

We are More than Our Joints: Predicting how 3D Bodies Move

Yan Zhang¹, Michael J. Black², Siyu Tang¹

¹ ETH Zürich, Switzerland

²Max Planck Institute for Intelligent Systems, Tübingen, Germany

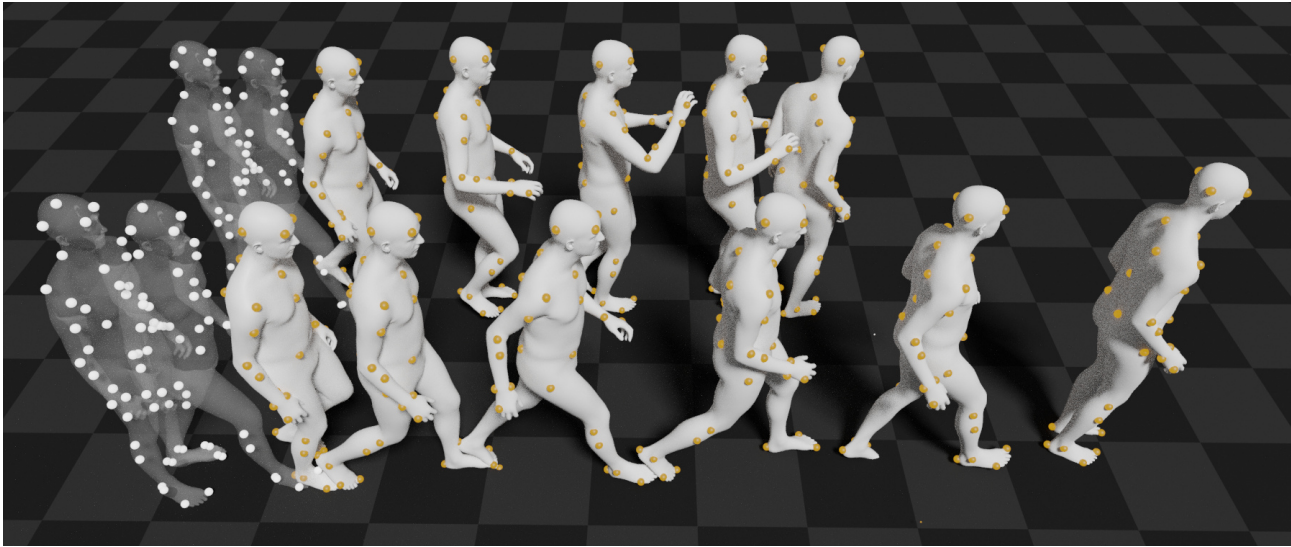


Figure 1: We represent the body in motion with a set of 3D markers on the body surface. Given a time sequence of markers from the past (white), MOJO predicts diverse marker sequences in the future (orange) with 3D bodies they represent (gray).

Abstract

A key step towards understanding human behavior is the prediction of 3D human motion. Successful solutions have many applications in human tracking, HCI, and graphics. Most previous work focuses on predicting a time series of future 3D joint locations given a sequence 3D joints from the past. This Euclidean formulation generally works better than predicting pose in terms of joint rotations. Body joint locations, however, do not fully constrain 3D human pose, leaving degrees of freedom (like rotation about a limb) undefined. Note that 3D joints can be viewed as a sparse point cloud. Thus the problem of human motion prediction can be seen as a problem of point cloud prediction. With this observation, we instead predict a sparse set of locations on the body surface that correspond to motion capture markers. Given such markers, we fit a parametric body model to recover the 3D body of the person. These sparse surface markers also carry detailed information about human movement that is not present in the joints, increasing the naturalness of the predicted motions. Using the AMASS dataset, we train MOJO (More than Our JOints), which is a novel variational autoencoder with a la-

tent DCT space that generates motions from latent frequencies. MOJO preserves the full temporal resolution of the input motion, and sampling from the latent frequencies explicitly introduces high-frequency components into the generated motion. We note that motion prediction methods accumulate errors over time, resulting in joints or markers that diverge from true human bodies. To address this, we fit the SMPL-X body model to the predictions at each time step, projecting the solution back onto the space of valid bodies, before propagating the new markers in time. Quantitative and qualitative experiments show that our approach produces state-of-the-art results and realistic 3D body animations. The code is available for research purposes at <https://yz-cnsdqz.github.io/MOJO/MOJO.html>.

1. Introduction

Human motion prediction has been extensively studied as a way to understand and model human behavior. Provided the recent past motion of a person, the goal is to predict either a deterministic or diverse set of plausible motions in the near future. While useful for animation, AR, and VR, predicting human movement is much more valuable be-

cause it means we have a *model* of how people move. Such a model is useful for applications in sports [59], pedestrian tracking [46], smart user interfaces [53], robotics [31] and more. While this is a variant of the well-studied time-series prediction problem, most existing methods are still not able to produce realistic 3D body motion.

To address the gap in realism, we make several novel contributions but start with a few observations. First, most existing methods for 3D motion prediction treat the body as a skeleton and predict a small set of 3D joints. While some methods represent the skeleton in terms of joint angles, the most accurate methods simply predict the 3D joint locations in Euclidean space. Second, given a sparse set of joint locations, animating a full 3D body is ambiguous because important degrees of freedom are not modeled, e.g. rotation about limb axes. Third, most papers show qualitative results by rendering skeletons and these often look fine to the human eye. We show, however, that, as time progresses, the skeletons can become less and less human in proportion so that, at the end of the sequence, the skeleton rarely corresponds to a valid human body. Fourth, the joints of the body cannot capture the nuanced details of how the surface of the body moves, limiting realism of any resulting animation.

We address these issues with a solution, called **MOJO (More than Our JOints)** the predicts realistic 3D body motion. MOJO incorporates a novel representation of the body in motion, a novel motion generative network, and a novel scheme for 3D body mesh recovery.

First, the set of 3D joints predicted by existing methods can be viewed as a sparse *point cloud*. In this light, existing human motion prediction methods preform *point cloud prediction*. Thus, we are free to choose a different point cloud that better satisfies the ultimate goal of animating bodies. Specifically, we model the body with a sparse set of surface markers corresponding to those used in motion capture (mocap) systems. We simply swap one type of sparse point cloud for another, but, as we will show, predicting surface markers has key advantages. For example, there exist methods to fit a SMPL body model [33] to such markers, producing realistic animations [32, 34]. Consequently this shift to predicting markers enables us to (1) leverage a powerful statistical body shape model to improve results, (2) immediately gives us realistic animations, (3) provides an output representation that can be used in many applications.

Second, to model fine-grained and high-frequency interactions between markers, we design a conditional variational autoencoder (CVAE) with a latent cosine space. It not only performs stochastic motion prediction, but also improves motion realism by incorporating high-frequency motion details. Compared to most existing methods that encode motion with a single vector (e.g. the last hidden state of an RNN), our model preserves full temporal resolution of the sequence, and decomposes motion into independent fre-

quency bands in the latent space via a discrete cosine transform (DCT). Based on the energy compaction property of the DCT [3, 44]¹, we train our CVAE with a robust Kullback–Leibler divergence (KLD) term [60], creating an implicit latent prior that carries most of the information at low frequency bands. To sample from this implicit latent prior, we employ diversifying latent flows (DLows) [56] in low-frequency bands to produce informative features, and from the standard normal distribution in high-frequency bands to produce white noise. Pieces of information from various frequencies are then fused to compose the output motion.

Third, in the inference phase, we propose a recursive projection scheme supported by our marker-based representation, in order to retain natural body shape and pose throughout the sequence. We regard the valid body space as a low-dimensional manifold in the Euclidean space of markers. When the CVAE decoder makes a prediction step, the predicted markers tend to leave this manifold because of error accumulation. Therefore, after each step we project the predicted markers back to the valid body manifold, by fitting an expressive SMPL-X [40] body mesh to the markers. On the fitted body, we know the true marker locations and pass these to the next stage of the RNN, effectively denoising the markers at each time instant. Besides keeping the solution valid, the recursive projection scheme directly yields body model parameters and hence realistic body meshes.

We exploit the **AMASS** [34] dataset for evaluation, as well as **Human3.6M** [25] and **HumanEva-I** [47] to compare our methods with the state-of-the-art in stochastic motion prediction. We show that our models with the latent DCT space outperform the state-of-the-art, and that the recursive projection scheme effectively eliminates unrealistic body deformation. We also evaluate realism of the generated motion with a foot skating measure and a perceptual study. Finally, we compare different body representations, in particular our solution with a traditional pipeline, which first predicts 3D joints and then fits a body to the joints. We find that they are comparable w.r.t. prediction diversity and accuracy, but the traditional pipeline can produce invalid body shapes.

Contributions. In summary, our contributions are: (1) We propose a marker-based representation for bodies in motion, which provides more constraints than the body skeleton and hence benefits 3D body recovery. (2) We design a new CVAE with a latent DCT space to improve motion modelling. (3) We propose a recursive projection scheme to preserve valid bodies at test time.

2. Related Work

Deterministic human motion prediction. Given an input human motion sequence, the goal is to forecast a de-

¹Most information will be concentrated at low frequency bands.

terministic future motion, which is expected to be close to the ground truth. This task has been extensively studied [5, 6, 10, 13, 15, 16, 17, 18, 28, 29, 35, 36, 41, 42, 48, 50, 62]. Martinez et al. [36] propose an RNN with residual connections linking the input and the output, and design a sampling-based loss to compensate for prediction errors during training. Cai et al. [10] and Mao et al. [35] use the discrete cosine transform (DCT) to convert the motion into the frequency domain. Then Mao et al. [35] employ graph convolutions to process the frequency components, whereas Cai et al. [10] use a transformer-based architecture.

Stochastic 3D human motion prediction. In contrast to deterministic motion prediction, stochastic motion prediction produces diverse plausible future motions, given a single motion from the past [7, 8, 14, 19, 30, 49, 55, 56, 60]. Yan et al. [55] propose a motion transformation VAE to jointly learn the motion mode feature and transition between motion modes. Barsoum et al. [7] propose a probabilistic sequence-to-sequence model, which is trained with a Wasserstein generative adversarial network. Bhattacharyya et al. [8] design a ‘best-of-many’ sampling objective to boost the performance of conditional VAEs. Gurumurthy et al. [19] propose a GAN-based network and parameterize the latent generative space as a mixture model. Yuan et al. [56] propose diversifying latent flows (DLow) to exploit the latent space of an RNN-based VAE, which generates highly diverse but accurate future motions.

Frequency-based motion analysis. Earlier studies like [38, 43] adopt a Fourier transform for motion synthesis and tracking. Akhter et al. [4] propose a linear basis model for spatiotemporal motion regularization, and discover that the optimal PCA basis of a large set of facial motion converges to the DCT basis. Huang et al. [24] employ low-frequency DCT bands to regularize motion of body meshes recovered from 2D body joints and silhouettes. The studies of [10, 35, 50] use deep neural networks to process DCT frequency components for motion prediction. Yumer et al. [58] and Holden et al. [23] handle human motions in the Fourier domain to conduct motion style transfer.

Representing human bodies in motion. 3D joint locations are widely used, e.g. [29, 36, 56]. To improve prediction accuracy, Mao et al. [35], Cui et al. [13], Li et al. [29] and others use a graph to capture interaction between joints. Askan et al. [6] propose a structured network layer to represent the body joints according to a kinematic tree. Despite their effectiveness, the skeletal bone lengths can vary during motion. To alleviate this issue, Hernandez et al. [22] use a training loss to penalize bone length variations. Gui et al. [17] design a geodesic loss and two discriminators to keep the predicted motion human-like over time. To remove the influence of body shape, Pavllo et al. [41, 42] use quaternion-based joint rotations to represent the body pose. When predicting the global motion, a walking path is first

produced and the pose sequence is then generated. Zhang et al. [60] represent the body in motion by the 3D global translation and the joint rotations following the SMPL kinematic tree [33]. When animating a body mesh, a constant body shape is added during testing. Although a constant body shape is preserved, foot skating frequently occurs due to the inconsistent relation between the body pose, the body shape and the global translation.

MOJO in context. Our solution not only improves stochastic motion prediction over the state-of-the-art, but also directly produces diverse future motions of realistic 3D bodies. Specifically, our latent DCT space represents motion with different frequencies, rather than a single vector in the latent space. We find that generating motions from different frequency bands significantly improves diversity while retaining accuracy. Additionally, compared to previous skeleton-based body representations, we propose to use markers on the body surface to provide more constraints on the body shape and DoFs. This marker-based representation enables us to design an efficient recursive projection scheme by fitting SMPL-X [40] at each prediction step. Recently, in the context of autonomous driving, Weng et al. [51, 52] forecast future LiDAR point clouds and then detect-and-track 3D objects in the predicted clouds. While this has similarities to MOJO, they do not address articulated human movements.

3. Method

3.1. Preliminaries

SMPL-X body mesh model [40]. Given a compact set of body parameters, SMPL-X produces a realistic body mesh including face and hand details. In our work, the body parameter set Θ includes the global translation $\mathbf{t} \in \mathbb{R}^3$, the global orientation $\mathbf{R} \in \mathbb{R}^6$ w.r.t. the continuous representation [61], the body shape $\beta \in \mathbb{R}^{10}$, the body pose $\theta \in \mathbb{R}^{32}$ in the VPoser latent space [40], and the hand pose $\theta^h \in \mathbb{R}^{24}$ in the MANO [45] PCA space. We denote a SMPL-X mesh as $\mathcal{M}(\Theta)$, which has a fixed topology with 10,475 vertices. MOJO is implemented using SMPL-X in our work, but any other parametric 3D body model could be used, e.g. [39, 54].

Diversifying latent flows (DLow) [56]. The entire DLow method comprises a CVAE to predict future motions, and a network \mathbf{Q} that takes the condition sequence as input and transforms a sample $\varepsilon \sim \mathcal{N}(0, \mathbf{I})$ to K diverse places in the latent space. To train the CVAE, the loss consists of a frame-wise reconstruction term, a term to penalize the difference between the last input frame and the first predicted frame, and a KLD term. Training the network \mathbf{Q} requires a pre-trained CVAE decoder, their training loss encourages diverse generated motion samples in which, at least one sample is close to the ground truth motion.

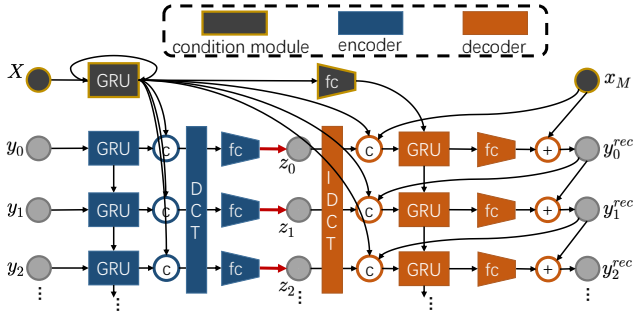


Figure 2: Illustration of our CVAE architecture. The red arrows denote sampling from the inference posterior. The circles with ‘c’ and ‘+’ denote feature concatenation and addition, respectively. The blocks with ‘fc’ denote a stack of fully-connected layers.

3.2. Human Motion Representation

Most existing methods use 3D joint locations or rotations to represent the body in motion. This results in ambiguities in recovering the full shape and pose [9, 40]. To obtain more constraints on the body, while preserving efficiency, we represent bodies in motion with markers on the body surface. Inspired by modern mocap systems, we follow the marker placements of either the CMU mocap dataset [1] or the SSM2 dataset [34], and select V vertices on the SMPL-X body mesh, which are illustrated in the Appendix.

The 3D markers are simply points in Euclidean space. Compared to representing the body with the global translation and the local pose, like in [42, 60], such a location-based representation naturally couples the global body translation and the local pose variation, and hence is less prone to motion artifacts like foot skating, which are caused by the mismatch between the global movement, the pose variation, and the body shape.

Therefore, in each frame the body is represented by a V -dimensional feature vector, i.e. the concatenation of the 3D locations of the markers, and the motion is represented by a time sequence of such vectors. We denote the input sequence to the model as $\mathbf{X} := \{\mathbf{x}_t\}_{t=0}^M$, and a predicted future sequence from the model as $\mathbf{Y} := \{\mathbf{y}_t\}_{t=0}^N$, where $\mathbf{y}_0 = \mathbf{x}_{M+1}$. With fitting algorithms like MoSh and MoSh++ [32, 34], it is much easier to recover 3D bodies from markers than from joints.

3.3. Motion Generator with Latent Frequencies

For a real human body, the motion granularity usually corresponds to the motion frequency because of the underlying locomotor system. For example, the frequency of waving hands is usually much higher than the frequency of a jogging gait. Therefore, we design a network with multiple frequency bands in the latent space, so as to better represent interactions between markers on the body surface and

to model motions at different granularity levels.

Architectures. Our model is visualized in Fig. 2, which is designed according to the CVAE in the DLow method [56]. The encoder with a GRU [12] preserves full temporal resolution of the input. Then, the motion information is decomposed onto multiple frequency bands via DCT. At individual frequency bands, we use the re-parameterization trick [26] to introduce randomness, and then use inverse DCT to convert the motion back to the temporal domain. To improve temporal smoothness and eliminate the first-frame jump reported in [36], we use residual connections at the output. We note that the CVAE in the DLow method, which does not have residual connections but has a loss to penalize the jump artifact, does not produce smooth and realistic marker motions. A visualization of this latent DCT space is shown in the Appendix.

Training with robust Kullback-Leibler divergence. Our training loss comprises three terms for frame-wise reconstruction, frame-wise velocity reconstruction, and latent distribution regularization, respectively;

$$\mathcal{L} = \mathbb{E}_{\mathbf{Y}}[\|\mathbf{Y} - \mathbf{Y}^{rec}\|] + \alpha \mathbb{E}_{\mathbf{Y}}[\|\Delta \mathbf{Y} - \Delta \mathbf{Y}^{rec}\|] + \Psi(KLD(q(\mathbf{Z}|\mathbf{X}, \mathbf{Y})||\mathcal{N}(0, \mathbf{I}))), \quad (1)$$

where the operation Δ computes the time difference, $q(\cdot|\cdot)$ denotes the inference posterior (the encoder), \mathbf{Z} denotes the latent frequency components, and α is a loss weight. We find the velocity reconstruction term can further improve temporal smoothness.

Noticeably, our distribution regularization term is given by the robust KLD [60] with $\Psi(s) = \sqrt{1 + s^2} - 1$ [11], which defines an implicit latent prior different from the standard normal distribution. During optimization, the gradients to update the entire KLD term become smaller when the divergence from $\mathcal{N}(0, \mathbf{I})$ becomes smaller. Thus, it expects the inference posterior to carry information, and prevents posterior collapse. More importantly, this term is highly suitable for our latent DCT space. According to the energy compaction property of DCT [3, 44], we expect that the latent prior deviates from $\mathcal{N}(0, \mathbf{I})$ at low-frequency bands to carry information, but equals $\mathcal{N}(0, \mathbf{I})$ at high-frequency bands to produce white noise. We let this robust KLD term determine which frequency bands to carry information automatically.

Sampling from the latent DCT space. Since our latent prior is implicit, sampling from the latent space is not as straightforward as sampling from the standard normal distribution, like for most VAEs. Due to the DCT nature, we are aware that motion information is concentrated at low-frequency bands, and hence we directly explore these informative low-frequency bands using the network \mathbf{Q} in DLow [56].

Specifically, we use $\{\mathbf{Q}_w\}_{w=1}^L$ to sample from the lowest L frequency bands, and sample from $\mathcal{N}(0, \mathbf{I})$ from $L + 1$




Figure 3: Illustration of our sampling scheme to generate K different sequences. Q_w denotes the network to produce the latent transformation T_w at the frequency band w . The red arrows denote the sampling operation.




Figure 4: Illustration of our prediction scheme with projection. The notation has the same meaning as before.

to the highest frequency bands. Since the cosine basis is orthogonal and individual frequency bands carry independent information, these L DLow models do not share parameters, but are jointly trained with the same losses as in [56], as well as the decoder of our CVAE with the latent DCT space. Our sampling approach is illustrated in Fig. 3. The influence of the threshold L is investigated in Sec. 4 and in the Appendix.

3.4. Recursive Projection to the Valid Body Space

Our generative model produces diverse motions in terms of marker location variations. Due to RNN error accumulation, predicted markers can gradually deviate from a valid 3D body, resulting in, e.g., flattened heads and twisted torsos. Existing methods with new losses or discriminators [17, 22, 27] can alleviate this problem, but may unpredictably fail due to the train-test domain gap.

Instead, we exploit the fact that valid bodies lie on a low-dimensional manifold in the Euclidean space of markers. Whenever the RNN performs a prediction step, the solution tends to leave this manifold. Therefore, at each prediction step, we project the predicted markers back to that manifold, by fitting a SMPL-X body mesh to the predicted markers. Since markers provide rich body constraints, and we start close to the solution, the fitting process is efficiently. Our recursive projection scheme is illustrated in Fig. 4. Note that we only apply recursive projection at the inference stage.

Following the work of MoSh [32] and MoSh++ [34], the fitting is optimization-based, and consists of three stages:

(1) optimizing the global configurations t and R , (2) additionally optimizing the body pose θ , and (3) additionally optimizing the hand pose θ^h . At each time t , we use the previous fitted result to initialize the current optimization process, so that the optimum can be reached with a small number of iterations. The loss of our optimization-based fitting at time t is given by

$$\mathcal{L}_f(\Theta_t) := |V_{\mathcal{M}(\Theta_t)} - \mathbf{y}_t^{pred}|^2 + \lambda_1 |\theta_t|^2 + \lambda_2 |\theta_t^h|^2, \quad (2)$$

in which λ s are the loss weights, V denotes the corresponding markers on the SMPL-X body mesh, and \mathbf{y}_t^{pred} denotes the markers predicted by the CVAE decoder. The recursive projection uses the body shape from the input sequence, and runs at 2.27sec/frame on average in our trials, which is comparable to the pose stage of MoSh++. From our recursive projection scheme, we not only obtain regularized markers, but also realistic 3D bodies as well as their characteristic parameters.

4. Experiment

MOJO has several components that we evaluate. First, to test the effectiveness of the MOJO CVAE architecture, in particular the benefits of the latent DCT space, we evaluate stochastic motion prediction in Sec. 4.3. Second, to test the effectiveness of MOJO with recursive projection, we evaluate the performance of MOJO w.r.t. realism of 3D body movements in Sec. 4.4. Finally, to test the advantage of the marker-based representation, we systematically compare different body representations in Sec. 4.5. We find that MOJO produces diverse realistic 3D body motions and outperforms the state-of-the-art.

4.1. Datasets

For training, we use the AMASS [34] dataset. Specifically, we train the models on CMU [1] and MPI HDM05 [37], and test models on ACCAD [2] and BML-handball [21, 20]. This gives 696.1 minutes of training motion from 110 subjects and 128.72 minutes of test motion from 30 subjects. The test sequences include a wide range of actions. To unify the sequence length and world coordinates, we canonicalize AMASS sequences as a pre-processing step. Details are in the Appendix.

To compare our method with SOTA stochastic motion prediction methods, we additionally perform skeleton-based motion prediction on the Human3.6M dataset [25] and the HumanEva-I dataset [47], following the experimental setting of Yuan et al. [56].

4.2. Baselines

MOJO predicts surface markers, and has several components. Unless mentioned, we use the CMU layout with 41 markers. ‘MOJO-DCT’ is the model without DCT, but

with the same latent space as the CVAE in DLow. ‘MOJO-proj’ is the model without recursive projection. Note that the suffixes can be concatenated; e.g. ‘MOJO-DCT-proj’ is the model without the latent DCT space and without the recursive projection scheme.

4.3. Evaluation of Stochastic Motion Prediction

4.3.1 Metrics

Diversity. We use the same diversity measure as [56], which is the average pair-wise distance between all generated sequences.

Prediction accuracy. As in [56], we use the average distance error (*ADE*) and the final distance error (*FDE*) to measure the minimum distance between the generated motion and the ground truth, w.r.t. frame-wise difference and the final frame difference, respectively. Additionally, we use *MMADE* and *MMFED* to evaluate prediction accuracy when the input sequence slightly changes; see Appendix.

Motion Frequency. Similar to [22], we compute the frequency spectra entropy (*FSE*) to measure motion frequency in the Fourier domain, which is given by the averaged spectra entropy minus the ground truth. A higher value indicates the generated motions contain more motion detail. Note that high frequency can also indicate noise, and hence this metric is jointly considered with the prediction accuracy.

4.3.2 Results

We generate 50 different future sequences based on each input sequence, as in [56]. Here we focus only on evaluating performance on motion prediction, and hence do not incorporate the body re-projection scheme. The results are shown in Tab. 1, in which we employ DLow on the 20% (i.e. the first 9) lowest frequency bands in ‘MOJO-proj’. We find that DCT consistently leads to better performance. Noticeably, higher motion frequency indicates that the generated motions contain more details, and hence are more realistic.

To further investigate the benefits of our latent DCT space, we add the latent DCT space into the DLow CVAE model [56] and train it with the robust KLD term. For sampling, we apply a set of DLow models $\{\mathbf{Q}_w\}$ on the lowest L bands, as in Sec. 3.3. We denote this modified model as ‘VAE+DCT+ L ’. Absence of the suffix ‘+ L ’ indicates sampling from $\mathcal{N}(0, \mathbf{I})$ in all frequency bands.

A comparison with existing methods is shown in Tab. 2. Overall, our latent DCT space effectively improves on the state-of-the-art. The diversity is improved by a large margin, while the prediction accuracies are comparable to the baseline. The performance w.r.t. *MMADE* and *MMFDE* is slightly inferior. A probable reason is VAE+DCT uses high-frequency components to generate motions, which makes

motion prediction sensitive to slight changes of the input. Moreover, by comparing ‘VAE+DCT’ and ‘VAE+DCT+ L ’, we can see that sampling from $\mathcal{N}(0, \mathbf{I})$ yields much worse results. This indicates that sampling from the standard normal distribution, which treats all frequency bands equally, cannot effectively exploit the advantage of the latent DCT space. Note that most information is in the low-frequency bands (see Appendix), and hence our proposed sampling method utilizes the latent frequency space in a more reasonable way and produces better results.

4.4. Evaluation of Motion Realism

4.4.1 Metrics

The motion prediction metrics cannot indicate whether a motion is realistic or not. Here, we employ MOJO with recursive projection to obtain 3D body meshes, and evaluate the body deformation, foot skating, and perceptual quality.

Body deformation. Body shape can be described by the pairwise distances between markers. As a body moves, there is natural variation in these distances. Large variations, however, indicate a deformed body that no longer corresponds to any real person. We use variations in pairwise marker distances for the head, upper torso, and lower torso as a measure of how distorted the predicted body is. See Appendix for the metric details.

Foot skating ratio. Foot skating is measured based on the two markers on the left and right foot calcaneus (‘LHEE’ and ‘RHEE’ in CMU [1]). We consider foot skating to have occurred, when both foot markers are close enough to the ground (within 5cm) and simultaneously exceed a speed limit (5mm between two consecutive frames or 75mm/s). We report the averaged ratio of frames with foot skating.

Perceptual score. We render the generated body meshes as well as the ground truth, and perform a perceptual study on Amazon Mechanical Turk. Subjects see a motion sequence and the statement “The human motion is natural and realistic.” They evaluate this on a six-point Likert scale from ‘strongly disagree’ (1) to ‘strongly agree’ (6). Each individual video is rated by three subjects. We report mean values and standard deviations for each method and each dataset.

4.4.2 Results

We randomly choose 60 different sequences from **ACCAD** and **BMLhandball**, respectively. Based on each sequence, we generate 15 future sequences. Figure 5 shows some generated 3D body motions. The motions generated by MOJO contain finer-grained body movements.

Body deformation. The results are shown in Tab. 3. With the recursive projection scheme, the body shape is preserved by construction and is close to the ground truth. Without the projection scheme, the shape of body parts can

Method	ACCAD [2]						BMLHandball [21, 20]					
	Diversity \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow	FSE \uparrow	Diversity \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow	FSE \uparrow
MOJO-DCT-proj	25.349	1.991	3.216	2.059	3.254	0.4	21.504	1.608	1.914	1.628	1.919	0.0
MOJO-proj	28.886	1.993	3.141	2.042	3.202	1.2	23.660	1.528	1.848	1.550	1.847	0.4

Table 1: Comparison between generative models for predicting marker-based motions. The symbol \downarrow (or \uparrow) denotes whether results that are lower (or higher) are better, respectively. Best results of each model are in boldface. The FSE scores are on the scale of 10^{-3} .

Method	Human3.6M [25]					HumanEva-I [47]				
	Diversity \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow	Diversity \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow
Pose-Knows [49]	6.723	0.461	0.560	0.522	0.569	2.308	0.269	0.296	0.384	0.375
MT-VAE [55]	0.403	0.457	0.595	0.716	0.883	0.021	0.345	0.403	0.518	0.577
HP-GAN [7]	7.214	0.858	0.867	0.847	0.858	1.139	0.772	0.749	0.776	0.769
Best-of-Many [8]	6.265	0.448	0.533	0.514	0.544	2.846	0.271	0.279	0.373	0.351
GMVAE [14]	6.769	0.461	0.555	0.524	0.566	2.443	0.305	0.345	0.408	0.410
DeLiGAN [19]	6.509	0.483	0.534	0.520	0.545	2.177	0.306	0.322	0.385	0.371
DSF [57]	9.330	0.493	0.592	0.550	0.599	4.538	0.273	0.290	0.364	0.340
DLow [56]	11.730	0.425	0.518	0.495	0.532	4.849	0.246	0.265	0.360	0.340
VAE+DCT	3.462	0.429	0.545	0.525	0.581	0.966	0.249	0.296	0.412	0.445
VAE+DCT+5	12.579	0.412	0.514	0.497	0.538	4.181	0.234	0.244	0.369	0.347
VAE+DCT+20	15.920	0.416	0.522	0.502	0.546	6.266	0.239	0.253	0.371	0.345

Table 2: Comparison between a baseline with our latent DCT space and the state-of-the-art. Best results are in boldface.




Figure 5: Visualization of 3D body motions. Bodies in gray-green and red denote the input and the generated motion, respectively. The solid and dash image borders denote the results from MOJO-DCT and MOJO, respectively.

Method	ACCAD [2]			BMLHandball [21, 20]		
	Head	Up. T.	Low. T.	Head	Up. T.	Low. T.
MOJO-DCT-proj	76.3	102.2	99.4	86.0	105.3	83.3
MOJO-proj	70.3	80.7	76.7	68.3	77.7	63.2
MOJO-DCT	1.32	34.0	6.97	1.32	43.1	6.97
MOJO	1.30	32.7	6.76	1.40	44.3	7.64
Ground truth	2.17	36.4	9.78	2.52	59.1	12.5

Table 3: Deformations of body parts (head, Upper Torso, Lower Torso). Scores are in millimeter. High values indicate distorted bodies.

Method	ACCAD [2]		BMLHandball [21, 20]	
	foot skate	percep. score	foot skate	percep. score
MOJO-DCT	0.341	4.15 \pm 1.38	0.077	4.00 \pm 1.26
MOJO	0.278	4.07 \pm 1.31	0.066	4.17 \pm 1.23
Ground truth	0.067	4.82 \pm 1.08	0.002	4.83 \pm 1.05

Table 4: Comparison between methods w.r.t. foot skating and the perceptual score, which is given by mean \pm std. Best results are in boldface.

drift significantly from the true shape, indicated here by high deformation numbers. MOJO is close to the ground truth but exhibits less deformation suggesting that some nuance is smoothed out by the VAE.

Foot skating and perceptual score. The results are presented in Tab. 4. The model with DCT produces fewer foot skating artifacts, indicating that high-frequency components in the DCT space can better model the foot movements. In the perceptual study, MOJO performs slightly worse than MOJO-DCT on ACCAD, but outperforms it on BMLhandball. A probable reason is that most actions in ACCAD are coarse-grained, whereas most actions in BMLhandball are fine-grained. The advantage of modelling finer-grained motion of the DCT latent space is more easily perceived in BMLhandball.

4.5. Comparison between Body Representations

The body in motion can be represented by locations of joints, markers with different placements, and their combinations. Here we perform a systematic comparison between them. For the joint-based representation, we use the

Method	ACCAD [2]						BMLHandball [21, 20]					
	Diversity \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow	BDF \downarrow	Diversity \uparrow	ADE \downarrow	FDE \downarrow	MMADE \downarrow	MMFDE \downarrow	BDF \downarrow
joints w/o proj.	21.363	1.184	2.010	1.291	2.067	0.185	19.091	0.930	1.132	1.000	1.156	0.205
joints	21.106	1.192	2.022	1.299	2.076	0	18.954	0.934	1.138	1.003	1.157	0
CMU 41	20.676	1.214	1.919	1.306	2.080	0	16.806	0.949	1.139	1.001	1.172	0
SSM2 67	24.373	1.124	1.699	1.227	1.838	0	18.778	0.924	1.099	0.975	1.149	0
joints + CMU 41	20.988	1.187	1.841	1.308	1.967	0	13.982	0.943	1.190	0.990	1.194	0
joints + SSM2 67	23.504	1.166	1.892	1.276	1.953	0	16.483	0.950	1.146	0.999	1.189	0

Table 5: Comparison between marker-based and joint-based representations. Evaluations are based on the joint locations. *BDF* denotes the bone deformation w.r.t. meter. The best results are in boldface.

SMPL [33] joint locations from CMU and MPI HDM05 to train a CVAE as in MOJO. A traditional pipeline is to first predict all joints in the future, and then fit the body mesh. Here we also test the performance when applying the recursive body re-projection scheme based on joints. For fair quantitative evaluation, the metrics are calculated based on the joints of the fitted body meshes. We re-calculate the diversity, the prediction accuracy metrics, and the eight limb bone deformation (BDF) (according to Eq. (6) in Appendix) w.r.t. the joint locations.

We randomly choose 60 sequences from each test set and generate 50 future motions based on each sequence. Results are presented in Tab. 5. The first two rows show that the motion naturalness is improved by the recursive projection, which eliminates the bone deformation. Although the result without projection is slightly better on other measures, the projection scheme completely removes bone deformation and does not produce false poses as in Fig. 6. Additionally, using more markers (SSM2 placement with 67 markers) significantly improves performance across the board. This shows that the marker distribution is important for motion prediction and that more markers is better. While MOJO works with markers, joints, or the combination of both, the combination of joints and markers does not produce better performance. Note that the joints are never directly observed, but rather are inferred from the markers by commercial mocap systems. Hence, we argue that the joints do not add independent information.

Figure 6 shows the risk of the traditional joint-based pipeline. While the skeletons may look fine to the eye, in the last frame the character cannot be fit to the joints due to unrealistic bone lengths.

5. Conclusion

In this paper, we propose MOJO, a new method to predict diverse plausible motions of 3D bodies. Instead of using joints to represent the body, MOJO uses a sparse set of markers on the body surface, which better constrain 3D body shape and pose recovery. In contrast to most existing methods that encode a motion into a single feature vector, we represent motion with latent frequencies, which can describe fine-grained body movements and improve motion prediction consistently. To produce valid 3D bodies in mo-




Figure 6: Fitting a character to predicted joints. The top row and the bottom row show the predicted skeletons and the fitted bodies, respectively. From left to right: The first predicted frame, the middle frame, and the last frame.

tion, MOJO uses a recursive projection scheme at test time. By fitting a SMPL-X body to the predicted markers at each frame, the 3D bodies stay valid over time and motion realism is improved. Compared to a traditional pipeline based on joints, MOJO thoroughly eliminates implausible body deformations and produces realistic 3D body movements.

Nevertheless, MOJO has some limitations to improve in the future. For example, the recursive projection scheme slows down the inference process. Also, the motion realism is still not comparable with the ground truth (see Tab. 4), indicating room to improve. Moreover, we will explore the performance of MOJO on other marker settings, or even real markers from mocap data.

Acknowledgement: We thank Nima Ghorbani for the advice on the body marker setting and the AMASS dataset. We thank Yinghao Huang, Cornelia Köhler, Victoria Fernández Abrevaya, and Qianli Ma for proofreading. We thank Xinchun Yan and Ye Yuan for discussions on baseline methods. We thank Shaofei Wang and Siwei Zhang for their help with the user study and the presentation, respectively.

Disclaimer: MJB has received research gift funds from Adobe, Intel, Nvidia, Facebook, and Amazon. While MJB is a part-time employee of Amazon, his research was performed solely at, and funded solely by, Max Planck. MJB has financial interests in Amazon Datagen Technologies, and Meshcapade GmbH.

References

- [1] CMU graphics lab. CMU graphics lab motion capture database. <http://mocap.cs.cmu.edu/>, 2000. 4, 5, 6, 1, 3
- [2] ACCAD MoCap System and Data. <https://accad.osu.edu/research/motion-lab/systemdata>, 2018. 5, 7, 8
- [3] Nasir Ahmed, T. Natarajan, and Kamisetty R. Rao. Discrete cosine transform. *IEEE transactions on Computers*, 100(1):90–93, 1974. 2, 4
- [4] Ijaz Akhter, Tomas Simon, Sohaib Khan, Iain Matthews, and Yaser Sheikh. Bilinear spatiotemporal basis models. *ACM Transactions on Graphics*, 31(2):1–12, 2012. 3
- [5] Emre Aksan, Peng Cao, Manuel Kaufmann, and Otmar Hilliges. Attention, please: A spatio-temporal transformer for 3D human motion prediction. *arXiv preprint arXiv:2004.08692*, 2020. 3
- [6] Emre Aksan, Manuel Kaufmann, and Otmar Hilliges. Structured prediction helps 3D human motion modelling. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7144–7153, 2019. 3
- [7] Emad Barsoum, John Kender, and Zicheng Liu. HP-GAN: Probabilistic 3D human motion prediction via GAN. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 1418–1427, 2018. 3, 7
- [8] Apratim Bhattacharyya, Bernt Schiele, and Mario Fritz. Accurate and diverse sampling of sequences based on a “best of many” sample objective. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8485–8493, 2018. 3, 7
- [9] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J. Black. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *European Conference on Computer Vision*, Oct. 2016. 4
- [10] Yujun Cai, Lin Huang, Yiwei Wang, Tat-Jen Cham, Jianfei Cai, Junsong Yuan, Jun Liu, Xu Yang, Yiheng Zhu, Xiaohui Shen, Liu Ding, Liu Jing, and Nadia Magnenat-Thalmann. Learning progressive joint propagation for human motion prediction. In *European Conference on Computer Vision*, 2020. 3
- [11] Pierre Charbonnier, Laure Blanc-Feraud, Gilles Aubert, and Michel Barlaud. Two deterministic half-quadratic regularization algorithms for computed imaging. In *Proceedings of the International Conference on Image Processing*, pages 168–172, 1994. 4
- [12] Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 1724–1734, Doha, Qatar, Oct. 2014. 4
- [13] Qiongjie Cui, Huaijiang Sun, and Fei Yang. Learning dynamic relationships for 3D human motion prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6519–6527, 2020. 3
- [14] Nat Dilokthanakul, Pedro AM Mediano, Marta Garnelo, Matthew CH Lee, Hugh Salimbeni, Kai Arulkumaran, and Murray Shanahan. Deep unsupervised clustering with gaussian mixture variational autoencoders. *arXiv preprint arXiv:1611.02648*, 2016. 3, 7
- [15] Partha Ghosh, Jie Song, Emre Aksan, and Otmar Hilliges. Learning human motion models for long-term predictions. In *International Conference on 3D Vision*, pages 458–466. IEEE, 2017. 3
- [16] Anand Gopalakrishnan, Ankur Mali, Dan Kifer, Lee Giles, and Alexander G Ororbia. A neural temporal model for human motion prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 12116–12125, 2019. 3
- [17] Liang-Yan Gui, Yu-Xiong Wang, Xiaodan Liang, and José MF Moura. Adversarial geometry-aware human motion prediction. In *European Conference on Computer Vision*, pages 786–803, 2018. 3, 5
- [18] Liang-Yan Gui, Yu-Xiong Wang, Deva Ramanan, and José MF Moura. Few-shot human motion prediction via meta-learning. In *European Conference on Computer Vision*, pages 432–450, 2018. 3
- [19] Swaminathan Gurumurthy, Ravi Kiran Sarvadevabhatla, and R Venkatesh Babu. DeLiGAN: Generative adversarial networks for diverse and limited data. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 166–174, 2017. 3, 7
- [20] Fabian Helm, Rouwen Cañal-Bruland, David L Mann, Nikolaus F Troje, and Jörn Munzert. Integrating situational probability and kinematic information when anticipating disguised movements. *Psychology of Sport and Exercise*, 46:101607, 2020. 5, 7, 8
- [21] Fabian Helm, Nikolaus F Troje, and Jörn Munzert. Motion database of disguised and non-disguised team handball penalty throws by novice and expert performers. *Data in brief*, 15:981–986, 2017. 5, 7, 8
- [22] Alejandro Hernandez, Jurgen Gall, and Francesc Moreno-Noguer. Human motion prediction via spatio-temporal inpainting. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7134–7143, 2019. 3, 5, 6
- [23] Daniel Holden, Ikhsanul Habibie, Ikuro Kusajima, and Taku Komura. Fast neural style transfer for motion data. *IEEE computer graphics and applications*, 37(4):42–49, 2017. 3
- [24] Yinghao Huang, Federica Bogo, Christoph Lassner, Angjoo Kanazawa, Peter V Gehler, Javier Romero, Ijaz Akhter, and Michael J Black. Towards accurate marker-less human shape and pose estimation over time. In *International Conference on 3D Vision*, pages 421–430. IEEE, 2017. 3
- [25] 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, 2014. 2, 5, 7
- [26] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *International Conference on Learning Representations*, 2014. 4
- [27] Muhammed Kocabas, Nikos Athanasiou, and Michael J. Black. VIBE: Video inference for human body pose and

- shape estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5252–5262, 2020. 5
- [28] Chen Li, Zhen Zhang, Wee Sun Lee, and Gim Hee Lee. Convolutional sequence to sequence model for human dynamics. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5226–5234, 2018. 3
- [29] 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 *IEEE Conference on Computer Vision and Pattern Recognition*, pages 214–223, 2020. 3
- [30] Hung Yu Ling, Fabio Zinno, George Cheng, and Michiel Van De Panne. Character controllers using motion VAEs. *ACM Transactions on Graphics*, 39(4):40–1, 2020. 3
- [31] Hongyi Liu and Lihui Wang. Human motion prediction for human-robot collaboration. *Journal of Manufacturing Systems*, 44:287–294, 2017. 2
- [32] Matthew Loper, Naureen Mahmood, and Michael J Black. Mosh: Motion and shape capture from sparse markers. *ACM Transactions on Graphics*, 33(6):1–13, 2014. 2, 4, 5
- [33] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics*, 34(6):248:1–248:16, Oct. 2015. 2, 3, 8
- [34] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *International Conference on Computer Vision*, pages 5442–5451, Oct. 2019. 2, 4, 5, 1
- [35] Wei Mao, Miaomiao Liu, Mathieu Salzmann, and Hongdong Li. Learning trajectory dependencies for human motion prediction. In *International Conference on Computer Vision*, pages 9489–9497, 2019. 3
- [36] Julieta Martinez, Michael J Black, and Javier Romero. On human motion prediction using recurrent neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2891–2900, 2017. 3, 4
- [37] Meinard Müller, Tido Röder, Michael Clausen, Bernhard Eberhardt, Björn Krüger, and Andreas Weber. Documentation mocap database HDM05. 2007. 5
- [38] Dirk Ormoneit, Hedvig Sidenbladh, Michael Black, and Trevor Hastie. Learning and tracking cyclic human motion. *Advances in Neural Information Processing Systems*, 13:894–900, 2000. 3
- [39] Ahmed A. A. Osman, Timo Bolkart, and Michael J. Black. STAR: Sparse trained articulated human body regressor. In *European Conference on Computer Vision*, volume LNCS 12355, pages 598–613, Aug. 2020. 3, 1
- [40] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 10975–10985, June 2019. 2, 3, 4
- [41] Dario Pavllo, Christoph Feichtenhofer, Michael Auli, and David Grangier. Modeling human motion with quaternion-based neural networks. *International Journal of Computer Vision*, pages 1–18, 2019. 3
- [42] Dario Pavllo, David Grangier, and Michael Auli. Quaternet: A quaternion-based recurrent model for human motion. In *British Machine Vision Conference*, 2018. 3, 4
- [43] K. Pullen and C. Bregler. Animating by multi-level sampling. In *Proceedings Computer Animation 2000*, pages 36–42, 2000. 3
- [44] K Ramamohan Rao and Ping Yip. *Discrete cosine transform: algorithms, advantages, applications*. Academic press, 2014. 2, 4
- [45] Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied Hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics, (Proc. SIGGRAPH Asia)*, 36(6), Nov. 2017. 3
- [46] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Darius M Gavrilu, and Kai O Arras. Human motion trajectory prediction: A survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020. 2
- [47] 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. 2, 5, 7
- [48] Yongyi Tang, Lin Ma, Wei Liu, and Wei-Shi Zheng. Long-term human motion prediction by modeling motion context and enhancing motion dynamic. In *International Joint Conference on Artificial Intelligence*, page 935–941, 2018. 3
- [49] Jacob Walker, Kenneth Marino, Abhinav Gupta, and Martial Hebert. The pose knows: Video forecasting by generating pose futures. In *International Conference on Computer Vision*, pages 3332–3341, 2017. 3, 7
- [50] Mao Wei, Liu Miaomiao, and Salzmann Mathieu. History repeats itself: Human motion prediction via motion attention. In *European Conference on Computer Vision*, 2020. 3
- [51] Xinshuo Weng, Jianren Wang, Sergey Levine, Kris Kitani, and Nick Rhinehart. 4D Forecasting: Sequential Forecasting of 100,000 Points. *Euro. Conf. Comput. Vis. Worksh.*, 2020. 3
- [52] Xinshuo Weng, Jianren Wang, Sergey Levine, Kris Kitani, and Nick Rhinehart. Inverting the Pose Forecasting Pipeline with SPF2: Sequential Pointcloud Forecasting for Sequential Pose Forecasting. *CoRL*, 2020. 3
- [53] Erwin Wu and Hideki Koike. Futurepong: Real-time table tennis trajectory forecasting using pose prediction network. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–8, 2020. 2
- [54] Hongyi Xu, Eduard Gabriel Bazavan, Andrei Zanfir, William T Freeman, Rahul Sukthankar, and Cristian Sminchisescu. GHUM & GHUML: Generative 3D human shape and articulated pose models. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6184–6193, 2020. 3, 1
- [55] Xinchen Yan, Akash Rastogi, Ruben Villegas, Kalyan Sunkavalli, Eli Shechtman, Sunil Hadap, Ersin Yumer, and Honglak Lee. MT-VAE: Learning motion transformations to generate multimodal human dynamics. In *European Conference on Computer Vision*, pages 265–281, 2018. 3, 7

- [56] Ye Yuan and Kris Kitani. DLow: Diversifying latent flows for diverse human motion prediction. *European Conference on Computer Vision*, 2020. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [1](#)
- [57] Ye Yuan and Kris M. Kitani. Diverse trajectory forecasting with determinantal point processes. In *International Conference on Learning Representations*, 2020. [7](#)
- [58] M Ersin Yumer and Niloy J Mitra. Spectral style transfer for human motion between independent actions. *ACM Transactions on Graphics*, 35(4):1–8, 2016. [3](#)
- [59] Jason Y Zhang, Panna Felsen, Angjoo Kanazawa, and Jitendra Malik. Predicting 3D human dynamics from video. In *International Conference on Computer Vision*, pages 7114–7123, 2019. [2](#)
- [60] Yan Zhang, Michael J Black, and Siyu Tang. Perpetual motion: Generating unbounded human motion. *arXiv preprint arXiv:2007.13886*, 2020. [2](#), [3](#), [4](#)
- [61] Yi Zhou, Connelly Barnes, Jingwan Lu, Jimei Yang, and Hao Li. On the continuity of rotation representations in neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5745–5753, 2019. [3](#)
- [62] Yi Zhou, Zimo Li, Shuangjiu Xiao, Chong He, Zeng Huang, and Hao Li. Auto-conditioned recurrent networks for extended complex human motion synthesis. In *International Conference on Learning Representations*, 2018. [3](#)