspired by the success of MLP-Mixer structure [36] in computer vision and its variants in downstream human motion modeling tasks [4, 6, 12], we propose an MLP-based version of our DeFeeNet, but is more lightweight yet effective enough. Compared to using multiple MLP-mixer blocks in [4, 6, 12, 36], our design discards block stacking and only contains one temporal-mixing MLP, one spatial-mixing MLP, and skip-connections (see Figure 3 left). Both MLP-mixer module are composed of two fully-connected layers, GELU activation function [13], and layer normalization operation [3].

For the GRU-based version, as is shown in Figure 3 right, we employ a one-layer GRU attached with one fully-connected layer in temporal dimension and one in spatial dimension for dimension alignment, to be inserted into the

original pipeline for following calculation.

To be specific, the original predictor is first fed with the observed segment $\mathbf{X}_1 = \mathbf{x}_{1,1:N}$ in \mathbf{R}_1 , and obtain its first round of prediction $\hat{\mathbf{Y}}_1 = \hat{\mathbf{y}}_{1,1:T}$ at timestep T , with DeFeeNet nonactivated and no deviation information involved. However, during the following rounds where \mathbf{X}_r becomes observable, the deviation between $v(\hat{\mathbf{y}}_{r-1,1:T})$ and $v(\mathbf{x}_{r,N-T:N})$ appears, and is fed into our DeFeeNet to capture its temporal and spatial information to be added into the main pipeline, enabling deviation feedback across adjacent prediction units.

3.4. Two-Round Training & Multi-Round Testing

The goal of our DeFeeNet is to improve the original prediction accuracy by consecutive perception and feedback on prediction deviation, as well as maintaining this priority in a sustained period of time. As the semi-overlapped structure offers us a natural way to learn prediction deviation, we accordingly present our resampling method and loss function. **Resampling.** Compared to existing prediction models that regard the “observe then predict” unit as a sample, we draw two rounds of unit as one training sample and multiple rounds of unit as one testing sample, i.e., $\mathbf{R}_{\text{train}} = [\mathbf{R}_1, \mathbf{R}_2]$ and $\mathbf{R}_{\text{test}} = [\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_{\max(r)}]$. For training phase, we just let DeFeeNet learn to produce useful information according to the previous adjacent deviation it percepts, so only two rounds of prediction (containing one overlapped region) are sufficient. While for testing phase, samples with multiple yet flexible rounds of prediction unit are necessary, as our ultimate goal is to evaluate whether DeFeeNet is able to generate stable improvement on the consecutive observation and prediction with unknown duration in advance.

Loss Function. Following current works [4, 8, 12, 27, 34], our loss is based on Mean Per Joint Position Error (MPJPE) [15], but differently, we calculate the loss on both rounds of prediction. In line with notation in Section 3.1, we express the loss at round r ($r = 1, 2$) as:

$$\mathcal{L}_r = \frac{1}{T} \sum_{t=1}^T \|\hat{\mathbf{y}}_{r,t} - \mathbf{y}_{r,t}\|_2, \quad (3)$$

frame num.	walking				eating				smoking				discussion				directions			
	2	4	8	10	2	4	8	10	2	4	8	10	2	4	8	10	2	4	8	10
LTD-GCN [27]	12.3	23.0	39.8	46.1	8.4	16.9	33.2	40.7	8.0	16.2	31.9	38.9	12.5	27.4	58.5	71.7	9.0	19.9	43.4	53.7
LTD-DeFee(MLP)-r1	11.7	21.2	36.4	43.5	8.1	16.2	31.9	39.8	7.7	15.3	29.9	36.8	12.7	26.8	56.4	69.4	9.1	19.3	41.7	52.0
LTD-DeFee(MLP)-r2	10.2	19.9	35.5	42.5	6.7	14.8	30.7	38.3	6.5	14.0	28.8	35.6	10.0	23.8	54.5	68.1	6.9	16.7	39.6	50.2
LTD-DeFee(GRU)-r1	11.8	22.2	36.8	44.0	8.2	16.9	32.2	40.2	7.8	15.8	30.2	37.5	12.5	26.7	57.1	70.4	8.8	19.5	42.4	52.8
LTD-DeFee(GRU)-r2	10.4	20.0	34.7	42.2	7.0	15.2	31.4	38.4	6.8	14.5	29.0	35.8	11.1	25.4	55.8	68.2	7.0	17.0	40.0	50.9
frame num.	greeting				phoning				posing				purchases				sitting			
	2	4	8	10	2	4	8	10	2	4	8	10	2	4	8	10	2	4	8	10
STS-GCN [34]	18.7	34.9	71.6	86.4	13.7	22.4	43.6	53.8	16.4	30.4	67.6	84.7	19.1	35.8	70.2	83.1	15.2	25.1	49.8	60.8
STS-DeFee(MLP)-r1	18.4	34.2	71.2	85.4	13.5	21.9	42.1	52.5	16.6	30.4	66.2	83.2	19.5	35.5	69.4	82.3	15.0	24.8	49.6	59.8
STS-DeFee(MLP)-r2	16.4	32.8	68.8	82.5	11.4	19.8	40.8	49.7	14.8	28.3	64.8	80.5	16.6	32.5	67.6	80.6	14.0	23.3	47.5	58.7
STS-DeFee(GRU)-r1	18.5	34.5	71.4	85.8	13.5	22.1	42.5	53.0	16.2	30.2	66.7	83.5	18.9	35.2	69.4	82.4	15.1	24.9	49.5	59.6
STS-DeFee(GRU)-r2	16.8	33.0	68.5	83.2	11.6	19.9	41.0	50.1	14.7	28.3	65.0	81.1	16.8	32.7	67.9	80.3	14.2	23.6	47.7	58.7
frame num.	sittingdown				takingphoto				waiting				walkingdog				walkingtogether			
	2	4	8	10	2	4	8	10	2	4	8	10	2	4	8	10	2	4	8	10
MotionMixer [4]	12.0	31.4	64.4	74.5	9.0	18.9	41.0	51.6	10.2	21.1	45.2	56.4	20.5	42.8	75.6	87.8	10.5	20.6	38.7	43.5
MotMix-DeFee(MLP)-r1	12.0	30.6	63.7	72.6	9.1	18.2	39.7	50.4	10.4	20.5	44.1	55.6	19.5	42.4	74.4	86.9	10.2	19.8	37.8	42.7
MotMix-DeFee(MLP)-r2	9.8	29.1	61.8	70.2	7.8	16.9	37.1	47.7	9.3	19.5	42.0	53.4	17.3	40.9	72.8	84.4	8.3	19.1	35.9	41.5
MotMix-DeFee(GRU)-r1	11.7	30.8	63.6	73.7	8.8	18.5	40.2	50.8	10.1	20.8	44.5	55.4	19.5	42.1	74.4	87.1	10.4	19.8	38.0	43.1
MotMix-DeFee(GRU)-r2	10.1	29.4	62.0	70.8	7.8	16.9	38.3	47.9	9.6	19.8	42.3	53.6	17.6	41.1	72.7	84.9	8.8	19.0	36.1	41.8

Table 1. Top to bottom: Prediction errors produced by the original baselines (isolated), baselines with MLP-based DeFeeNet inserted (round 1 and 2), and with GRU-based DeFeeNet inserted (round 1 and 2). Values in bold indicate lower errors and prove the deviation feedback is valid. For each baseline, we present the performance on 5 actions out of 15 in Human3.6M. We present prediction errors at frame 2, 4, 8, and 10.

where $\hat{\mathbf{y}}_{r,t}$ denotes the predicted pose at frame t at round r , and $\mathbf{y}_{r,t}$ as the corresponding ground truth. Note that, for fair comparison, Eq. 3 will have some slight changes in the specific details when our DeFeeNet is added on different baseline models and evaluated on different datasets (may refer to the supplemental material).

Then the total loss can be calculated as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_1 + \mathcal{L}_2. \quad (4)$$

In addition to supervising both rounds of prediction as accurate as possible, \mathcal{L}_1 could further ensure that the inclusion of DeFeeNet does not encumber the main pipeline too much (as the model parameters have been trained to accommodate the existence of DeFeeNet, but no deviation is involved in round 1), while \mathcal{L}_2 could keep the deviation feedback valid.

4. Experiments

In this section, we evaluate our DeFeeNet on the task of consecutive human motion prediction. Two datasets (Human3.6M [15] and the more recent and challenging BABEL [32]) are preprocessed to meet the task requirements, and we conduct experiments on three different motion prediction baselines with DeFeeNet inserted. We additionally provide ablation studies for further analysis.

4.1. Datasets

Human3.6M [15] is one of the most widely used datasets for human motion modeling, which includes seven actors performing actions under 15 categories, such as *walking*,

eating, *smoking*, and *discussion*, with downsampling rate 25 fps to generate the sequence data. Each human skeletal pose is represented by 32 joints, and we follow existing works [4, 8, 12, 27, 34, 40] that evaluate on 22 of these joints. Currently, they aim to produce 10 or 25 frames (i.e., 400 ms short-term or 1000 ms long-term prediction) based on 10 frames of observation, within the isolated “unit”.

For consecutive prediction task, we use our resampling method (Section 3.4) for further preprocessing. For fair comparison, we also set each prediction round with 10 frames observed and the following 10 to be predicted. Training samples are composed of two adjacent rounds, wherein the second round is 10 frames further compared to the first one. Testing samples are first arranged with 2 rounds to prove the deviation feedback between adjacent units is valid, then we arrange 10 rounds to evaluate the stability of DeFeeNet along a relatively long horizon.

BABEL [32] is a recent proposed dataset with language labels on actions. Different from other datasets for human motion prediction wherein actions contained in each sequence are of the same category, BABEL allows for sequence with multiple kinds of actions. Currently, the most suitable preprocessing method for our task is [26] that exports two dataset files: (a) single-action sequences under 20 action labels (such as *stand*, *walk*, *step* and *stretch*); (b) sequences with two categories of actions and their in-between transitions that contain these 20 actions, both of which are of 30 fps, with only pose parameters for evaluation, while SMPL shape parameters discarded (may refer to the link provided in [26]).

frame num.	LTD-GCN [27]				LTD-GCN [27]		STS-GCN [34]		MotionMixer [4]	
	2	4	8	10	avg		avg		avg	
isolated	12.69	26.06	52.28	63.53	38.64		41.16		35.52	
	-DeFee(MLP)				-D(MLP)	-D(GRU)	-D(MLP)	-D(GRU)	-D(MLP)	-D(GRU)
r1	12.59	25.16	50.02	61.47	37.31	37.65	40.76	40.89	34.12	34.35
r2	10.30	22.60	48.17	59.79	35.21	35.80	38.15	38.77	32.38	32.69
r3	10.43	22.74	48.44	60.14	35.44	35.67	38.33	38.67	32.16	32.74
r4	10.39	22.69	48.31	59.95	35.33	35.66	38.46	38.72	32.05	32.75
r5	10.38	22.68	48.27	59.91	35.31	35.91	38.45	38.90	32.17	32.59
r6	10.38	22.67	48.25	59.87	35.30	36.10	38.37	39.05	32.34	32.69
r7	10.37	22.66	48.24	59.88	35.29	35.88	38.41	38.95	32.33	32.66
r8	10.39	22.68	48.27	59.89	35.31	35.66	38.45	38.89	32.47	32.77
r9	10.40	22.71	48.30	59.92	35.33	35.75	38.38	38.93	32.41	32.71
r10	10.41	22.72	48.33	59.96	35.35	35.70	38.49	38.78	32.44	32.75

Table 2. Consecutive 10 rounds of prediction errors on Human3.6M. Values in bold from round 2 to 10 indicate that deviation-aware prediction *stably* yields improved performance. Detailed errors of LTD-DeFee(MLP) at frame 2, 4, 8, 10 are on the left. Average errors (avg) of these four testpoints produced by other baselines with DeFeeNet (abbreviated as *D*) are on the right.

As no experiments of deterministic motion prediction has been conducted on BABEL, we first cut isolated unit samples from the above sequence files to train the three baselines, with observed and predicted frames both set to 10. Then, similar to Human3.6M, we rearrange our training samples with 2 rounds while testing samples with 2 rounds for effectiveness validation and 5 rounds for stability validation. Moreover, we remove the actions with less than 200 sequence samples and leave samples with 11 action categories for evaluation.

Evaluation Metric. Following [4, 8, 12, 24, 34, 40], we evaluate our network by Mean Per Joint Position Error (MPJPE) on 3D human joint coordinates, which calculates the average L_2 -norm on discrepancy of different joints between prediction and the corresponding ground truth.

4.2. Baselines

Our DeFeeNet is agnostic to its basic prediction models, and can be added on many existing networks to realize consecutive human motion prediction. Here we introduce three baselines for our experiments.

LTD-GCN [27] is a GCN-based model composed of 12 GCN blocks depicting the spatial correlations of different joint dimensions with residual connections. Discrete Cosine Transformation (DCT) is used to extract temporal features. **STS-GCN [34]** designs the single-graph structure which allows for space-time cross-talk, which factorizes the graph into separable space and time matrices to fully learn the joint-joint and time-time interactions.

MotionMixer [4] proposes efficient motion prediction model that only adopts MLPs by sequentially mixing both spatial and temporal dependencies, and with squeeze-and-excitation blocks enhanced.

4.3. Implementation Details

We conduct our experiments under Pytorch [29] framework with Adam optimizer [19] on a single NVIDIA RTX

frame num.	w/o transi				w/ transi			
	3	6	8	10	3	6	8	10
LTD-GCN [27]	0.2111	0.3828	0.4751	0.5617	0.2179	0.3961	0.4940	0.5802
LTD-D(MLP)-r1	0.2032	0.3781	0.4689	0.5559	0.2054	0.3805	0.4811	0.5681
LTD-D(MLP)-r2	0.1971	0.3683	0.4566	0.5464	0.1988	0.3699	0.4614	0.5481
LTD-D(GRU)-r1	0.2009	0.3778	0.4762	0.5601	0.2022	0.3792	0.4799	0.5614
LTD-D(GRU)-r2	0.1921	0.3637	0.4583	0.5470	0.1945	0.3708	0.4607	0.5434
STS-GCN [34]	0.2404	0.4283	0.5330	0.6306	0.2394	0.4323	0.5536	0.6438
STS-D(MLP)-r1	0.2277	0.4063	0.5164	0.6063	0.2274	0.4064	0.5061	0.5960
STS-D(MLP)-r2	0.2101	0.3809	0.4933	0.5745	0.2147	0.3855	0.4839	0.5760
STS-D(GRU)-r1	0.2291	0.4138	0.5119	0.6140	0.2279	0.4126	0.5107	0.6028
STS-D(GRU)-r2	0.2143	0.3903	0.4917	0.5885	0.2163	0.3942	0.4927	0.5866
MotionMixer [4]	0.1975	0.3688	0.4401	0.5172	0.1968	0.3754	0.4591	0.5386
MotMix-D(MLP)-r1	0.1926	0.3615	0.4372	0.5119	0.1905	0.3644	0.4426	0.5178
MotMix-D(MLP)-r2	0.1763	0.3443	0.4140	0.4961	0.1791	0.3511	0.4268	0.4989
MotMix-D(GRU)-r1	0.1884	0.3617	0.4407	0.5136	0.1899	0.3685	0.4422	0.5151
MotMix-D(GRU)-r2	0.1743	0.3512	0.4223	0.4881	0.1768	0.3541	0.4248	0.4896

Table 3. Comparisons of prediction errors between the original baselines (isolated) and baselines with MLP/GRU-based DeFeeNet (round 1 and round 2). We conduct experiments on BABEL with (a) only single-action samples (i.e., w/o transi) and (b) two-action samples with transitions between different actions (i.e., w/ transi). Lower errors in round 2 (marked in bold) show that the deviation produced previously is effective to improve the subsequent prediction accuracy. We abbreviate our DeFeeNet as *D*.

2080Ti. In MLP-based version, the hidden layer size of temporal-mixing and spatial-mixing is set to the same as the corresponding dimensions of the basic pipeline. In GRU-based version, the hidden size is set to 256 regardless of the basic model, while fully-connected layers are used for dimension alignment. Details are provided in the supplemental material. For LTD-GCN, STS-GCN, and MotionMixer, we jointly train each baseline with our DeFeeNet with learning rate as 0.0005, 0.01 and 0.01, respectively. We train 50 epochs on Human3.6M and 100 on BABEL.

4.4. Results on Human3.6M

Table 1 presents the prediction errors of models with and without DeFeeNet inserted on Human3.6M dataset. Without DeFeeNet, current baselines observe and predict human motions within the isolated short unit (i.e., length of 10 frames is 400 ms). However, we jointly train two rounds of prediction process, wherein the prediction deviation generated in round 1 (r1) is detected and learned by our inserted DeFeeNet, in order to assist the original pipeline to predict with this previous deviation as feedback to improve the performance of round 2 (highlighted in purple background color). Values in bold indicate lower errors, which prove that this deviation information between adjacent prediction rounds is valid for prediction improvement.

To demonstrate the stability of DeFeeNet for consecutive human motion prediction, we further show the results of consecutive 10 rounds of prediction in Table 2. From errors of round 2 to 10, we observe that, compared to baselines which treat motion prediction as a task within the isolated unit, our deviation-aware prediction manner could *stably* produce improved results during a sustained period of

frame num.	w/o transi						w/ transi									
	LTD-GCN [27]		STS-GCN [34]		MotionMixer [4]		LTD-GCN [27]		STS-GCN [34]		MotionMixer [4]		STS-GCN [34]			
	avg		avg		avg		avg		avg		avg		3	6	8	10
isolated	0.4077		0.4581		0.3809		0.4221		0.4673		0.3925		0.2394	0.4323	0.5536	0.6438
	-D(MLP)	-D(GRU)	-D(MLP)	-D(GRU)	-D(MLP)	-D(GRU)	-D(MLP)	-D(GRU)	-D(MLP)	-D(GRU)	-D(MLP)	-D(GRU)	-DeFee(MLP)			
r1	0.4015	0.4038	0.4392	0.4421	0.3758	0.3761	0.4089	0.4055	0.4413	0.4385	0.3785	0.3766	0.2372	0.4112	0.5139	0.6029
r2	0.3842	0.3901	0.4185	0.4239	0.3590	0.3589	0.3982	0.3963	0.4270	0.4265	0.3547	0.3599	0.2342	0.3965	0.4949	0.5824
r3	0.3810	0.3890	0.4122	0.4203	0.3611	0.3566	0.3815	0.3802	0.4205	0.4189	0.3502	0.3582	0.2332	0.3920	0.4845	0.5723
r4	0.3808	0.3826	0.4132	0.4277	0.3582	0.3540	0.3877	0.3845	0.4223	0.4192	0.3610	0.3570	0.2316	0.3904	0.4897	0.5777
r5	0.3855	0.3865	0.4166	0.4359	0.3545	0.3578	0.3840	0.3866	0.4290	0.4133	0.3579	0.3613	0.2330	0.3933	0.4952	0.5946

Table 4. The consecutive 5 rounds of prediction errors on BABEL (w/o or w/ transi). The detailed errors of STS-DeFee(MLP) at frame 3, 6, 8, 10 are presented on the right. Average errors (avg) of these four testpoints produced by other baselines with DeFeeNet (abbreviated as *D*) are also provided. From round 2 to 5, our method could *stably* yield lower prediction errors than isolated baselines (marked in bold).




Figure 4. Visualized comparisons of consecutive motion prediction w/o or w/ deviation feedback on a sample *walking together* from Human3.6M. Due to the limited space, we only provide 3 prediction rounds (r6, r7, r8) out of 10, with each round marked by a unique color (details are provided in Section 4.4). Here the original predictor is LTD-GCN [27], and we add our MLP-based DeFeeNet on it. From top to bottom: observation/GT (covering a relatively long horizon of motions for consecutive prediction), w/o deviation (directly employing the existing predictor round by round to roughly implement consecutive prediction), w/ deviation (deviation-aware consecutive prediction with DeFeeNet). Poses drawn in dashed lines indicate GT, and our DeFeeNet helps to yield improved prediction.

time. Figure 4 shows visualized comparisons of consecutive motion prediction w/o or w/ deviation feedback on a sample of action *walking together*. We provide r6, r7 and r8 out of the total 10 rounds that marked with different colors. For example, in r8, the existing predictor only takes as input the poses in the yellow box from *line GT* as observation, and then predict the poses in the yellow box from *line w/o deviation*, to roughly implement one round of prediction. However, our DeFeeNet is able to detect the deviation between current observation and previous prediction of r7 (poses in the mint green from *line w/ deviation*), and therefore use this deviation feedback to predict more accurately (see poses in the yellow box from *line w/deviation*). Meanwhile, these predicted poses could also help the next round of prediction (r9) in the same manner as above. Note that we only draw 4 poses in each observed/predicted box for space saving, and the actual pose number is 10.

Additionally, in both Table 1 and 2, we find that round 1 is also improved although there is no deviation involved. This is reasonable as the total loss for each sample is calculated by adding the prediction loss of round 1 and 2, which

means that the second loss would also have an effect on the first prediction. To remove such effect, ablation study is given in supplemental material, containing comparison between deviation feedback enabled/disabled, to further show the effectiveness of DeFeeNet on accuracy improvement.

4.5. Results on BABEL

We provide the prediction errors of models with and without DeFeeNet inserted on BABEL in Table 3. We conduct experiments on (a) BABEL w/o transi (only single-motion samples of 11 action categories) and (b) BABEL w/ transi (involving two-action samples and the in-between action transitions within the 11 categories). From the table, our DeFeeNet is effective for error reduction whether action changing is involved or not. Similar to Human3.6M, the deviation produced by round 1 is learned by DeFeeNet and therefore can be used to improve the next round of prediction, which lays the foundation for consecutive motion prediction that requires deviation feedbacks round by round.

We also present the stability validation results in Table 4, which shows the errors of consecutive 5-round predic-

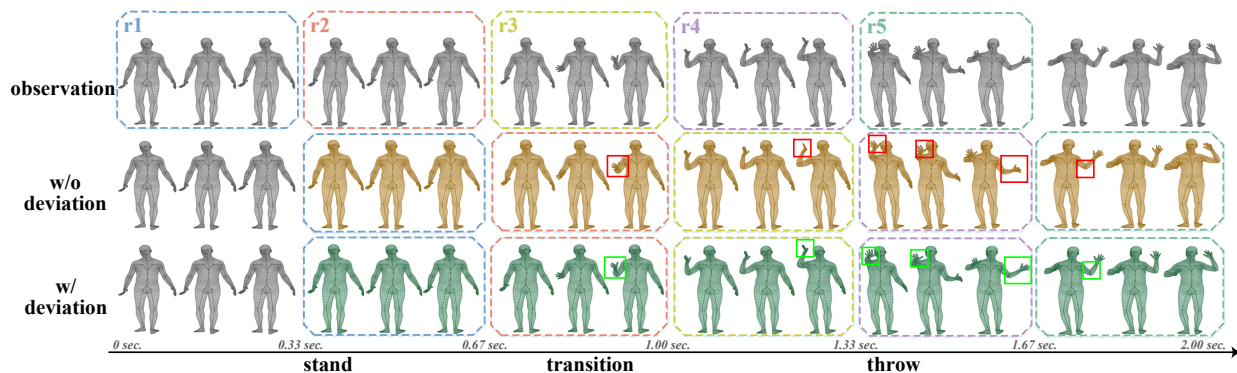


Figure 5. Visualized comparisons of consecutive motion prediction w/o or w/ deviation feedback on a sample *stand-transition-throw* from BABEL. Each round is marked with a unique color. Here the original predictor is STS-GCN [34] to roughly implement consecutive prediction. We add our GRU-based DeFeeNet on it, and obtain the improved prediction results with more accurate and elegant motions.

tion. Although slightly fluctuating, our DeFeeNet is still able to constantly get lower errors than isolated baselines. We provide the visualized comparisons on a sample named *stand-transition-throw* in Figure 5. As shown in the figure, our deviation-aware consecutive prediction could yield accurate prediction results even when faced with challenging action-changing samples. This is reasonable, as DeFeeNet is more sensitive to local information, i.e., the prediction deviation that *just happened*, which is a potential factor of action state changes. In our experiments, 5 rounds of prediction with observation part included accounts for 60 frames (i.e., 2 seconds). We do not select a longer horizon (like Human3.6M) due to that many motion sequences in [26] are shorter than 50 frames and cannot be processed as test samples for our multi-round prediction.

4.6. Ablation Study

To further validate our pipeline-inserted design is effective yet tailored to human motion prediction, and prove its essential difference against the “residual correction” spirit (mentioned in Section 2.2), we modify our DeFeeNet as DeFee-corr: a post-processing step of existing human motion predictors that aims to correct the already-produced prediction. The main architecture is retained but moved to the end of existing predictors, which detects *previous* prediction deviation, and then maps it into an estimated deviation embedding to be directly added on *current* round of prediction as “correction”. From Table 5, DeFee-corr fails to yield the improvement like ours, as baselines inserted with DeFeeNet allows for joint decoding on both current observation and previous deviation information, which leverages the decoding power of the existing predictors and therefore predict naturally improved motions, while simply adding the deviation vector produced by the simple yet small DeFee-corr for prediction correction is not suitable for delicate human motion data.

	LTD-GCN [27]		STS-GCN [34]		MotionMixer [4]	
isolated	0.4221		0.4673		0.3925	
	MLP	GRU	MLP	GRU	MLP	GRU
DeFee-r2	0.3945	0.3923	0.4150	0.4224	0.3640	0.3613
DeFee-corr-r2	0.4187	0.4122	0.4472	0.4603	0.3838	0.3892

Table 5. Comparisons of round 2 prediction errors (average on frame 3, 6, 8, 10) between DeFeeNet and DeFee-corr on BABEL w/ transi. Values in bold indicate lower errors.

5. Conclusion

We reformulate current human motion prediction task from the consecutive perspective which covers multiple rounds of “observe then predict”, and argue that each round of prediction accuracy can be improved if the prediction *deviation* generated in previous round is well detected and learned by robots/machines for practical use. In this paper, we propose DeFeeNet with MLP-based version and GRU-based version, both of which are simple yet effective, and can be inserted to existing human motion prediction models to realize deviation perception and feedback across adjacent prediction units. Our DeFeeNet is able to produce stably improved prediction during a relatively long horizon on Human3.6M as well as the more recent and challenging BABEL that contains samples with different actions and the in-between transitions.

Limitations. In real world, observations obtained by devices may contain noises and occlusions, which harm the deviation perception process, whose failure would further harm prediction quality. To utilize prediction deviation under imperfect condition still requires follow-up research.

Acknowledgements. This work was supported in part by the National Natural Science Foundation of China (NO. 62176125, 61772272).

References

- [1] Emre Aksan, Manuel Kaufmann, Peng Cao, and Otmar Hilliges. A spatio-temporal transformer for 3d human motion prediction. In *International Conference on 3D Vision (3DV)*, pages 565–574. IEEE, 2021. 2
- [2] Sadegh Aliakbarian, Fatemeh Sadat Saleh, Mathieu Salzmann, Lars Petersson, and Stephen Gould. A stochastic conditioning scheme for diverse human motion prediction. In *CVPR*, pages 5223–5232, 2020. 2
- [3] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 4
- [4] Arij Bouazizi, Adrian Holzbock, Ulrich Kressel, Klaus Dietmayer, and Vasileios Belagiannis. Motionmixer: Mlp-based 3d human body pose forecasting. *arXiv preprint arXiv:2207.00499*, 2022. 1, 3, 4, 5, 6, 7, 8
- [5] Yujun Cai, Lin Huang, Yiwei Wang, Tat-Jen Cham, Jianfei Cai, Junsong Yuan, Jun Liu, Xu Yang, Yiheng Zhu, Xiaohui Shen, et al. Learning progressive joint propagation for human motion prediction. In *ECCV*, pages 226–242. Springer, 2020. 2
- [6] Pablo Cervantes, Yusuke Sekikawa, Ikuro Sato, and Koichi Shinoda. Implicit neural representations for variable length human motion generation. In *ECCV*, pages 356–372. Springer, 2022. 4
- [7] Enric Corona, Albert Pumarola, Guillem Alenyà, and Francesc Moreno-Noguer. Context-aware human motion prediction. In *CVPR*, pages 6992–7001, 2020. 1, 2
- [8] Lingwei Dang, Yongwei Nie, Chengjiang Long, Qing Zhang, and Guiqing Li. Msr-gcn: Multi-scale residual graph convolution networks for human motion prediction. In *ICCV*, pages 11467–11476, 2021. 1, 2, 4, 5, 6
- [9] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *ICCV*, pages 4346–4354, 2015. 1, 2
- [10] Haifeng Gong, Jack Sim, Maxim Likhachev, and Jianbo Shi. Multi-hypothesis motion planning for visual object tracking. In *ICCV*, pages 619–626, 2011. 1
- [11] Chuan Guo, Xinxin Zuo, Sen Wang, Shihao Zou, Qingyao Sun, Annan Deng, Minglun Gong, and Li Cheng. Action2motion: Conditioned generation of 3d human motions. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2021–2029, 2020. 2
- [12] Wen Guo, Yuming Du, Xi Shen, Vincent Lepetit, Xavier Alameda-Pineda, and Francesc Moreno-Noguer. Back to mlp: A simple baseline for human motion prediction. *arXiv preprint arXiv:2207.01567*, 2022. 3, 4, 5, 6
- [13] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 4
- [14] Qian Huang, Horace He, Abhay Singh, Ser-Nam Lim, and Austin R Benson. Combining label propagation and simple models out-performs graph neural networks. *arXiv preprint arXiv:2010.13993*, 2020. 3
- [15] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013. 2, 4, 5
- [16] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *CVPR*, pages 5308–5317, 2016. 1, 2
- [17] Junteng Jia and Austion R Benson. Residual correlation in graph neural network regression. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 588–598, 2020. 3
- [18] Daejin Kim, Youngin Cho, Dongmin Kim, Cheonbok Park, and Jaegul Choo. Residual correction in real-time traffic forecasting. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*, pages 962–971, 2022. 3
- [19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [20] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *CVPR*, pages 5226–5234, 2018. 1, 2
- [21] Maosen Li, Siheng Chen, Yangheng Zhao, Ya Zhang, Yanfeng Wang, and Qi Tian. Dynamic multiscale graph neural networks for 3d skeleton based human motion prediction. In *CVPR*, pages 214–223, 2020. 1, 2
- [22] Binbing Liao, Jingqing Zhang, Ming Cai, Siliang Tang, Yifan Gao, Chao Wu, Shengwen Yang, Wenwu Zhu, Yike Guo, and Fei Wu. Dest-resnet: A deep spatiotemporal residual network for hotspot traffic speed prediction. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 1883–1891, 2018. 3
- [23] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *ACM transactions on graphics (TOG)*, 34(6):1–16, 2015. 3
- [24] Tiezheng Ma, Yongwei Nie, Chengjiang Long, Qing Zhang, and Guiqing Li. Progressively generating better initial guesses towards next stages for high-quality human motion prediction. In *CVPR*, pages 6437–6446, 2022. 1, 6
- [25] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. Generating smooth pose sequences for diverse human motion prediction. In *ICCV*, pages 13309–13318, 2021. 2
- [26] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. Weakly-supervised action transition learning for stochastic human motion prediction. In *CVPR*, pages 8151–8160, 2022. 2, 3, 5, 8
- [27] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *ICCV*, pages 9489–9497, 2019. 1, 2, 3, 4, 5, 6, 7, 8
- [28] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *CVPR*, pages 2891–2900, 2017. 1, 2
- [29] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 6

- [30] Dario Pavllo, Christoph Feichtenhofer, Michael Auli, and David Grangier. Modeling human motion with quaternion-based neural networks. *International Journal of Computer Vision*, 128(4):855–872, 2020. [2](#)
- [31] Mathis Petrovich, Michael J Black, and Gül Varol. Action-conditioned 3d human motion synthesis with transformer vae. In *ICCV*, pages 10985–10995, 2021. [2](#), [3](#)
- [32] Abhinanda R Punnakkal, Arjun Chandrasekaran, Nikos Athanasiou, Alejandra Quiros-Ramirez, and Michael J Black. Babel: Bodies, action and behavior with english labels. In *CVPR*, pages 722–731, 2021. [2](#), [5](#)
- [33] Alessio Sampieri, Guido Maria D’Amely di Melendugno, Andrea Avogaro, Federico Cunico, Francesco Setti, Geri Skenderi, Marco Cristani, and Fabio Galasso. Pose forecasting in industrial human-robot collaboration. In *European Conference on Computer Vision*, pages 51–69. Springer, 2022. [1](#), [2](#)
- [34] Theodoros Sofianos, Alessio Sampieri, Luca Franco, and Fabio Galasso. Space-time-separable graph convolutional network for pose forecasting. In *ICCV*, pages 11209–11218, 2021. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [35] Xiaoning Sun, Qiongjie Cui, Huaijiang Sun, Bin Li, Weiqing Li, and Jianfeng Lu. Overlooked poses actually make sense: Distilling privileged knowledge for human motion prediction. In *ECCV*, pages 678–694. Springer, 2022. [2](#)
- [36] Ilya O Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. *NeurIPS*, 34:24261–24272, 2021. [4](#)
- [37] Dong Wei, Huaijiang Sun, Bin Li, Jianfeng Lu, Weiqing Li, Xiaoning Sun, and Shengxiang Hu. Human joint kinematics diffusion-refinement for stochastic motion prediction. In *AAAI*, 2022. [2](#)
- [38] Julian Wiederer, Arij Bouazizi, Ulrich Kressel, and Vasileios Belagiannis. Traffic control gesture recognition for autonomous vehicles. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 10676–10683. IEEE, 2020. [1](#)
- [39] Ye Yuan and Kris Kitani. Dlow: Diversifying latent flows for diverse human motion prediction. In *ECCV*, pages 346–364, 2020. [2](#)
- [40] Chongyang Zhong, Lei Hu, Zihao Zhang, Yongjing Ye, and Shihong Xia. Spatio-temporal gating-adjacency gen for human motion prediction. In *CVPR*, pages 6447–6456, 2022. [1](#), [2](#), [5](#), [6](#)
- [41] Yiqi Zhong, Zhenyang Ni, Siheng Chen, and Ulrich Neumann. Aware of the history: Trajectory forecasting with the local behavior data. In *ECCV*, pages 393–409. Springer, 2022. [3](#)