

Multi-Person Extreme Motion Prediction with Cross-Interaction Attention

Wen Guo¹, Xiaoyu Bie¹, Xavier Alameda-Pineda¹, Francesc Moreno-Noguer²

¹Inria, Univ. Grenoble Alpes, CNRS, Grenoble INP, LJK, 38000 Grenoble, France

²Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain

¹{wen.guo,xiaoyu.bie,xavier.alameda-pineda}@inria.fr, ²{fmoreno}@iri.upc.edu

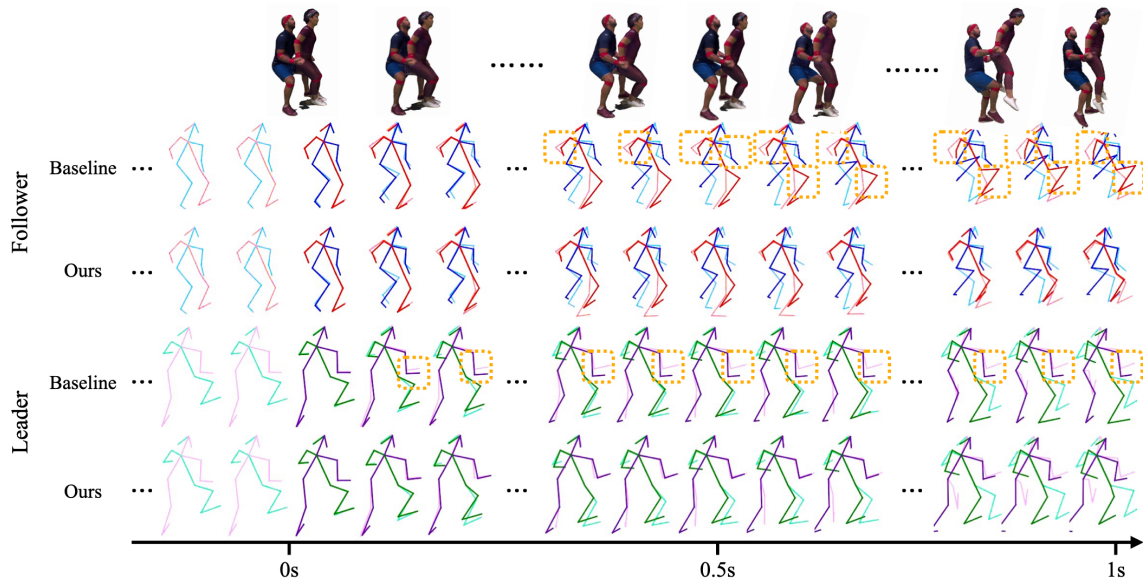


Figure 1: **Top row:** Samples of 3D meshes from our Extreme Pose Interaction Dataset(ExPI Dataset) proposed in this paper. **Bottom:** Motion prediction results. The first two rows (red/blue) are the follower, the last two rows (purple/green) are the leader. Light colors indicate the input sequence (before 0 s) and the ground-truth (after). The dark colors show the prediction results for the baseline model (top) and our method (bottom). The dashed boxes point to failures of the baseline method, where our proposed method succeeds. In this paper, we use only the 3D pose labels.

Abstract

Human motion prediction aims to forecast future human poses given a sequence of past 3D skeletons. While this problem has recently received increasing attention, it has mostly been tackled for single humans in isolation. In this paper we explore this problem from a novel perspective, involving humans performing collaborative tasks. We assume that the input of our system are two sequences of past skeletons for two interacting persons, and we aim to predict the future motion for each of them. For this purpose, we devise a novel cross interaction attention mechanism that exploits historical information of both persons and learns to predict cross dependencies between self poses and the poses of the other person in spite of their spatial or temporal dis-

tance. Since no dataset to train such interactive situations is available, we have captured ExPI (Extreme Pose Interaction), a new lab-based person interaction dataset of professional dancers performing acrobatics. ExPI contains 115 sequences with 30k frames and 60k instances with annotated 3D body poses and shapes. We thoroughly evaluate our cross interaction network on this dataset and show that both in short-term and long-term predictions, it consistently outperforms baselines that independently reason for each person. We plan to release our code jointly with the dataset and the train/test splits to spur future research on the topic.

1. Introduction

Human motion prediction aims to predict future motions from previous observations. With the huge development of human pose estimation from single images by deep learning methods [3, 8, 16, 24, 34, 35, 38–41, 47], motion prediction begins to draw an increasing attention [1, 2, 7, 11, 12, 15, 18, 21, 25, 28, 31–33, 42]. Most existing works formulate motion prediction as a sequence-to-sequence task, in which past observations represented as 3D skeleton data are injected to a deep network which forecasts skeleton movements in a future temporal window.

In any event, a common denominator of the approaches [11, 15, 18, 21, 31, 33, 45] for human motion prediction is that they treat the human as an independent and isolated entity: The motion predicted for a person relies solely on his past motion. However, in the real world scenario, people interact a lot with each other and the motion of one person is typically dependent to the motion of the person that he is interacting with. For example, in Acroyoga, we could infer the motion of one person by seeing only the motion of his partner, this can be very useful especially when one person is occluded. Thus, we could potentially improve the performance of motion prediction by exploiting the human interaction.

In this paper we present a novel task: The collaborative motion prediction, which aims at predicting joint motion of two persons performing highly interactive actions. To the best of our knowledge, only [7] addressed a related problem in which weak person-object interactions were explored. Here we aim to step further and analyse situations with very strong person-person interactions, which we may encounter in *e.g.* team sports or collaborative assembly tasks in factories.

With the goal to foster research on these new task, we collected a large dataset of professional dancers performing aerial Lindy Hop¹ steps. Our dataset, dubbed ExPI (Extreme Pose Interaction), contains 115 sequences of 2 professional couples performing 16 different aeriels. ExPI has been recorded in a multiview motion capture studio, and both the 3D poses and 3D shapes of the two persons are annotated for each all the 30K frames. We have carefully created train/test splits, and proposed three different extensions of the pose evaluation metrics to adapt them to the collaborative motion prediction task. We plan to release this dataset to the community. Interestingly, to perform these aeriels, the two dancers perform different movements that require a high level of synchronisation. These aeriels are composed of extreme poses and require strict and close cooperation between the two persons, which is highly suitable for the study of collaborative motion prediction. Some examples

of this highly-interacted dataset could be see in Figure 2.

To model such strong human-to-human interactions, we introduce a novel Cross-Interaction Attention (XIA) mechanism that allows to simultaneously exploit historical motion data of the two persons. XIA builds upon a standard multi-attention head [46], in which the attention module is updated so as to consider bidirectional mappings between the queries-keys-values of the two persons.

A thorough evaluation of the proposed XIA mechanism on the ExPI dataset shows that our method consistently improves (5 ~ 10% accuracy improvement) other baselines that perform independent motion prediction per person. We overcome these approaches both for short (≤ 500 ms) and long term prediction (500 ms ~ 1000 ms).

Our key contributions can be summarized as follows:

- We introduce the task of collaborative motion prediction, to focus on the estimation of future poses of people in highly interactive setups.
- We collect and will make publicly available ExPI, a large dataset of highly interacted extreme dancing poses, annotated with 3D joint locations and body shapes. Additionally we define the benchmark with carefully selected train/test splits and evaluation protocols.
- We propose a novel cross-interaction attention (XIA) mechanism that exploits historical motion of two persons to predict future movements. We demonstrate the interest of XIA, and more generally of the collaborative motion prediction task and the ExPI dataset.

2. Related Work

2.1. Human Motion Prediction

As the dominant architecture in 3D human motion prediction, RNNs based models have been widely studied in recent years. For instance, Fragkiadaki *et al.* [11] propose an Encoder-Recurrent-Decoder (ERD) model where human motion is embedded and reconstructed via the encoder-decoder framework and an LSTM cell is used to update the latent space. To avoid error accumulation, Fragkiadaki *et al.* [11] use curriculum learning by corrupting the mocap vectors with zero-mean Gaussian with progressively increasing variance. Jain *et al.* [21] combine the spatial-temporal graph with RNN to build a structured latent space, however, this manually designed structure will restrict the capacity of long-term spatial modeling. Martinez *et al.* [33] outperforms [11, 21] by introducing a residual connection to model the velocities instead of the poses themselves. To alleviate the discontinuities in the previous work, Gui *et al.* [15] turns to use adversarial training to smooth the results.

Although RNNs have achieved great success in sequence-to-sequence prediction, they suffer from contain-

¹The Lindy Hop is an African-American couple dance born in the 1930's in Harlem, New York, see [37]. The dancing action is called aerial.

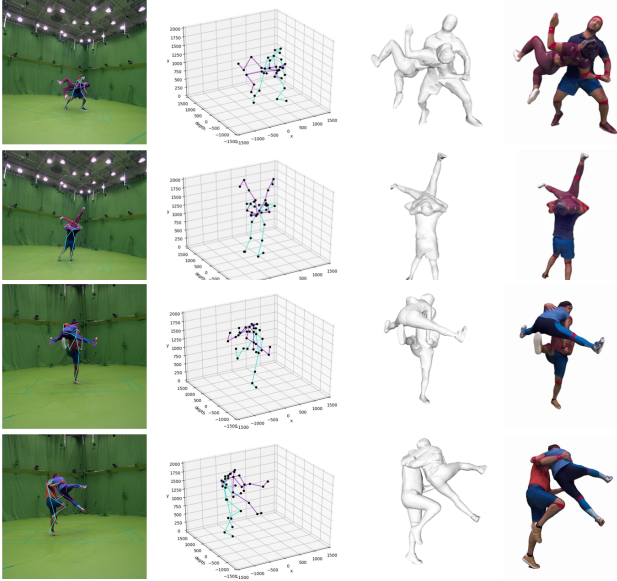


Figure 2: Samples of the ExPI dataset: raw image, 3D pose, mesh and textured mesh.

ing the entire history with a fix-size hidden state in long-term prediction. In such consideration, several works employ feed-forward models, like sliding windows [4, 5] or convolutional models [18, 19, 26]. Later, Mao *et al.* [32] proposed to use Discrete Cosine Transform (DCT) to represent human motion in trajectory space instead of pose space. In [31] this approach is further developed by following a graph convolutional network (GCN) and introducing an attention mechanism. Very recently, Aksan *et al.* [1] propose a spatio-temporal transformer (ST-Transformer), which replaces the GCN with a fully autoregressive model.

2.2. Contextual Information in Human Interaction

Recent works demonstrate the benefit of aggregating contextual information in 3D pose estimation [16, 17, 22, 27, 48]. Precisely, Li *et al.* [27] encode interaction information by dividing human relations into different levels and use 3 extra losses to enhance ordering of different people and joints. Hassan *et al.* [17] exploit static 3D scene structure to better estimate human pose from monocular images. Zanfiri *et al.* [48] integrate the scene constraints for multiple interacting people. Jiang *et al.* [22] introduce a loss to penalize interpenetration of reconstructed people and promote depth ordering. Guo *et al.* [16] improve the performance by looking at the whole group of humans. More recently, Corona *et al.* [7] expand the use of contextual information into motion prediction problem by proposing an architecture with a semantic-graph model, where nodes parameterize the human and objects in the scene and the edges their interactions. Weak human-to-human or human-to-object correlations are modeled using this approach.

Table 1: Comparison of ExPI with other datasets on 3D human modeling. ExPI will be made publicly available.

Dataset	AMASS [30]	H3.6m [20]	MuPoTS [36]	CHI3D [10]	ExPI
3D joints	✓	✓	✓	✓	✓
Video	✓	✓		✓	✓
Shape	✓	✓		✓	✓
Multi-person			✓	✓	✓
Extreme poses					✓
Released	✓	✓	✓		✓
Multi-view				✓	✓

In this paper we use [31] as baseline. We retain the global attention as well as the GCN predictor following the original settings, but on top of this we introduce human-human interaction as supplementary contextual information. Experimental results on our newly collected ExPI dataset verify that our approach outperforms the baseline.

2.3. Datasets

Using deep learning methods to study 3D human pose tasks relies on high-quality datasets. Most previous 3D human datasets are single person [20, 30, 44] or made of pseudo 3D poses [36, 47]. Other datasets which contain label-based 3D data usually do not explicitly have close interactions [14, 29, 36, 43]. Recently, some works start to pay attention to the importance of interaction and propose newly collected dataset to model interaction of artificial persons with real scenes [6, 49]. Besides, Fieraru *et al.* [9] created a dataset of human interaction with a contact-detection-aware framework. However, this dataset just contains several daily interaction scenarios, and the data is not released yet.

Thus, the new ExPI dataset we propose which is designed with extreme highly interacted poses will be helpful to several tasks, including human pose estimation, human motion prediction and human shape reconstruction. Table 1 shows a comparison of our dataset with some of the most popular current 3D datasets on 3D human pose/motion.

3. New task: Collaborative Motion Prediction

As discussed in the introduction, the task of human motion prediction, while fairly recent in the literature, is well established. In this paper we want to propose and start investigating a new task: collaborative motion prediction. We first recall the human motion prediction task, to later on formalise the new task.

The task of human motion prediction is defined as predicting future poses $P_{t:t_E}$ (from time t to t_E) from a previous observation $P_{t_1:t-1}$ (from t_1 to $t-1$), where t_1 denote the initial sequence frames and t_E denotes the ending sequence frames. The standard motion prediction task aims to learn a mapping \mathcal{M} to estimate the future movements of one person:

$$\mathcal{M} : P_{t_1:t-1} \rightarrow P_{t:t_E}, \quad (1)$$

where P_t denotes the pose of the person at time t .

We are interested in extending the problem formulation to two persons within the same physical interaction. While our formulation is general and could work for any kind of interactions, for the sake of consistency throughout the paper, we will here denote by ℓ and f variables corresponding to the leader and the follower respectively (see Section 4 on the dataset description). Therefore, the mapping we would like to learn for *collaborative motion prediction* is written as:

$$\mathcal{M}_{collab} : P_{t_i:t-1}^\ell, P_{t_i:t-1}^f \rightarrow P_{t:t_E}^\ell, P_{t:t_E}^f. \quad (2)$$

In other words, since the two persons are involved in the same interaction, we believe it is possible to exploit the pose information of the interaction partner to better predict the motion of any of the two persons interacting.

In addition to defining the collaborative motion prediction task, we provide an experimental framework – consisting of a dataset and evaluation metrics – to foster research in this direction, and the first results obtained with our new method based on cross-interaction attention.

4. The Extreme Pose Interaction Dataset

We present the Extreme Pose Interaction (ExPI) Dataset, a new person interaction dataset of Lindy Hop aerial steps. In Lindy Hop, the two dancers have different roles, referred to as *leader* and *follower*.² We recorded two couples of dancers, one at a time, in a multi-camera environment, equipped also with a motion-capture system. In this section we first describe the recording procedure and the recorded data, to later on discuss the data post-processing procedure.

4.1. Dataset Structure

In the ExPI dataset 16 different aeriels are performed, some by both dancer couples, some by only one of the couples, as shown in Table 2. Each of the aeriels was repeated five times to account for variability. Overall, ExPI contains 115 short sequences, each one depicting an execution of one of the aeriels. The data were collected in a multi-camera platform³ equipped with 60+ synchronised and calibrated color cameras and a motion capture system. More precisely, for each recorded sequence ExPI provides:

- Multi-view image sequences at 50FPS, from all the available cameras.
- Mocap data: 3D position of all joints per frame. We track 18 joints (3 in the head one in the neck and both shoulders, elbows, wrists, hips, knees, heels and toes) per person. Mocap data is recorded at 100FPS.
- Camera calibration information.

²This is the standard gender-neutral terminology.

³A detailed description of the recording platform could compromise the double-blind review process, and will be provided in the final version of the manuscript.

Table 2: Composition of the ExPI Dataset. The seven first aeriels are performed by both couples. Six more aeriels are performed by Couple 1, while three others by Couple 2.

Aerial	Name	Couple 1	Couple 2
A_1	A-frame	✓	✓
A_2	Around the back	✓	✓
A_3	Coochie	✓	✓
A_4	Frog classic	✓	✓
A_5	Noser	✓	✓
A_6	Toss out	✓	✓
A_7	Cartwheel	✓	✓
A_8	Back flip	✓	
A_9	Big ben	✓	
A_{10}	Chandelle	✓	
A_{11}	Check the change	✓	
A_{12}	Frog-turn	✓	
A_{13}	Twisted toss	✓	
A_{14}	Crunch-toast		✓
A_{15}	Frog-kick		✓
A_{16}	Ninja-kick		✓

- 3D shapes as a texture mesh for each frame.

Overall, the dataset contains almost 30k visual frames for each view point and 60k 3D instances annotated. The dataset will be released to the community to foster research in this direction.

4.2. Data Post-processing

When collecting the motion capture data, many points are missed by the system. This is often due to occlusions or tracking losses. To overcome this issue, we manually post-processed the missing points. We have designed and implemented a 3D hand labelling toolbox. For each missed value, we choose two orthogonal views among the several view-points, and label the missed keypoints by hand on these two frames to get two image coordinates. We then use the camera calibration to back project these two image coordinates into the 3D world coordinate, obtaining two straight lines. Ideally, the intersection of these two lines is the world coordinate of this missing point. Since these two lines do not always intersect, we find the nearest point, in the least-squares sense, to these two lines to represent the approximated intersection. In this procedure we did not use the distortion parameters, since we observed that the distortion error is negligible on the views we choose for labeling. The intersection is project into 3D and various 2D images to confirm the quality of the approximation by visual inspection.

5. Method

We next introduce the proposed model for collaborative motion prediction. We first discuss the problem formulation, then discuss the overall pipeline and finally provide the details on the cross interaction attention (XIA) module.

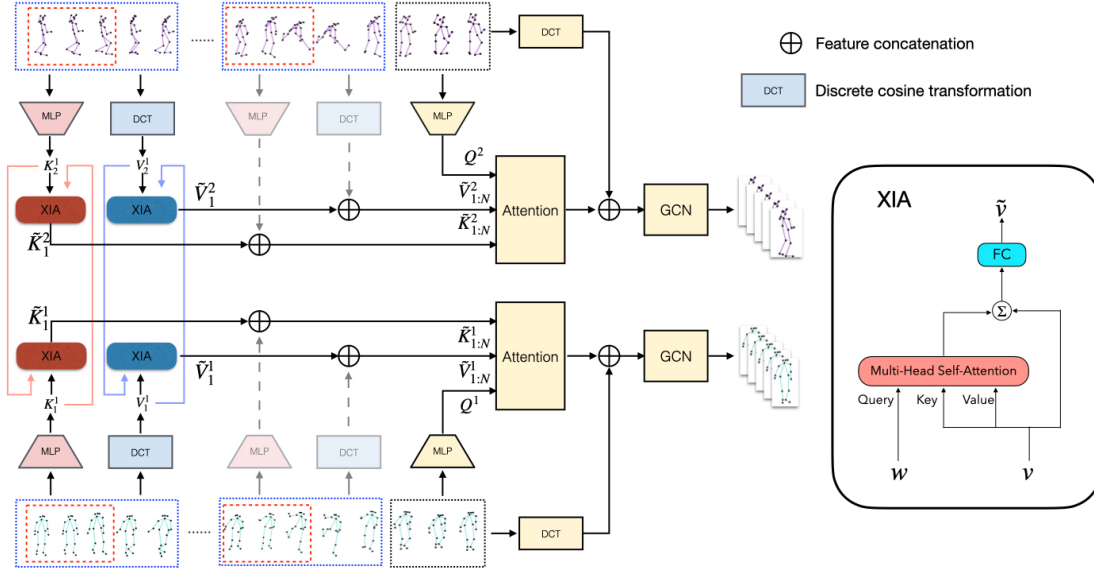


Figure 3: **Left:** Computing flow of the proposed method. Two parallel pipelines – for the leader and the follower – are implemented. The keys and values are refined with XIA modules (different for key and value, leader and follower). **Right:** Cross-interaction attention (XIA) mechanism. The multi-head attention is queried with w , and the key and value are the vector v to be refined. The output of the MHA is a residual value to add to v before feeding a FC layer.

5.1. The Base Method

We build on a recent state-of-the-art motion prediction model [31], and extend it to account for interaction between pose sequences of two persons. The model in [31] – referred to as *base* method – consists on two parts: an attention model used for learning temporal attention w.r.t. the past motions, and a predictor based on GCN [23] used to model the spatial attention among joints by the adjacency matrix. The temporal attention model aims at finding the most relative sub-sequence in the past by measuring the similarity between the last observed sub-sequence $P_{t-1-M:t-1}$ and a set of past sub-sequences $P_{t_i:t_i+M+T}$, where t_i with $i \in \{1, \dots, N\}$ indicates the start frame of each past sub-sequence in observation. The base model uses a transformer, where the query and the keys are learned from a multi-layer perceptron from sequences of length M . More precisely, Q denotes the representation associated to $P_{t-1-M:t-1}$ and K_i denotes the representation of the start chunk of the i -th subsequence, which is $P_{t_i:t_i+M}$. For each key K_i there is an associated value V_i consisting of a DCT representation built from the sub-sequence $P_{t_i:t_i+M+T}$. See [31] for a detailed discussion.

As discussed in Section 3, the task of collaborative motion prediction consists in learning a mapping that can predict the motion of two persons interacting. As a task, it is a natural extension of the original human motion prediction task to two persons, which means a couple of dancers (a leader and a follower, see Eq. (2)) in our case.

One straightforward application of the base method to

the collaborative human prediction task, is to learn a leader-specific and a follower-specific human motion prediction mappings. In other words, we train [31] to predict the future motion of the leader separately from the follower. We refer to this as *independent motion prediction* (IMP), and use it as our baseline. While this strategy could exploit the attention and prediction modules proposed in [31], it does not account for any interaction between the two dancing partners. In the next section we describe how to address this issue with the proposed cross-interaction attention module.

5.2. Cross-Interaction Attention

To exploit the interactions to predict future human motions of both the follower and the leader, we will build upon the base method. In particular, we will have one query for the leader Q^ℓ and one for the follower Q^f , as well as several key-value pairs for the leader (K_i^ℓ, V_i^ℓ) and for the follower (K_i^f, V_i^f) . We then naturally cast the collaborative human motion prediction task in how to jointly exploit the information in (K_i^ℓ, V_i^ℓ) and (K_i^f, V_i^f) when querying with Q^ℓ or Q^f . Our intuition is that the pose information (key-value pairs) of one dancer can be used to transform the pose information of the other dancer for better motion prediction. We implement this intuition with the help of the proposed *cross-interaction attention modules* (XIA). Such a module takes as input a pair of vectors v, w and uses multi-head self attention to refine the first vector as follows (see Fig. 3-right):

$$\tilde{v} = \text{XIA}(v, w) = \text{FC}(\text{MHA}(w, v, v) + v), \quad (3)$$

Table 3: Aerial-wise results on **CA** split with the three evaluation metrics (in *mm*). Model is trained and tested on common aerials performed by different actors.

Dancer		Leader							Follower								
Dancer		A_1	A_2	A_3	A_4	A_5	A_6	A_7	AVG	A_1	A_2	A_3	A_4	A_5	A_6	A_7	AVG
JME	Base	91.5	124.4	83.5	61.5	91.7	88.1	76.9	88.2	113.0	268.5	133.6	99.4	131.5	155.1	223.0	160.6
	XIA	78.0	114.3	73.2	55.5	82.6	82.7	63.6	78.6	108.6	250.0	128.3	91.6	123.5	137.8	208.5	149.8
SME	Base	91.5	124.4	83.5	61.6	91.7	88.1	76.9	88.2	101.6	144.9	106.7	89.4	100.1	99.8	239.7	126.0
	XIA	78.0	114.4	73.2	55.5	82.6	82.7	63.6	78.6	93.9	139.8	101.0	88.5	93.4	91.7	241.2	121.3
AME	Base	68.4	90.9	57.6	44.3	53.2	68.1	55.4	62.6	71.3	106.0	73.8	67.2	72.6	72.2	154.0	88.2
	XIA	61.7	82.8	52.4	39.6	48.2	63.4	48.6	56.7	67.3	102.4	69.8	64.9	65.2	66.1	150.0	83.7

Table 4: Joint-wise (averaged) results on the **CA** split using the three evaluation metrics (in *mm*). We bold the results when the improvement is larger than *5mm* in order to show the mostly improved joints.

Dancer		Leader								Follower									
Joint		Head	Back	Sho.	Elbow	Wrist	Hip	Knee	Heel	Toes	Head	Back	Sho.	Elbow	Wrist	Hip	Knee	Heel	Toes
JME	Base	44.4	13.2	40.5	93.2	139.8	5.2	110.9	155.9	175.5	134.5	131.9	129.3	129.4	135.5	151.7	180.2	218.5	233.1
	XIA	42.4	13.4	39.5	86.0	127.0	5.2	95.7	134.0	149.5	122.0	120.3	119.0	122.5	127.8	145.9	166.7	204.7	218.0
SME	Base	44.4	13.2	40.5	93.2	139.8	5.2	110.9	155.9	175.6	83.9	30.5	66.1	108.3	142.0	14.5	166.3	231.5	264.6
	XIA	42.4	13.4	39.5	86.0	127.0	5.2	95.7	134.0	149.5	79.5	28.2	63.4	104.7	137.3	14.4	158.8	224.9	255.2
AME	Base	40.6	31.8	39.2	67.0	95.9	53.3	66.4	79.7	85.0	60.9	51.8	57.5	77.1	103.0	73.6	105.6	119.6	139.8
	XIA	34.4	27.4	34.3	61.1	90.3	45.7	63.6	73.4	76.5	57.4	49.4	55.1	73.1	97.4	73.0	98.4	112.6	132.4

where $\text{MHA}(a, b, c)$ stands for multi-head attention with query a , key b and value c , and FC for fully connected.

The proposed XIA module is integrated at several stages of the computing flow as shown in Figure 3. More precisely, we use a XIA module to refine all keys:

$$\tilde{K}_i^\ell = \text{XIA}(K_i^\ell, K_i^f), \quad \tilde{K}_i^f = \text{XIA}(K_i^f, K_i^\ell), \quad (4)$$

and analogously for the values. In our implementation, the multi-head attention has 8 attention heads for the key, and a single attention head for the value. Moreover, we add a skip-connection for the MHA module following by two FC layers with identical dimension as input.

5.3. Pose Normalization

The raw 3D data of the ExPI dataset is represented in an arbitrary world coordinate. However the position and orientation of the dancers change from one sequence to another. Therefore we need a way to normalise the data so it is possible for the base and proposed method to learn the motion patterns and successfully predict future motions.

We apply a rigid transformation (translation and rotation) to the couple together on every frame, so that the pose errors can be computed directly in this new representation. More precisely, at every frame, we translate all joints so that the middle point of the hips of the leader moves to the coordinate origin. Furthermore, the rotation transforms the space in a way that the unit vector from the hip-center to the left hip of the leader become the x -axis, and the neck is used to determine XOZ plane. For each time t , a transformation

matrix $T_{\tilde{P}_t^\ell}$ is computed to align the leader, then we apply it both on the raw position of leader \tilde{P}_t^ℓ and follower \tilde{P}_t^f :

$$P_t^\ell = T_{\tilde{P}_t^\ell}(\tilde{P}_t^\ell), \quad P_t^f = T_{\tilde{P}_t^\ell}(\tilde{P}_t^f). \quad (5)$$

6. Experimental Evaluation

This section describes the experimental protocol for ExPI dataset, various dataset splits and evaluation metrics.

6.1. Splitting the ExPI Dataset

As shown in Table 2, ExPI Dataset contains a set of common aerials and various extra (couple-specific) aerials (16 in total). $\mathcal{A}_c^j = \{A_i^j\}_{i=1}^7$ is the set of 7 common aerials performed by couple j . We denote the reunion of these two subsets as: $\mathcal{A}_c = \mathcal{A}_c^1 \cup \mathcal{A}_c^2$. Regarding the extra aerials, \mathcal{A}_e^1 and \mathcal{A}_e^2 denote the set of aerials performed only by the first (A_8 to A_{13}) and the second (A_{14} to A_{16}) couples. With this notation we propose three dataset splits.

Single Aerial (SA) split. In this split, we train one model for each common aerial. More precisely, the training set is A_j^2 and the test set is A_j^1 for $j = 1, \dots, 7$, we will have 7 action specific models in the end.

Common Aerial (CA) split. Similar to [20], we consider the common actions done by different couple of actors as train and test data. More precisely, \mathcal{A}_c^2 is the train dataset and \mathcal{A}_c^1 is the test dataset. By doing so, train and test data contain same aerials but performed by different dancers.

Extra Aerial (EA) split. In this split, the train set is the

Table 5: Results on different predicting time window (CA).

Dancer	Leader				Follower				
Time (ms)	80	400	720	1000	80	400	720	1000	
# Frames	2	10	18	25	2	10	18	25	
JME	Base	83.3	84.8	86.8	88.2	151.2	153.9	157.4	160.6
	XIA	73.8	75.2	77.1	78.6	141.2	143.1	146.7	149.8
SME	Base	83.3	84.8	86.8	88.2	117.3	119.9	123.5	126.0
	XIA	73.7	75.2	77.1	78.6	112.8	115.3	118.7	121.3
AME	Base	60.1	60.8	61.8	62.6	82.6	84.4	86.3	88.2
	XIA	54.1	54.9	55.9	56.7	77.6	79.6	81.6	83.7

entire set of common aerials performed by both couples (*i.e.* \mathcal{A}_c), and the test set is the rest extra aerials of both couples \mathcal{A}_e^1 and \mathcal{A}_e^2 . By doing so, the train data and test data contain aerials of same actors, but the aerials for testing are not seen in the training data.

6.2. Evaluation metrics

The most common metric for evaluating joint position in pose estimation and motion prediction tasks is the mean per joint position error (MPJPE):

$$\text{MPJPE} = \frac{1}{TJ} \sum_{t,j=1}^{T,J} \left\| \hat{P}_{t,j} - P_{t,j} \right\|_2, \quad (6)$$

where T and J are the number of frames and joints respectively, and $\hat{P}_{t,j}$ and $P_{t,j}$ are the estimated and ground truth position of joint j at frame t .

JME: A straightforward extension of the standard metric to our collaborative human motion prediction problem consists on averaging through the joints of the leader and the follower by considering them in a same coordinate. We call this the joint mean per joint position error, and denote it as JME for simplicity:

$$\text{JME} = \frac{1}{2TJ} \sum_{t,j=1}^{T,J} \left\| \hat{P}_{t,j}^\ell - P_{t,j}^\ell \right\|_2 + \left\| \hat{P}_{t,j}^f - P_{t,j}^f \right\|_2. \quad (7)$$

JME provides an overall idea of the performance of the methods in the collaborative motion prediction task by measuring both the error of poses and their relative position.

SME: If the position of the dancers is not correctly estimated, then JME will be large even if the relative pose is correct. To complete this, we propose a separated version of JME where each pose is further normalised independently. More precisely, we define $\bar{P}_t^\ell = T_{\hat{P}_t^\ell}(\hat{P}_t^\ell)$ and $\bar{P}_t^f = T_{\hat{P}_t^f}(\hat{P}_t^f)$, where T_P is the normalisation transformation computed from the pose P as defined in Section 5.3. Very importantly, the new normalisation erase the bias of the relative position between the two dancers. The new met-

ric, called SME, writes like:

$$\text{SME} = \frac{1}{2TJ} \sum_{t,j=1}^{T,J} \left\| \bar{P}_{t,j}^\ell - P_{t,j}^\ell \right\|_2 + \left\| \bar{P}_{t,j}^f - P_{t,j}^f \right\|_2. \quad (8)$$

AME: While SME provides a position-free error metric, its reliability is conditioned to providing good estimates of the joints that are used to compute the normalisation transform T . If these joints are incorrectly estimated, so will be the transformation T , and the associated error will be unreliable. To mitigate this effect, we further compute the best rigid alignment between the estimated pose $\hat{P}_{t,j}^f$ and the ground-truth $P_{t,j}^f$ using [13], obtaining $\tilde{P}_{t,j}^f$. The aligned metric AME writes like:

$$\text{AME} = \frac{1}{2TJ} \sum_{t,j=1}^{T,J} \left\| \tilde{P}_{t,j}^\ell - P_{t,j}^\ell \right\|_2 + \left\| \tilde{P}_{t,j}^f - P_{t,j}^f \right\|_2. \quad (9)$$

To summarize, **JME** encodes the two relative poses together with the position bias, hence we propose **SME** to see the pure pose error. However the latter can be unreliable if the estimation of the hips and back is inaccurate, which lead us to complement these two metrics with **AME**.

6.3. Implementation Details

Since this is the first time the Collaborative Motion Prediction task is presented in the literature, there are no available methods to compare with. As we build our architecture from [31], we believe the fairest comparison is to train two different models – for the leader and the follower – since our architecture will also have a leader- and a follower-dedicated branch (see Figure 3): this will be denoted as “Base”. The proposed cross-interaction attention architecture will be denoted as “XIA”. The ExPI dataset is collected at 50 FPS, we downsample our data to 25 FPS. We train our model using 50 input frames and calculate the loss on 10 predicted frames. When testing for longer futures, we regard the prediction as the lasted observation to iterate the future frames. Besides, in order to avoid high variability, we follow [31] and take 64 sub-sequences for each testing sequence. So we have 320 testing sub-sequences for each aerial. Overall, we have 6k and 2k sub-sequences for training and testing respectively in the CA and SA splits. In the EA split, the training set has 12k sub-sequences, and we have 2k testing sequences.

6.4. Results and Discussion

CA split. Results of the Base and XIA method on the CA split are reported in Tables 3 (per aerial), 4 (per joint) and 5 (different prediction times). Beside of Table 5, all results are for 1s predictions. From Table 3 we observe that the proposed XIA method outperforms the Base method systematically almost for all aerials and in all metrics. Similar conclusions can be extracted when reporting the per-joint error

Table 6: Aerial-wise results on **SA** split with the three evaluation metrics (in *mm*). Seven action-wise models are trained independently.

Dancer		Leader							Follower								
Aerial		A_1	A_2	A_3	A_4	A_5	A_6	A_7	AVG	A_1	A_2	A_3	A_4	A_5	A_6	A_7	AVG
JME	Base	114.5	163.9	94.3	90.9	101.9	121.3	75.2	108.8	159.2	385.6	188.4	159.5	177.9	284.2	286.7	234.5
	XIA	94.3	160.0	80.3	75.8	90.1	123.4	79.1	100.4	148.3	370.6	146.1	143.9	160.4	186.7	247.4	200.5
SME	Base	114.8	163.5	95.3	90.2	102.6	122.5	75.2	109.1	121.7	185.8	131.4	115.4	118.5	125.4	272.7	153.0
	XIA	93.6	159.4	79.4	75.7	90.0	121.4	79.6	99.9	111.1	186.6	119.1	110.1	106.4	107.3	258.0	142.7
AME	Base	93.5	125.1	69.7	66.8	62.7	100.6	64.1	83.2	99.3	148.2	94.7	90.8	89.7	101.7	192.2	116.7
	XIA	73.6	122.7	61.2	57.7	56.6	97.1	63.6	76.1	84.8	141.0	81.3	83.4	79.9	79.0	173.9	103.3

Table 7: Aerial-wise results on **EA** split with the three evaluation metrics (in *mm*). Models are tested on unseen aerials.

Dancer		Leader									Follower										
Aerial		A_8	A_9	A_{10}	A_{11}	A_{12}	A_{13}	A_{14}	A_{15}	A_{16}	AVG	A_8	A_9	A_{10}	A_{11}	A_{12}	A_{13}	A_{14}	A_{15}	A_{16}	AVG
JME	Base	96.6	112.6	67.4	171.7	73.2	54.9	66.1	78.9	77.5	88.8	197.0	207.8	138.4	226.9	184.9	136.6	160.8	150.6	153.4	172.9
	XIA	93.5	112.6	65.8	168.9	73.1	53.8	64.7	76.9	73.4	87.0	203.3	190.8	135.5	228.7	180.1	137.7	151.9	147.7	148.2	169.3
SME	Base	96.6	112.6	67.4	171.7	73.2	54.9	66.1	78.9	77.5	88.8	153.5	109.8	91.7	116.7	99.6	104.4	156.8	116.1	122.3	119.0
	XIA	93.6	112.6	65.8	169.0	73.1	53.8	64.7	76.9	73.4	87.0	155.7	111.0	85.9	116.1	91.5	98.4	147.5	113.5	119.9	115.5
AME	Base	77.9	84.5	38.3	126.4	58.9	44.4	52.5	49.5	54.5	65.2	114.6	85.7	68.9	92.9	74.5	75.5	101.1	79.0	89.3	86.8
	XIA	74.4	85.0	37.7	124.0	56.7	42.6	51.5	47.4	50.5	63.3	115.1	84.0	62.2	90.4	67.4	71.5	93.9	76.2	86.1	83.0

in Table 4. More specially, the keypoints of the limbs (joints on the arms and legs) are improved largely. Arguably, the joints on the torso have little influence on the interaction between people, which comes mostly through the limbs. Hence, it is reasonable that using the cross-interaction gives the most significant boost in these joints. We also notice that the follower is more difficult to predict than the leader, as the poses associated to the follower are more extreme. Also, we observe that the lower legs are more difficult to estimate than the rest, and the very same reason holds. Furthermore, table 5 shows that the improvement is also systematic for different predicting horizons, both in short and long term.

EA and SA split. We push the analysis of the proposed XIA method by reporting the results on the **SA** and **EA** splits in Tables 6 and 7, respectively. For **SA** split, we notice that XIA outperforms the Base method also using the aerial-specific models. Regarding **EA** split, we can see that for the unseen actions, XIA still outperforms the Base method, demonstrating that our approach can better generalize to unseen aerials than the baseline model. More results and discussion can be found in the supplementary material.

Qualitative result. Figure 1 shows an example of our visualisation result comparing with the prediction results of the baseline model and the groundtruth. We can see that the poses estimated by XIA are obviously closer to the ground truth than those estimated by the Base method.

Ablations. Table 8 reports our ablation study. The ablation study is conducted on the **CA** split, and we report average values per dancer, on the three metrics. Additionally we

Table 8: **Ablation study results** (in *mm*). Different rows in the table correspond to: the base method (Base), Base method trained on the concatenation of the two poses (2P-cat), removing the residual connexion in XIA (XIA w/o res.), and using the XIA module with self-attention (XIA self-att) instead of cross-attention.

Dancer	Leader			Follower			Model
Metric	AME	JME	SME	AME	JME	SME	size
Base	88.2	62.6	88.2	126.0	88.2	160.6	6.4M
2P-cat	80.8	60.4	80.8	125.9	85.7	160.6	4.27M
XIA w/o res.	79.7	57.5	79.6	121.5	83.4	151.0	7.18M
XIA self-att.	78.6	56.8	78.7	124.6	85.2	152.2	7.18M
XIA	78.6	56.7	78.6	121.3	83.7	149.8	7.18M

report the model size. We demonstrate the interest of the design of XIA, as the three versions of XIA have the same amount of parameters but the proposed one is the best in performance, and our method could give a significant improvement to the baseline with a small adding complexity.

7. Conclusion

We have moved beyond existing approaches for 3D human motion prediction by considering an scenario with two people performing highly interactive activities. Current motion prediction formulations are restricted to a single person. In order to learn coupled motion dynamics, we have introduced a new mechanism that exploits historical information of both people in an attention-like fashion. This model is trained with ExPI, a new dataset we collected of professional dancers performing aerial steps. ExPI is annotated

with sequences of 3D body poses and shape, opening the door to not only being applied for interactive motion prediction but also for single frame pose estimation or multiview 3D reconstruction. The results of cross attention mechanism show consistent improvement compared to baselines that independently predict the motion of each person.

References

- [1] Emre Aksan, Peng Cao, Manuel Kaufmann, and Otmar Hilliges. A spatio-temporal transformer for 3d human motion prediction. *arXiv e-prints*, pages arXiv–2004, 2020. 2, 3
- [2] Emad Barsoum, John Kender, and Zicheng Liu. Hp-gan: Probabilistic 3d human motion prediction via gan. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1418–1427, 2018. 2
- [3] Abdallah Benzine, Bertrand Luvison, Quoc Cuong Pham, and Catherine Achard. Deep, robust and single shot 3d multi-person human pose estimation from monocular images. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 584–588. IEEE, 2019. 2
- [4] Judith Butepage, Michael J Black, Danica Kragic, and Hedvig Kjellström. Deep representation learning for human motion prediction and classification. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6158–6166, 2017. 3
- [5] Judith Bütepage, Hedvig Kjellström, and Danica Kragic. Anticipating many futures: Online human motion prediction and generation for human-robot interaction. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 4563–4570. IEEE, 2018. 3
- [6] Zhe Cao, Hang Gao, Karttikeya Mangalam, Qi-Zhi Cai, Minh Vo, and Jitendra Malik. Long-term human motion prediction with scene context. In *European Conference on Computer Vision*, pages 387–404. Springer, 2020. 3
- [7] Enric Corona, Albert Pumarola, Guillem Alenya, and Francesc Moreno-Noguer. Context-aware human motion prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6992–7001, 2020. 2, 3
- [8] Rishabh Dabral, Nitesh B Gundavarapu, Rahul Mitra, Abhishek Sharma, Ganesh Ramakrishnan, and Arjun Jain. Multi-person 3d human pose estimation from monocular images. In *2019 International Conference on 3D Vision (3DV)*, pages 405–414. IEEE, 2019. 2
- [9] Mihai Fieraru, Mihai Zanfir, Elisabeta Oneata, Alin-Ionut Popa, Vlad Olaru, and Cristian Sminchisescu. Three-dimensional reconstruction of human interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7214–7223, 2020. 3
- [10] Mihai Fieraru, Mihai Zanfir, Elisabeta Oneata, Alin-Ionut Popa, Vlad Olaru, and Cristian Sminchisescu. Three-dimensional reconstruction of human interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7214–7223, 2020. 3
- [11] Katerina Fragkiadaki, Sergey Levine, Panna Felsen, and Jitendra Malik. Recurrent network models for human dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4346–4354, 2015. 2
- [12] Partha Ghosh, Jie Song, Emre Aksan, and Otmar Hilliges. Learning human motion models for long-term predictions. In *2017 International Conference on 3D Vision (3DV)*, pages 458–466. IEEE, 2017. 2
- [13] John C Gower. Generalized procrustes analysis. *Psychometrika*, 40(1):33–51, 1975. 7
- [14] CMU graphics lab. Cmu graphics lab motion capture database., 2009. <http://mocap.cs.cmu.edu/>. 3
- [15] Liang-Yan Gui, Yu-Xiong Wang, Xiaodan Liang, and José MF Moura. Adversarial geometry-aware human motion prediction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 786–803, 2018. 2
- [16] Wen Guo, Enric Corona, Francesc Moreno-Noguer, and Xavier Alameda-Pineda. Pi-net: Pose interacting network for multi-person monocular 3d pose estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2796–2806, 2021. 2, 3
- [17] Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J Black. Resolving 3d human pose ambiguities with 3d scene constraints. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2282–2292, 2019. 3
- [18] Alejandro Hernandez, Jurgen Gall, and Francesc Moreno-Noguer. Human motion prediction via spatio-temporal inpainting. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7134–7143, 2019. 2, 3
- [19] Daniel Holden, Jun Saito, Taku Komura, and Thomas Joyce. Learning motion manifolds with convolutional autoencoders. In *SIGGRAPH Asia 2015 Technical Briefs*, pages 1–4. 2015. 3
- [20] 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. 3, 6
- [21] Ashesh Jain, Amir R Zamir, Silvio Savarese, and Ashutosh Saxena. Structural-rnn: Deep learning on spatio-temporal graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5308–5317, 2016. 2
- [22] Wen Jiang, Nikos Kolotouros, Georgios Pavlakos, Xiaowei Zhou, and Kostas Daniilidis. Coherent reconstruction of multiple humans from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5579–5588, 2020. 3
- [23] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 5
- [24] Laxman Kumarapu and Prerana Mukherjee. Animepose: Multi-person 3d pose estimation and animation. *arXiv preprint arXiv:2002.02792*, 2020. 2
- [25] Anliang Li, Shuang Wang, Wenzhu Li, Shengnan Liu, and Siyuan Zhang. Predicting human mobility with federated

- learning. In *Proceedings of the 28th International Conference on Advances in Geographic Information Systems*, pages 441–444, 2020. 2
- [26] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5226–5234, 2018. 3
- [27] Jiefeng Li, Can Wang, Wentao Liu, Chen Qian, and Cewu Lu. Hmor: Hierarchical multi-person ordinal relations for monocular multi-person 3d pose estimation. *arXiv preprint arXiv:2008.00206*, 2020. 3
- [28] Yong-Lu Li, Siyuan Zhou, Xijie Huang, Liang Xu, Ze Ma, Hao-Shu Fang, Yan-Feng Wang, and Cewu Lu. Transferable interactiveness prior for human-object interaction detection. *arXiv preprint arXiv:1811.08264*, 2018. 2
- [29] Jun Liu, Amir Shahroudy, Mauricio Perez, Gang Wang, Ling-Yu Duan, and Alex C Kot. Ntu rgb+ d 120: A large-scale benchmark for 3d human activity understanding. *IEEE transactions on pattern analysis and machine intelligence*, 42(10):2684–2701, 2019. 3
- [30] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5442–5451, 2019. 3
- [31] Wei Mao, Miaomiao Liu, and Mathieu Salzmann. History repeats itself: Human motion prediction via motion attention. In *European Conference on Computer Vision*, pages 474–489. Springer, 2020. 2, 3, 5, 7
- [32] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9489–9497, 2019. 2, 3
- [33] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2891–2900, 2017. 2
- [34] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Mohamed Elgharib, Pascal Fua, Hans-Peter Seidel, Helge Rhodin, Gerard Pons-Moll, and Christian Theobalt. Xnect: Real-time multi-person 3d human pose estimation with a single rgb camera. *arXiv preprint arXiv:1907.00837*, 2019. 2
- [35] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *2018 International Conference on 3D Vision (3DV)*, pages 120–130. IEEE, 2018. 2
- [36] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *3D Vision (3DV), 2018 Sixth International Conference on*. IEEE, sep 2018. 3
- [37] Terry Monaghan. Why study the lindy hop? *Dance Research Journal*, 33(2):124–127, 2001. 2
- [38] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10133–10142, 2019. 2
- [39] Francesc Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2823–2832, 2017. 2
- [40] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net: Localization-classification-regression for human pose. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3433–3441, 2017. 2
- [41] Grégory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net++: Multi-person 2d and 3d pose detection in natural images. *IEEE transactions on pattern analysis and machine intelligence*, 2019. 2
- [42] Caner Sahin, Guillermo Garcia-Hernando, Juil Sock, and Tae-Kyun Kim. Instance-and category-level 6d object pose estimation. In *RGB-D Image Analysis and Processing*, pages 243–265. Springer, 2019. 2
- [43] Amir Shahroudy, Jun Liu, Tian-Tsong Ng, and Gang Wang. Ntu rgb+ d: A large scale dataset for 3d human activity analysis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1010–1019, 2016. 3
- [44] Leonid Sigal, Alexandru O Balan, and Michael J Black. Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *International journal of computer vision*, 87(1-2):4, 2010. 3
- [45] Yongyi Tang, Lin Ma, Wei Liu, and Wei-Shi Zheng. Long-term human motion prediction by modeling motion context and enhancing motion dynamics. In *IJCAI*, 2018. 2
- [46] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017. 2
- [47] Timo von Marcard, Roberto Henschel, Michael J Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *Proceedings of the European Conference on Computer Vision*, pages 601–617, 2018. 2, 3
- [48] Andrei Zanfir, Elisabeta Marinoiu, and Cristian Sminchisescu. Monocular 3d pose and shape estimation of multiple people in natural scenes-the importance of multiple scene constraints. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2148–2157, 2018. 3
- [49] Xucong Zhang, Seonwook Park, Thabo Beeler, Derek Bradley, Siyu Tang, and Otmar Hilliges. Eth-xgaze: A large scale dataset for gaze estimation under extreme head pose and gaze variation. In *European Conference on Computer Vision*, pages 365–381. Springer, 2020. 3