

3D Human Motion Estimation via Motion Compression and Refinement

Zhengyi Luo^[0000-0002-1842-7622], S. Alireza Golestaneh^[0000-0002-3230-2797],
and Kris M. Kitani^[0000-0002-9389-4060]

Carnegie Mellon University, Pittsburgh, PA 15213, USA
{zluo2,sgolesta,kkitani}@cs.cmu.edu

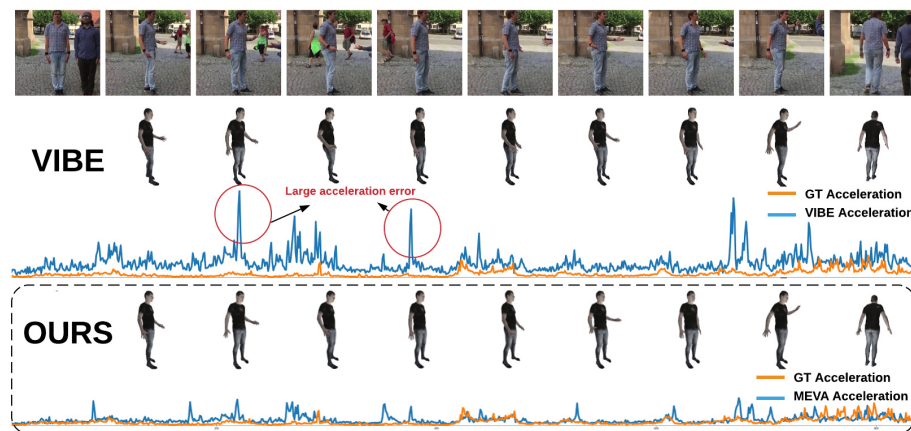


Fig. 1. Given an in-the-wild video, state-of-the-art methods (*e.g.*, VIBE [1]) can achieve high per joint accuracy, but also suffers from high acceleration error. To tackle this, we develop a two-stage motion estimation method, *MEVA*, that is able to produce both accurate and smooth human motion.

Abstract. We develop a technique for generating smooth and accurate 3D human pose and motion estimates from RGB video sequences. Our method, which we call **M**otion **E**stimation via **V**ariational **A**utoencoder (*MEVA*), decomposes a temporal sequence of human motion into a smooth motion representation using auto-encoder-based motion compression and a residual representation learned through motion refinement. This two-step encoding of human motion captures human motion in two stages: a general human motion estimation step that captures the coarse overall motion, and a residual estimation that adds back person-specific motion details. Experiments show that our method produces both smooth and accurate 3D human pose and motion estimates.

1 Introduction

Estimating the 3D pose sequence of a person from a single video requires a computational model that can extract the underlying kinematics of human motion while also preserving motions that are unique to the person being captured. Since people share a similar body structure (*e.g.*, same number of joints) and similar physical constraints (*e.g.*, joint limitations), it is possible to learn a generalized kinematic model that can be matched against the image to infer the general motion of a person. However, since generalized models of motion can also fail to model person-specific motions, it may also be necessary to ‘add back in’ or refine the general motion estimates using image evidence. In this work, we present a two-stage 3D motion estimation method that first extracts coarse kinematic sequences of a person in a video and then refines that sequence to produce a more accurate 3D motion estimate. We show that by decomposing the inference process into (1) a general model of motion and (2) a person-specific model of motion, we are able to obtain more accurate and smooth estimates.

Over the past years, significant progress has been made on improving the accuracy of 3D human pose estimation. Impressive results have been obtained through human mesh recovery from single images [2,3,4,5,6] and videos [7,1,8,9]. The main metric used to evaluate these methods is the Mean Per Joint Position Error (MPJPE), which measures the performance in terms of the relative joint positions computed for each frame of a video. However, less emphasis has been given to the temporal smoothness of the estimated motion. Optimizing for this metric, the tendency is to generate pose estimates that ‘jitter’ near the true pose. This is expected as the MPJPE only penalizes for spatial errors and is not designed to account for temporal consistency. As the methods for 3D pose estimation have improved in recent years, the ‘jitter’ has become less pronounced, especially when applied to dynamic scenes with vibrant moving backgrounds and camera motion. However, by rendering the estimated 3D pose of state-of-the-art methods on a plain background, the ‘jitter’ can still be observed, resulting in an overall unnatural motion estimation.

The issue of temporal smoothness is a known problem and has been addressed in part by prior work. Large-scale motion datasets such as Archive of Motion Capture as Surface Shapes (AMAAS) [10] and adversarial loss have enabled methods to improve both pose accuracy and temporal smoothness [1,7]. Other methods have been developed to enforce temporal smoothness [2] by letting the model predict frame ordering. However, using prior knowledge only in the loss function, it is hard to find the balance between smoothness and accuracy.

In this work, we argue that striking such balance between smoothness and accuracy can be done through an explicit breakdown of coarse and fine motion. First, we learn a coarse motion model by observing a large dataset of human motion—since human motion is typically smooth (*e.g.*, we usually do not shake as we walk), if one were to fit a model to a large set of human motions, most of the data would lie in a sub-space in which motions are smooth. This implies that if we were to compress human motion data, it should learn a latent subspace in which the motions are inherently smooth and coarse. Using this latent space

as the regression target, we can directly infer coarse human motion from the input videos. The problem of using such a human motion subspace, of course, is that rare motions (*e.g.*, sudden motions) are removed from the motion data. To retain such motion, we also argue that producing the final 3D motion estimate can be treated as a refinement step to "add back" the fine details to the coarse motion sequence.

To validate our arguments, we propose a two-stage 3D human motion estimation method that first estimates a coarse human pose sequence through data compression using a Variational Autoencoder (VAE), which we call the Variational Motion Estimator (VME). Then we take the output of the VME and refine the pose estimate using the image evidence with a pose regressor, which we call the Motion Refinement Regressor (MRR).

In summary, we propose a video-based 3D human motion estimation method that focuses on producing smooth and accurate human motion sequences. Our main contributions are as follows: (1) we propose a two-stage motion estimation method for ensuring temporal smoothness and accurate pose estimates; (2) we describe a procedure to learn a robust Variational Autoencoder that can serve as a latent human motion subspace for estimating coarse 3D human motion from videos; (3) we demonstrate state-of-the-art pose/motion estimation performance on challenging in-the-wild dataset such as 3DPW [11], reducing acceleration error by 54.3% while achieving state-of-the-art MPJPE results.

2 Related Works

In this section, we will first review the relevant work in human shape and pose recovery from a single image and from videos—human motion recovery can be treated as a subset of human pose estimation, as human motion is a sequence of human poses. Then we will review how existing methods use human motion as a prior and how popular methods map motion sequences to a low-dimensional space.

2.1 Recovering 3D human pose and shape from a single image

Here we focus on model-based methods that jointly recover human shape and pose. We choose to use a parametric 3D human body model [12,13,14] since it can be easily turned into a 3D human mesh that is usable for downstream tasks such as animation. Directly fitting a parametric 3D human body to image input has gained substantial traction over the years, morphing from methods that require silhouette or human input [15,16,17,18,19], to ones that can directly fit model parameters to 2D joint positions [20], and to ones that can directly estimate shape and pose from images [8,21,22,12,23,24,3]. Due to the lack of ground truth 3D labels, these methods use a weakly supervised approach to fit the 3D human body to 2D joint positions [8,23,24], body part segmentation [21,25], or dense pixel correspondence [6]. Although these methods achieve amazing results, their extracted motion tends to be unstable due to the lack of temporal information.

2.2 Recovering 3D human pose and shape from video

Using temporal information to aid 3D human pose estimation is a natural extension to single frame methods. [9,26,27,28] focus on "lifting" predicted joint positions from 2D to 3D, and uses LSTM [9], Temporal Convolution [26], and fully connected layers [27] to exploit temporal information. [29,30], on the other hand, predict 3D joint positions directly from images and use a temporal filter to postprocess the motion sequence. For methods that jointly recover shape and pose. HMMR [7], Sun *et al.* [2] and VIBE [1] are the best performing models that exploit temporal information. HMMR [7] proposes to enforce temporal consistency by letting the model predict future and past frames of motion. Sun *et al.* [2] learns temporal information by predicting the ordering of shuffled frames. VIBE [1] utilizes temporal information by employing a Gated Recurrent Unit (GRU) to convert the input frames of features into a sequence of temporally correlated latent features.

2.3 Human Pose and Motion Prior

Using prerecorded human motion sequences as a prior has also been explored in various tasks related to human motion. Earlier work like [31] tries to quantify unnaturalness in animated human motion by statistically analyzing existing motion capture (MoCap) sequences, and [32,33] propose to use learned MoCap motions to aid 3D motion tracking. More recently, [8,7] use an adversarial discriminator at a per-frame level to ensure that the recovered pose is a valid human pose. [12] uses a pretrained pose VAE's latent space for a similar purpose. [1] proposes to use the discriminator at a temporal level, discriminating against a whole motion sequence. All above methods use a pose or motion prior in an adversarial way, utilizing the prior knowledge in the loss function.

2.4 Human Motion Representation

Compressing human motion into a compact latent representation plays an important role in tasks such as human motion generation [34,35,36], human motion generation across different modalities [37,38,39,40], and trajectory forecasting [41,42,43,44]. Existing methods leverage different generative models such as VAE [42,43], generative flow [45], or generative adversarial networks [34] to achieve a compact motion representation in the latent space.

3 Approach

As discussed in Sections 1 and 2, existing human motion estimation methods often find it difficult to achieve a balance between temporal smoothness and accuracy. To tackle this, we propose *MEVA*, **M**otion **E**stimation via **V**ariational **A**utoencoder, a framework that learns the overall coarse motion first then adds back fine detailed motion as a residual. *MEVA* processes the inputs in three steps:

it first extracts correlated temporal features using a spatio-temporal feature extractor (STE), and then captures the overall coarse motion through a variational motion estimator (VME), and finally uses a motion residual regressor (MRR) to add back the fine motion details. Our overall framework can be shown in Fig. 2. In this section, we will first setup the overall problem and then present the details of our framework. Finally, we will discuss the training procedures.

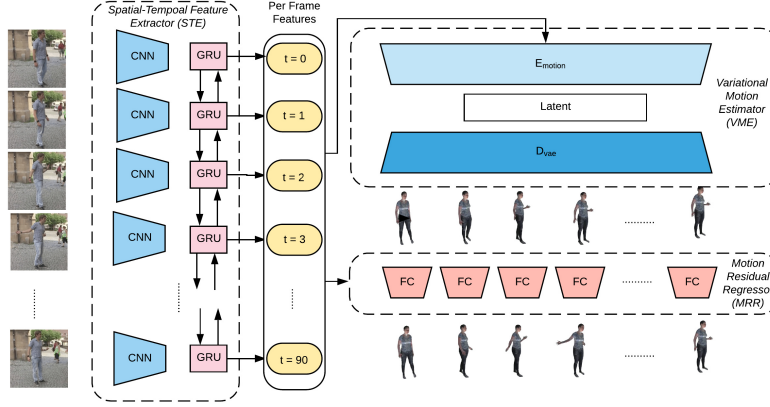


Fig. 2. Overall Architecture. *MEVA* estimates motion from videos by first extracting temporal features using Spatio-Temporal Feature Extractor (STE) and then estimates the overall coarse motion inside the video with Variational Motion Estimator (VME). Finally, a Motion Residual Regressor (MRR) is used to refine the motion estimates.

3.1 Problem Formulation

Given an input video $V_T = \{I_t\}_{t=1}^T$, where I_t denotes the t^{th} frame, our task is to recover coherent human motion sequences $M_T = \{\theta_t\}_{t=1}^T$ where each θ_t represents the human pose for the t^{th} frame. To represent the human motion, we utilize the SMPL 3D human mesh model [13]. We choose SMPL parametrization over 3D joint positions or other human models due to its versatility: SMPL parameters can be easily converted to 3D joint positions and human mesh. Specifically, a human body is represented by its shape β and pose θ , denoted by $\Theta = \{\beta, \theta\}$. Given θ and β , let S denote the pretrained SMPL function, where $S(\Theta) : \beta, \theta \rightarrow R^{6890 \times 3}$ (6890 is the number of vertices of the resulting triangular human mesh). The pose parameter $\theta \in R^{24 \times N}$ stands for the joint angles for the 23 joints plus the root orientation. N is the dimension of the chosen rotation representation ($N = 3$ for axis/euler angle, $N = 4$ for quaternions, $N = 6$ for a 6 degrees-of-freedom rotation representation [46]). The shape parameters $\beta \in R^{10}$ represent the linear coefficient for the principal component of the parametric human shape space. Given a set of β and θ , S can recover the 3D joint positions

through a pretrained mesh vertex regressor $P : jp^{3d} = P\{S(\beta, \theta)\} \in R^{N \times 3}$. To project the 3D joint positions back to 2D images, a weak perspective camera $\pi = \{s, t_x, t_y\}$ needs to be estimated: $jp^{2d} = \Pi(P\{S(\beta, \theta)\}) \in R^{N \times 2}$.

Intuitively, recovering human motion from video frames does not require recovering the human shape: one can directly learn a mapping from input video frames to the estimated human motion if sufficient paired ground truth data exist. However, videos paired with ground truth motion annotation (SMPL sequences) require professional capture equipment such as a motion capture (MoCap) rig, which is still rather rare compared to annotated 2D pose datasets. In the absence of 3D data samples, it is critical for our model to learn motion sequences in a semisupervised fashion (*i.e.*, from videos with 3D or 2D labeled joint positions following the approach utilized in [8,7,1,3,2]).

Overall, our motion estimation objective is to learn a function $\text{MEVA}(V) : V_T \rightarrow M_T$ where $V_T = \{I_t\}_{t=1}^T$ and $M_T = \{\theta_t\}_{t=1}^T$.

3.2 Spatio-Temporal Feature Extractor (STE)

Human motion is inherently temporal and correlated, and past movement can give cues about future motion. Thus, instead of extracting per-frame visual features using a feed-forward convolutional network independently, we can produce temporally correlated features that lead to better motion sequence modeling. Similar to [1], we use a GRU-based temporal feature extractor (STE) that encodes the input video frames $I_1, I_2, I_3, \dots, I_T$ into a sequence of temporally correlated features $f'_1, f'_2, f'_3, \dots, f'_T$.

3.3 Variational Motion Estimator (VME)

Human Motion VAE To learn a human motion subspace that can encapsulate a broad spectrum of human motion, we choose to use a Variational Autoencoder (VAE). VAEs can effectively capture a large number of possible data modes by explicitly mapping each data point to a latent code, and by imposing a Gaussian prior on the learned latent space, similar motions' latent codes will be near each other. Thus, the VAE's latent space allows for more overlap between codes and therefore enforces smoothness in the latent space. Having a smooth latent code is essential in improving the generalizability of the model since the space of possible human motion is highly correlated and limited. Formally, following the previous work on VAEs [43,47,48,48], the objective is to maximize the evidence lower bound of the log-likelihood $p_\lambda(x)$ (λ and ϕ denotes the function parametrization):

$$L_{VAE} = E_{q_\phi}[\log p(x|z)] - \text{KLD}(q_\phi(z|x)||p_\phi(z)), \quad (1)$$

where x is the input and the latent code $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

In the context of encoding human motion via VAE, the encoder E_{vae} takes in a sequence of W frames of human motion represented in terms of SMPL pose parameters: $x = M_W = [\theta_{w1}, \theta_{w2}, \theta_{w3}, \dots] \in R^{W \times 144}$ and outputs the latent code z . A single frame of SMPL pose is represented in joint rotations,

resulting by a $24 \times 6 = 144$ dimensional input (a 6 degrees-of-freedom rotation representation[46] is used for continuity purpose). The decoder D_{vae} takes in the latent code z and reconstructs the motion \hat{M}_W . Both the encoder E_{vae} and decoder D_{vae} are implemented as GRUs, and the detailed architectures are given in Fig.3. Based on the Gaussian parameterization of the VAE, the objective function of Eq.(1) can be written Eq.(2)




Fig. 3. Motion VAE Architecture.

$$\mathcal{L}_{vae}(\mathbf{x}; \theta, \phi) = -\frac{1}{S} \sum_{s=1}^S \|\tilde{\mathbf{x}} - \mathbf{x}\|^2 + \beta \cdot \frac{1}{S_z} \sum_{j=1}^{S_z} (1 + 2 \log \sigma_j - \mu_j^2 - \sigma_j^2), \quad (2)$$

where S is the number of samples for the current batch, S_z is the dimension of the current latent variable, and β is the weighing parameter. Once the VAE is trained and converged to a desirable reconstruction accuracy, the decoder D_{vae} is frozen for later use. During inference, given a latent code $z \in R^{1 \times S_z}$, D_{vae} can decode it back into a sequence of human motion: $M_W \in R^{w \times 144}$.

Human Motion Data augmentation Our learned VAE should be able to generalize to unseen human motion sequences and achieve high reconstruction accuracy to ensure that the learned latent space can indeed serve as a comprehensive human motion subspace. Using an already large-scale human motion dataset AMASS [10] (13k motion samples with varying length), our trained VAE still has poor generalizability on unseen sequences (for details refer to Sec.4.3). Thus, we devise an elaborate data augmentation scheme that can augment the existing motion and produce viable yet distinct human motion sequences. While data augmentation has been studied extensively in image processing, to the best of our knowledge, few attempts have been done in augmenting a human motion dataset. When used in trajectory forecasting [43] and human motion generation [34], the generalizability of the VAE latent space has not been discussed extensively since the models only need to generate new motion sequences and do not emphasize on the ability to encode unseen motion sequences.

Given a T frame human motion sequence in SMPL parameters $M_T \in R^{T \times 144}$ with a frame-rate F_{amass} , we employ the following data augmentation scheme:

- **Speeding up and slowing down:** based on F_{amass} , we can uniformly up-sample or downsample the frames and produce novel sequences that are still plausible and natural human motion.
- **Flipping left and right:** The same action, performed using either the left or right hand, will remain a valid human motion. Thus, we can follow the kinematic tree of the SMPL model and mirror the motion across the left and right and generate a new motion sequence.
- **Random root rotation:** We randomly sample a root rotation from a unit sphere to capture different root orientations for possible human motion. Different pose estimators may assume different ground planes and coordinate systems, so SMPL parameters usually come in different root orientations. Sampling random root rotation helps the model cope with different possible coordinate frame choices.

Learning Smooth Motion from videos After learning a comprehensive human motion subspace using the VAE, we learn an additional encoder E_{motion} that can directly extract coarse motion sequences from video features, mapping to the same latent space as E_{vae} . Given an input sequence of video features $f_W = \{f_w\}_{w=1}^W$, the encoder E_{motion} ’s task is to compress the input features into a latent code z that best summarizes the current observation as a coarse human motion sequence. We use the pretrained decoder D_{vae} from the motion VAE to force E_{motion} to sample from our pretrained motion subspace. Constraining the latent space of the E_{motion} to a pretrained human motion subspace provides a strong human motion prior that greatly aids the learning process of E_{motion} . Combining E_{motion} and D_{vae} , we form our Variational Motion Estimator (VME).

3.4 Motion Residual Regressor (MRR)

As noted in the previous section, the learned motion sequences using the VAE’s latent space as the target are inherently smooth and coarse, capturing the overall motion signature of the current video frames through information compression. To capture the details, we utilize a SMPL regressor from [8] that can iteratively refine the estimated poses. The regressor takes in an initial pose and shape estimation Θ_t and the visual feature f_t for a single frame to calculate its estimation Θ'_t for k iterations. Notice that though [8,7,1,3] all utilize the same regressor, *MEVA* uses it in a fundamentally different way—in [8,7,1,3], the regressor is initialized with mean SMPL pose Θ_{mean} . At a sequence level, a regressor initialized uniformly with the mean pose Θ_{mean} is trying to capture the overall motion in one pass, while in *MEVA*, the regressor is initialized with the computed poses from VME. Thus, the regressor is only tasked to do small cosmetic changes to the coarse estimation, adding back the fine details of the motion lost during our compression step. Similar to [1], the input visual features f_t to the regressor are

encoded using a temporal visual encoder, so even though each frame’s estimation is calculated separately at this stage, the visual feature is already temporally correlated. The VME computes the overall coarse motion from videos by using a general model of motion, and the regressor jointly refines motion and human shape estimates, which amounts to adding back the person-specific motion details at a per frame level. We call this regressor the Motion Residual Regressor (MRR). MRR completes the overall framework of our proposed method, as shown in Fig. 2.

3.5 Training and Losses

Our framework is trained in two stages. At first, the motion VAE is pretrained. Then STE, VME, and MRR are trained jointly end-to-end. Using videos with various levels of annotation (2D joint positions, 3D joint positions, SMPL parameters), similar to in [7,1,3], the networks are trained with losses consisting of L_{2D} , L_{3D} , L_{SMPL} , as long as respective data is available. Specifically:

$$L_{meva} = L_{3D} + L_{2D} + L_{SMPL}, \quad (3)$$

$$L_{3D} = \sum_{t=1}^T \|jp_t^{3d} - \hat{jp}_t^{3d}\|_2, \quad (4)$$

$$L_{2D} = \sum_{t=1}^T \|jp_t^{2d} - \hat{jp}_t^{2d}\|_2, \quad (5)$$

$$L_{SMPL} = \|\beta - \hat{\beta}\|_2 + \sum_{t=1}^T \|\theta_t - \hat{\theta}_t\|_2. \quad (6)$$

For implementation details, please refer to the supplementary material.

4 Experiments

To demonstrate the effectiveness of our proposed method, we evaluate our method in terms of the overall accuracy and smoothness of the estimated motion on MPI-INF-3DPH [49], 3DPW [11], and human 3.6M [50]. In the following sections, we will first describe the main datasets used to train and evaluate *MEVA*. Then in Section 4.2 we provide extensive evaluation results. Finally, in Section 4.3, we provide the ablation studies for our proposed method.

4.1 Datasets

For the motion VAE, we use motion sequences from AMASS [10] for training and sequences from 3DPW [11] for evaluation. For training with videos, in addition to the train split of MPI-INF-3DPH [49], 3DPW [11], and human 3.6M [50], which have 3D joint annotation, we also use InstaVariety [7] and PennAction [51] which contain 2D joint annotation.

For training our motion VAE, we use **AMASS** [10]. It is a recent dataset that contains a large sample of human motion sequences in SMPL parameters. These sequences are fitted from MoCap sequences using Mosh++ [10]. There

are in total 13k motion sequences with varying length. We use this dataset only for training our motion VAE.

For training with videos, **3DPW** is the only dataset that contains paired SMPL parameters and video sequences (which provide direct supervision to *MEVA*). The videos from this dataset are mostly outdoors and in-the-wild. It uses paired IMU sensors and video input to compute the near ground truth SMPL pose and shape parameters. This is a relatively small dataset and we use the official split in [11] for train, validation, and test. There are in total 60 videos with varying length (24 train, 24 test, 12 val). **MPI-INF-3DHP** is a dataset that contains 3D joint position annotation. It is captured using a multiview camera setup, and the 3D joint annotation is calculated through multiview methods [49]. There are 8 subjects and 16 videos per subject, in total 128 videos with varying length. We use the official test and train split. **H3.6M** is a popular pose estimation dataset that contains 3D joint position annotations, captured indoors with MoCap markers. Notice that a number of previous works [8,7,3,1,2] had access to a near-ground truth SMPL pose and shape parameters calculated by the Mosh [52] method. However, this annotation has since been removed from public access due to legal issues. SMPL parameters provide the best supervision, as noted in [53], so for a fair comparison we retrain some of the state-of-the-art methods without such supervision. There are 840 videos in total across 7 subjects in H3.6M and we use the official train/test subject split ([S1, S5, S6, S7, S8] vs [S9, S11]). During preprocessing, we subsample every 5 frames from the dataset. The **PennAction** dataset contains human annotated ground truth 2D keypoints paired with video sequences. There are in total 2326 videos with varying length. **InstaVariety** dataset contains human annotated pseudo ground truth 2D keypoints paired with video sequences. The 2D keypoints are estimated using openpose[54]. There are in total 28,272 videos with varying length.

4.2 Evaluation Results and Analysis

Metrics To best capture human motion, we utilize three popular metrics that measure the overall accuracy and smoothness of the motion. Mean per joint position error (MPJPE) and MPJPE after Procrustes Alignment (PA-MPJPE) measure the 3D joint positional discrepancy between the predicted and ground truth 3D joint positions in millimeters (*mm*), and are calculated after aligning the root position (human pelvis) of the 3D joint positions. Both MPJPE and PA-MPJPE serve as the accuracy indicator of the motion estimator. Acceleration error (ACC-ERR), proposed in [7], measures the difference between the predicted and ground truth 3D acceleration for each keypoint in mm/s^2 . ACC-ERR serves as the major smoothness indicator for the estimated motion sequences. Acceleration is calculated using the finite difference between individual frames. It is imperative to view these metrics jointly: a low position error indicates overall correctness in motion capture and a better acceleration error marks a smooth and natural estimation of human motion.

Generalization of Motion VAE In this section, we study the generalizability of our learned motion VAE. We report the reconstruction error of the VAE on unseen motion sequences from the 3DPW dataset. Table 1 shows VAE motion reconstruction error over the different splits of 3DPW. The Motion VAE model is the best performing model that is trained with all three forms of data augmentation techniques. Detailed analysis about the effects of data augmentation can be found in 4.3. The result shows that our VAE generalizes well to unseen sequences and the learned subspace can represent the human motion space with reasonable quality.

Table 1. VAE Reconstruction Test Error on 3DPW dataset Here, the VAE is tasked to encode and decode unseen motion sequences from the 3DPW dataset, and we calculate our metrics between the ground truth and reconstructed sequences. Motion sequences from different splits of 3DPW have varying difficulties, but are all unseen by our VAE.

	3DPW Train Split			3DPW Val Split			3DPW Test Split		
	MPJPE ↓	PA-MPJPE ↓	ACC-ERR ↓	MPJPE ↓	PA-MPJPE ↓	ACC-ERR ↓	MPJPE ↓	PA-MPJPE ↓	ACC-ERR ↓
Motion VAE	72.7	52.3	8.9	72.1	52.9	9.3	58.7	43.9	8.2

Quantitative Results The result in Table 2 shows that our method obtains state-of-the-art results on video motion estimation across all three test datasets. Overall, our method achieves comparable results in terms of position error (MPJPE and PA-MPJPE) while significantly improving the smoothness (acceleration error), signifying a smoother and more natural motion estimation without sacrificing accuracy. Notice that all methods in italics have access to SMPL parameter annotation to the H3.6M dataset, which has since been removed from the web due to legal reasons. The SMPL parameters provide the most direct supervision for the task, so the performance gain is significant especially on the H3.6M test set. For a more direct comparison, we retrain the previous state-of-the-art method, VIBE [1], using the official implementation with the exact same datasets as ours. On 3DPW, under the same training condition, *MEVA* outperforms VIBE on almost all three metrics while reducing the acceleration error by 54.3%, 59.3%, and 41.3 %, respectively. Even compared to VIBE trained with extra data, our model achieves comparable results in accuracy while sporting a great reduction in acceleration error, except for the H3.6M dataset. Note that the H3.6M dataset contains mainly indoor scenes with limited background variation and models trained with direct SMPL supervision tend to perform well on this dataset. Compared to HMMR [7], which is the state-of-the-art on smoothness, our model still achieves a smoother result (23.7% reduction in acceleration error) while improving greatly in MPJPE by 25.4%.

Qualitatively Results As motions are best seen in videos, please refer to [the supplementary video](#) for qualitative results. Overall, our model achieves better

Table 2. Testing error of state-of-the-art models on 3DPW, MPI-INF-3DHP, and H3.6M. Here we compare with state-of-the-art methods on video pose estimation, and report metrics on positional accuracy (MPJPE and PA-MPJPE) as well as acceleration error. Notice that since an important annotation of H3.6M dataset has since been made unavailable, we put all methods that are trained with such supervision in italics. The most fair comparison is between our method and VIBE [1], the previous best-performing model (trained using the same datasets).

	3DPW			MPI-INF-3DHP			H3.6M		
	MPJPE ↓	PA-MPJPE ↓	ACC-ERR ↓	MPJPE ↓	PA-MPJPE ↓	ACC-ERR ↓	MPJPE ↓	PA-MPJPE ↓	ACC-ERR ↓
<i>HMR (w/ H3.6M SMPL) [8]</i>	<i>130.0</i>	<i>76.7</i>	<i>37.4</i>	<i>124.2</i>	<i>89.8</i>	-	<i>88</i>	<i>56.8</i>	-
<i>HMMR (w/ H3.6M SMPL) [7]</i>	<i>116.5</i>	<i>72.6</i>	<i>15.2</i>	-	-	-	-	<i>56.9</i>	-
<i>SPIN (w/ H3.6M SMPL) [3]</i>	<i>96.9</i>	<i>59.2</i>	<i>29.8</i>	<i>105.2</i>	<i>67.5</i>	-	-	<i>41.1</i>	-
<i>Sun et al. (w/ H3.6M SMPL) [2]</i>	-	<i>69.5</i>	-	-	-	-	<i>59.1</i>	<i>42.4</i>	-
<i>VIBE (w/ H3.6M SMPL) [1]</i>	<i>82.9</i>	<i>51.9</i>	<i>23.4</i>	<i>96.6</i>	<i>64.6</i>	<i>31.2</i>	<i>65.6</i>	<i>41.4</i>	<i>27.3</i>
VIBE (w/o H3.6M SMPL) [1]	91.9	57.6	25.4	103.9	68.9	27.3	78.0	53.3	27.3
MEVA (w/o H3.6M SMPL)	86.9	54.7	11.6	96.4	65.4	11.1	76.0	53.2	15.3

acceleration error while preserving high joint position accuracy, resulting in an overall smooth and natural motion.

4.3 Ablation Experiments

Effect of data augmentation for training the motion VAE As mentioned in Sec.3.3, data augmentation performed on the AMASS dataset significantly improves the generalizability of our motion VAE. Table 3 demonstrates the VAE reconstruction error on the unseen sequences from the 3DPW dataset (train/test/val), with various levels of data augmentation techniques. Overall, RR (random root) rotation is essential since the motion sequences in AMASS dataset, captured mostly in MoCap studio, have a single fixed initial root rotation. Model trained only on AMASS would suffer greatly when it encounters variation in root rotation. Changing the sampling frame rate (FR) and flip left and right (LR) also provide a significant boost to generalizability. Using all three augmentation techniques results in our best performing motion VAE.

Table 3. Ablation of Data augmentation for VAE. Here we show the effects of data augmentation on the VAE. RR: random root rotation, FR: different sampling framerate, LR: left and right flip

	Train			Val			Test		
	MPJPE ↓	PA-MPJPE ↓	ACC-ERR ↓	MPJPE ↓	PA-MPJPE ↓	ACC-ERR ↓	MPJPE ↓	PA-MPJPE ↓	ACC-ERR ↓
No Aug	200.2	86.0	15.9	181.7	90.5	15.8	172.8	75.2	14.1
RR	146.3	78.0	12.1	133.9	81.2	13.0	115.0	70.4	11.0
FR + LR	291.0	63.03	14.7	308.8	66.1	16.1	261.9	58.9	13.3
FR + LR + RR	72.7	52.3	8.9	72.1	52.9	9.3	58.7	43.9	8.2

Coarse motion vs fine motion retrieval *MEVA* benefits from an explicit breakdown of coarse and fine motion retrieval, using a temporal compressive step that captures the overall motion in a given human motion sequence. Just how

much coarse/smooth motion information is retrieved in the framework? Table 4 shows the result of *MEVA* if trained only using the VME (capturing only coarse motion). Notice that *MEVA* with only VME achieves result similar to HMMR [7], the previous state-of-the-art in producing low acceleration error motion estimation. An illustrative visualization of coarse and fine motion decomposition can be found in Fig. 4.




Fig. 4. Breakdown of coarse and fine human motion. The first row of estimated human is the coarse part of the motion (output of VME), while the second row adds back the fine details (output of MRR).

Table 4. Ablation of MEVA. Here we show the *MEVA* model trained with only Variational Motion Estimator (without Motion Residual Regressor) or without using pretrained VAE.

	3DPW		
	MPJPE ↓	PA-MPJPE ↓	ACC-ERR ↓
MEVA-VME only	118.1	73.7	15.4
MEVA-without using pretrained VAE	95.9	59.7	14.1
MEVA	86.9	54.7	11.6

Effects of the pretrained Motion VAE *MEVA* benefits from using a pretrained motion VAE’s latent space. As argued in Sec. 3.3, using a pretrained VAE provides a human motion subspace that assists in constraining the estimated motion sequences to be natural and plausible human motion. Table 4 shows the result of *MEVA* trained without using a pretrained VAE (not using the pretrained D_{vae}). In this case, the whole framework is trained end-to-end from scratch. Here we observe that the model performed relatively well in both accuracy and smoothness, demonstrating the power of our two-stage estimation framework. However, upon visual inspection, as shown in Fig. 5, a few kinematically invalid human poses are estimated during the sequence, resulting in an overall accurate but flawed estimation.




Fig. 5. MEVA results without using pretrained VAE. Here we show an ablative study where we do not use a pretrained VAE. The second row shows **MEVA** result without using a pretrained VAE and the first row shows our full model estimation. Notice that the model estimates unnatural human poses in the first few frames.

5 Conclusion

We have shown that to achieve temporally smooth and accurate 3D human pose estimates, it is important to learn a compressive model that encodes the smoothness of general human motion while also learning an image-based regression model that can capture person-specific motion. We propose a two-stage model that first trains a Variational Autoencoder to model the general statistics of coarse/smooth human motion and then learns a person-specific motion refinement regression module to retain motions not captured by the general motion model. Through comprehensive experiments, we demonstrate that our method produces both smooth and accurate motion.

Acknowledgements: This project was sponsored in part by IARPA (D17PC00340), and JST AIP Acceleration Research Grant (JPMJCR20U1).

References

1. Kocabas, M., Athanasiou, N., Black, M.J.: Vibe: Video inference for human body pose and shape estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2020) 5253–5263
2. Xu, Y., Zhu, S.C., Tung, T.: Denserac: Joint 3d pose and shape estimation by dense render-and-compare. 2019 IEEE/CVF International Conference on Computer Vision (ICCV) (2019) 7759–7769
3. Kolotouros, N., Pavlakos, G., Black, M.J., Daniilidis, K.: Learning to reconstruct 3d human pose and shape via model-fitting in the loop. 2019 IEEE/CVF International Conference on Computer Vision (ICCV) (2019) 2252–2261
4. Georgakis, G.V., Li, R., Karanam, S., Chen, T., Kosecka, J., Wu, Z.: Hierarchical kinematic human mesh recovery. ArXiv **abs/2003.04232** (2020)
5. Güler, R.A., Neverova, N., Kokkinos, I.: Densepose: Dense human pose estimation in the wild. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (2018) 7297–7306
6. Xu, Y., Zhu, S.C., Tung, T.: Denserac: Joint 3d pose and shape estimation by dense render-and-compare. 2019 IEEE/CVF International Conference on Computer Vision (ICCV) (2019) 7759–7769

7. Kanazawa, A., Zhang, J.Y., Felsen, P., Malik, J.: Learning 3d human dynamics from video. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 5607–5616
8. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (2018) 7122–7131
9. Hossain, M.R.I., Little, J.: Exploiting temporal information for 3d human pose estimation. In: ECCV. (2018)
10. Mahmood, N., Ghorbani, N., Troje, N.F., Pons-Moll, G., Black, M.J.: Amass: Archive of motion capture as surface shapes. 2019 IEEE/CVF International Conference on Computer Vision (ICCV) (2019) 5441–5450
11. von Marcard, T., Henschel, R., Black, M.J., Rosenhahn, B., Pons-Moll, G.: Recovering accurate 3D human pose in the wild using IMUs and a moving camera. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) **11214 LNCS** (2018) 614–631
12. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 10967–10977
13. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: a skinned multi-person linear model. ACM Trans. Graph. **34** (2015) 248:1–248:16
14. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: SCAPE: Shape Completion and Animation of People. ACM Transactions on Graphics **24** (2005) 408–416
15. Grauman, K., Shakhnarovich, G., Darrell, T.: Inferring 3D structure with a statistical image-based shape model. Proceedings of the IEEE International Conference on Computer Vision **1** (2003) 641–648
16. Agarwal, A., Triggs, B.: Recovering 3D human pose from monocular images. IEEE Transactions on Pattern Analysis and Machine Intelligence **28** (2006) 44–58
17. Sigal, L., Balan, A., Black, M.J.: Combined discriminative and generative articulated pose and non-rigid shape estimation. Advances in Neural Information Processing Systems 20 - Proceedings of the 2007 Conference (2009) 1–8
18. Zhou, S., Fu, H., Liu, L., Cohen-Or, D., Han, X.: Parametric reshaping of human bodies in images. ACM SIGGRAPH 2010 Papers, SIGGRAPH 2010 **29** (2010) 1–10
19. Peng Guan, Weiss, A., Blan, A.O., Black, M.J.: Estimating human shape and pose from a single image. In: 2009 IEEE 12th International Conference on Computer Vision. (2009) 1381–1388
20. Bogo, F., Kanazawa, A., Lassner, C., Gehler, P., Romero, J., Black, M.J.: Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics) **9909 LNCS** (2016) 561–578
21. Omran, M., Lassner, C., Pons-Moll, G., Gehler, P., Schiele, B.: Neural body fitting: Unifying deep learning and model based human pose and shape estimation. Proceedings - 2018 International Conference on 3D Vision, 3DV 2018 (2018) 484–494
22. Guler, R.A., Kokkinos, I.: Holopose: Holistic 3D human reconstruction in-the-wild. Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition **2019-June** (2019) 10876–10886
23. Tan, J.K.V., Budvytis, I., Cipolla, R.: Indirect deep structured learning for 3D human body shape and pose prediction. British Machine Vision Conference 2017, BMVC 2017 (2017) 1–11

24. Tung, H.Y.F., Tung, H.W., Yumer, E., Fragkiadaki, K.: Self-supervised learning of motion capture. *Advances in Neural Information Processing Systems* **2017-December** (2017) 5237–5247
25. Pavlakos, G., Zhu, L., Zhou, X., Daniilidis, K.: Learning to Estimate 3D Human Pose and Shape from a Single Color Image. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (2018) 459–468
26. Pavlo, D., Feichtenhofer, C., Grangier, D., Auli, M.: 3D human pose estimation in video with temporal convolutions and semi-supervised training. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* **2019-June** (2019) 7745–7754
27. Dabral, R., Mundhada, A., Kusupati, U., Afaq, S., Sharma, A., Jain, A.: Learning 3D human pose from structure and motion. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **11213 LNCS** (2018) 679–696
28. Xu, J., Yu, Z., Ni, B., Yang, J., Yang, X., Zhang, W.: Deep kinematics analysis for monocular 3d human pose estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. (2020)
29. Mehta, D., Sridhar, S., Sotnychenko, O., Rhodin, H., Shafiei, M.M.h.M.M.h.M.M.h., Seidel, H.p., Casas, D., Theobalt, C., Shafiei, M.M.h.M.M.h.M.M.h., Seidel, H.p., Xu, W.: VNect Real-time 3D Human Pose Estimation with a Single RGB Camera ACM Reference format VNect Real-time 3D Human Po.pdf. *Tog* **36** (2017) 1–13
30. Mehta, D., Sotnychenko, O., Mueller, F., Xu, W., Elgharib, M.A., Fua, P., Seidel, H.P., Rhodin, H., Pons-Moll, G., Theobalt, C.: Xnect: Real-time multi-person 3d motion capture with a single rgb camera. In: *SIGGRAPH 2020*. (2020)
31. Ren, L., Patrick, A., Efros, A.A., Hodgins, J.K., Rehg, J.M.: A data-driven approach to quantifying natural human motion. In: *SIGGRAPH 2005*. (2005)
32. Urtasun, R., Fleet, D.J., Fua, P.: 3D people tracking with Gaussian process dynamical models. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* **1** (2006) 238–245
33. Ormoneit, D., Sidenbladh, H., Black, M.J., Hastie, T.: Learning and tracking cyclic human motion. *Advances in Neural Information Processing Systems* (2001)
34. Wang, Z., Yu, P., Zhao, Y., Zhang, R., Zhou, Y., Yuan, J., Chen, C.: Learning diverse stochastic human-action generators by learning smooth latent transitions. *ArXiv abs/1912.10150* (2020)
35. Cai, H., Bai, C., Tai, Y.W., Tang, C.K.: Deep video generation, prediction and completion of human action sequences. In: *The European Conference on Computer Vision (ECCV)*. (2018)
36. Walker, J., Marino, K., Gupta, A., Hebert, M.: The pose knows: Video forecasting by generating pose futures. In: *The IEEE International Conference on Computer Vision (ICCV)*. (2017)
37. Plappert, M., Mandery, C., Asfour, T.: Learning a bidirectional mapping between human whole-body motion and natural language using deep recurrent neural networks. *Robotics and Autonomous Systems* **109** (2018) 13–26
38. Ahuja, C., Morency, L.P.: Language2pose: Natural language grounded pose forecasting. *2019 International Conference on 3D Vision (3DV)* (2019) 719–728
39. Yamada, T., Matsunaga, H., Ogata, T.: Paired recurrent autoencoders for bidirectional translation between robot actions and linguistic descriptions. *IEEE Robotics and Automation Letters* **3** (2018) 3441–3448
40. Lee, H.Y., Yang, X., Liu, M.Y., Wang, T.C., Lu, Y.D., Yang, M.H., Kautz, J.: Dancing to music. In: *NeurIPS*. (2019)

41. Fragkiadaki, K., Levine, S., Felsen, P., Malik, J.: Recurrent network models for human dynamics. *Proceedings of the IEEE International Conference on Computer Vision* **2015 Inter** (2015) 4346–4354
42. Yan, X., Rastogi, A., Villegas, R., Sunkavalli, K., Shechtman, E., Hadap, S., Yumer, E., Lee, H.: MT-VAE: Learning Motion Transformations to Generate Multimodal Human Dynamics. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **11209 LNCS** (2018) 276–293
43. Yuan, Y., Kitani, K.: Diverse trajectory forecasting with determinantal point processes. In: *International Conference on Learning Representations*. (2020)
44. Butepage, J., Black, M.J., Kragic, D., Kjellstrom, H.: Deep representation learning for human motion prediction and classification. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2017)
45. Yuan, Y., Kitani, K.: Dlow: Diversifying latent flows for diverse human motion prediction. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. (2020)
46. Zhou, Y., Barnes, C., Lu, J., Yang, J., Li, H.: On the continuity of rotation representations in neural networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019) 5738–5746
47. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *2nd International Conference on Learning Representations, ICLR 2014 - Conference Track Proceedings* (2014) 1–14
48. Walker, J., Marino, K., Gupta, A., Hebert, M.: The Pose Knows: Video Forecasting by Generating Pose Futures. *Proceedings of the IEEE International Conference on Computer Vision* **2017-October** (2017) 3352–3361
49. Mehta, D., Rhodin, H., Casas, D., Fua, P., Sotnychenko, O., Xu, W., Theobalt, C.: Monocular 3d human pose estimation in the wild using improved cnn supervision. *2017 International Conference on 3D Vision (3DV)* (2017) 506–516
50. Ionescu, C., Papava, D., Olaru, V., Sminchisescu, C.: Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **36** (2014) 1325–1339
51. Zhang, W., Zhu, M., Derpanis, K.G.: From actemes to action: A strongly-supervised representation for detailed action understanding. *Proceedings of the IEEE International Conference on Computer Vision* (2013) 2248–2255
52. Loper, M., Mahmood, N., Black, M.J.: MoSh: Motion and shape capture from sparse markers. *ACM Transactions on Graphics* **33** (2014)
53. Wang, Z., Yu, P., Zhao, Y., Zhang, R., Zhou, Y., Yuan, J., Chen, C.: Learning diverse stochastic human-action generators by learning smooth latent transitions. *ArXiv abs/1912.10150* (2020)
54. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. In: *CVPR*. (2017)