3D Kinematics Estimation from Video with a Biomechanical Model and Synthetic Training Data

Zhi-Yi Lin¹ Bofan Lyu² Judith Cueto Fernandez² Eline van der Kruk² Ajay Seth² Xucong Zhang¹

¹Computer Vision Lab, Department of Intelligent Systems, Delft University of Technology ²Department of BioMechanical Engineering, Faculty of Mechanical Engineering (3me), Delft University of Technology

Abstract

Accurate 3D kinematics estimation of human body is crucial in various applications for human health and mobility, such as rehabilitation, injury prevention, and diagnosis, as it helps to understand the biomechanical loading experienced during movement. Conventional markerbased motion capture is expensive in terms of financial investment, time, and the expertise required. Moreover, due to the scarcity of datasets with accurate annotations, existing markerless motion capture methods suffer from challenges including unreliable 2D keypoint detection, limited anatomic accuracy, and low generalization capability. In this work, we propose a novel biomechanics-aware network that directly outputs 3D kinematics from two input views with consideration of biomechanical prior and spatio-temporal information. To train the model, we create synthetic dataset ODAH with accurate kinematics annotations generated by aligning the body mesh from the SMPL-X model and a full-body OpenSim skeletal model. Our extensive experiments demonstrate that the proposed approach, only trained on synthetic data, outperforms previous stateof-the-art methods when evaluated across multiple datasets, revealing a promising direction for enhancing video-based human motion capture.

1. Introduction

Kinematics estimation is the process of capturing the relative position of human body segments through time. In biomechanics, accurate estimation of joint loading of the human body is crucial for applications in healthcare and sports, as it helps in understanding the biomechanical stresses experienced by the joints during movement. Optoelectronic motion capture systems are the established standard in Kinematics estimation [32]. These systems necessitate the attachment of reflective markers to the bony landmarks of the subject. As a result, this method is time-



Figure 1. The proposed biomechanics-aware network consists of a frame feature encoder and a spatio-temporal feature refinement module, which collectively infer 3D kinematics from two-view real-world video inputs. To train the model, we create a synthetic RGB video dataset ODAH by combining the kinematics skeleton from the OpenSim model, the body mesh from the SMPL-X model, and motions from the AMASS dataset to provide accurate ground truth data. Particularly, the end-to-end biomechanicsaware 3D kinematics estimation model is exclusively trained on this self-created synthetic data. Examples of real person images are from the OpenCap dataset [31], and faces were pixelated for privacy reasons.

consuming, requires a substantial financial investment and the markers also restrict the natural movement of the individuals [8, 20]. Consequently, many researchers have focused on developing markerless kinematics estimation [1, 23, 29, 31], which usually utilize 3D human pose estimation methods, such as OpenPose [2], as the backbone to infer 3D poses from detected 2D joints. The recent development of deep learning methods has propelled substantial advancements in the field of 3D human pose estimation from detected 2D joints [3, 7, 16, 38]. It provides great potential through the integration of deep learning 3D human pose estimation techniques and biomechanics models.

Unfortunately, there are significant gaps between the two research fields. The primary issue is the 2D joint annotations for conventional human pose estimation [2] are often anatomically wrong, resulting in imprecise 3D kinematics estimation. Many 3D human pose estimation methods overlook the incorporation of biomechanical constraints, leading to predictions that deviate from anatomically realistic movements. This discrepancy is especially problematic for 3D kinematics estimation, which demands a higher level of biomechanical fidelity for clinical diagnosis and sports science domains.

Another challenge is the scarcity of large-scale datasets with accurate 3D biomechanics annotations. The gold standard way to measure human joint kinematics is through the X-ray scan of the person [13, 35], which is not realistic to conduct in daily life settings. Although some datasets incorporate maker-based motion capture systems [9, 28, 31], they suffer from sensor noise and lack synchronization with the recording videos. Additionally, the expensive and timeconsuming nature of marker-based motion capturing further impedes the construction of large-scale datasets necessary for deep-learning-based markerless motion capture models [1].

In this work, we propose a novel markerless motion capture framework comprising a biomechanics-aware network, as illustrated in Figure 1. The model leverages two-view RGB videos as inputs and performs feature extraction [5], followed by a feature aggregation step to generate frame features. There is no explicit human body joint detection, instead, we sample the feature point on the input frame. Given the dynamic nature of human motion, the frame features are further refined by integrating temporal information through a transformer-based U-Net architecture [36]. Finally, the model outputs a sequence of joint angles and a set of body segment scales.

To address the limited availability of datasets with accurate kinematics annotations, we propose to create animated humans by aligning the SMPL-X model [24] from the computer vision community and OpenSim [6, 27] from the biomechanics community. We subsequently animate the SMPL-X model using OpenSim skeletal model and joint angles derived from AMASS dataset [19], which includes diverse motions captured by marker-based motion capture systems. To cover variant situations in real-world settings, we augment the data in terms of clothing, lighting, and camera positions, resulting in a synthetic dataset. We name this dataset as **O**penSim **D**riven Animated **H**uman (ODAH).

Please note we train the proposed biomechanics-aware network only on our self-created synthetic dataset, which has accurate ground truth labels. We perform an extensive evaluation of our synthetic data as well as two real-world datasets directly. we demonstrate that our framework outperforms three state-of-the-art markerless motion capture methods on average joint angle error and joint position error across all datasets, which shows the strong generalization of the proposed combination of biomechanics-aware network and synthetic data.

In summary, the main contributions of this paper are:

- We introduce an end-to-end 3D kinematics estimation model that predicts joint kinematics and body segment scales with an underlying OpenSim skeletal model.
- We create a synthetic video dataset ODAH with accurate kinematics annotations, varied subject appearance, motions, and scene settings.
- We demonstrate that exclusively trained on synthetic data, that the proposed biomechanics-aware network can achieve superior performance in average joint angle error across synthetic and real-world datasets, indicating its potential for improving kinematics estimation and domain generalization.

2. Related Works

2.1. Markerless Motion Capture with OpenSim

Markerless motion capture draws people's attention with its cost efficient nature and its ability to yield comparable results to marker-based motion capturing [4, 29]. Advancements in 3D human motion estimation have enabled the integration of 3D human motion estimation techniques with biomechanical models, allowing for a comprehensive analysis of human biomechanics [12]. Most of the existing markerless motion capture methods use multi-step processing for kinematics estimation [23, 31]. The process starts with deriving 3D joint positions by triangulating the detected 2D landmarks [2] from multiple views. Next, the 3D joint positions are treated as marker positions in the Inverse Kinematics (IK) tool and scaling tool in OpenSim software to derive joint kinematics and body segment scales. The mapping from the detected 3D joints to the real 3D joints can be encoded in the marker offsets defined in the Open-Sim model [23]. Alternatively, it is possible to train a model to learn this mapping function [31]. To further improve the performance, the landmark confidence scores are considered to remove low-confidence landmarks [23, 31].

However, the necessity for operator input and interaction in the intermediate steps (e.g. model scaling and inverse kinematics) of multi-step approaches introduces variability and inconsistency. Therefore, end-to-end solutions are preferred for reliable and accurate 3D kinematics estimation [1, 22]. In addition, the end-to-end method can leverage the raw information contained in the input image and not rely on landmark detection accuracy. D3KE [1], an endto-end method, utilizes CNNs to estimate joint kinematics and body segment scales for each frame from monocular videos. A lifting transformer encoder [15] is included to refine the predicted joint angles and body segment scales by incorporating temporal information.

Following the end-to-end approach, the proposed model simultaneously estimates joint kinematics and body scales based on visual inputs. Nevertheless, unlike D3KE, the backbone of the proposed network is particularly designed for human pose estimation, two views are utilized to better handle occlusions, and the model is trained on a large-scale synthetic dataset with accurate annotations and varied augmentations.

2.2. 3D Human Pose Estimation

The two primary approaches are 2D-to-3D lifting and direct 3D estimation. The 2D-to-3D lifting approach requires feature learning from a sequence of 2D poses. This allows for the consideration of temporal information inherent in human motions. Additionally, it compensates for the loss of 3D information from monocular inputs. Temporal dilated Convolution Networks (TCNs) are widely used [3, 18, 25, 34] because they can effectively learn spatial and temporal features that are essential for lifting 2D joints and improving motion coherence. In recent years, transformers have gained popularity in handling long-range sequential data [16, 17, 38, 39]. These methods typically involve lifting the 2D pose to 3D ones through transformer-based networks, followed by spatial and temporal refinements.

The direct human motion estimation approach eliminates the reliance on 2D landmark detection by directly predicting 3D poses, offering advantages such as avoiding biases towards specific camera views and mitigating ambiguities associated with 2D landmarks [33, 37]. Without guidance from 2D landmarks, multi-view inputs, and temporal feature learning become necessary in such frameworks.

With the direct human pose estimation strategy, the proposed method uses a frame feature encoder to map the input frame into the frame features with sampling on the input video frames. Different from previous works [5], we largely improve the architecture to work with two input views to perform 3D kinematic estimation. In addition, we propose a spatio-temporal refinement module with a transformerbased U-Net architecture adapted from [36], which was originally proposed for image denoise.



Figure 2. Architecture of the frame feature encoder. Image features are extracted by a stacked hourglass network. The locations to extract the local image features are calculated by projecting the 3D sampled point on two views. Subsequently, point features are derived by concatenating the local image features and the 3D coordinates of the sampled 3D points. Finally, MLP encodes all point features into one compact frame feature.

3. Method

3.1. Overview

The proposed biomechanics-aware network first encodes each frame into a frame feature using a stacked hourglass network [21]. Similar to previous works [23,31], we utilize two different views for each frame to provide an absolute 3D kinematic scale, which is essential for many clinical and sports science applications. Then a spatio-temporal feature refinement is performed to jointly consider spatial and temporal information to optimize the features across frames. The final output is the joint parameters of the OpenSim model [6]. To achieve accurate kinematics annotations, we create a synthetic dataset ODAH with the combination of skeleton from OpenSim [6,27] model, meshes from SMPL-X [24], and acquire the body shape and motion for the mesh from the subject-specific OpenSim skeletal model and joint angles.

3.2. Biomechanics-aware network

3.2.1 Frame Feature Encoder

We first segment the human body using the method proposed in [11] and crop the tight bounding box containing the human body. The biomechanics-aware network comprises a frame feature encoder and a spatio-temporal refinement module. The frame feature encoder is in charge of generating frame features from two-view frames. More specifically, the image feature of each view in a frame is extracted by a stacked hourglass network, which has demonstrated exceptional performance in human pose estimation tasks [21]. To effectively combine the image features of two views, we first sample candidate 3D points based on the camera rays and the camera parameters. Then, N points are randomly selected from candidates whose 2D projections



Figure 3. The overview of the proposed spatio-temporal feature refinements. With the sequence of frame features from the frame feature encoder across frames, this refinement architecture treats the feature as a 2D image to process. This process results in a sequence of joint angles and a set of body segments scales with global optimization across frames.

on both views fall within the human segmentation masks detected by YOLO [11]. Last, the point feature, denoted as $\mathbf{z}_i^{point} \in \mathbb{R}^L$ for point *i*, is obtained by concatenating the 3D point coordinates and the image local features allocated from two views.

To generate one compact frame feature based on the given N point features, we perform a two-stage feature encoding. In the first stage, each point feature is transformed into a more compact representation, denoted as $\tilde{\mathbf{z}}_{i}^{point} \in \mathbb{R}^{L'}$ for point *i*, using a shared MLP. In the second stage, the resulting N compact point features in frame *j* are concatenated and further encoded into a frame feature $\mathbf{Z}_{j}^{frame} \in \mathbb{R}^{D}$, for frame *j*, using a second MLP. The architecture of the frame feature encoder is shown in Figure 2.

3.2.2 Spatio-temporal Feature Refinement

The spatio-temporal feature refinement is applied to refine a sequence of frame features with temporal information. As shown in Figure 3, this module adopts a transformer-based U-Net architecture to extract multi-range spatio-temporal features from a given sequence of frame features. The adaptations include the downsizing of the U-Net, the temporalonly downsampling and upsampling, and the removal of the skip connection from input to output.

The input to the transformer-based U-Net is generated by concatenating a sequence of frame features along the temporal axis to create a feature map, denoted as $\mathbf{Z}_0^{seq} \in \mathbb{R}^{D \times T \times 1}$, where *T* is the number of frames in the sequence. The feature map then undergoes processing through the transformer blocks and the U-Net architecture to extract multi-resolution spatio-temporal features along the contracting and expanding paths of the U-Net. During the contracting path, the feature map at level l, $\mathbf{Z}_l^{seq} \in \mathbb{R}^{D \times t \times c}$, is downsampled to $\mathbf{Z}_{l+1}^{seq} \in \mathbb{R}^{D \times (t/r) \times rc}$, where *r* represents the downsampling factor from level *k* to level l + 1, *t* is the temporal length, and c is the number of channels. During the expanding path, refined features are generated by hierarchically combining the latent features from different levels using multiple transformer blocks. Finally, regression heads are employed to output a sequence of per-frame joint angles $\hat{\theta} \in \mathbb{R}^C$ and the per-sequence body segment scales $\hat{s} \in \mathbb{R}^{B \times 3}$, where C denotes the number of predicted joint angles and B represents the number of body segments.

3.3. Loss Function

For supervision, we utilize a range of factors such as joint angles, body segment scales, biomechanical constraints, and keypoint positions. The final loss function is written as:

$$\mathcal{L}_{total} = \mathcal{L}_{angle} + \mathcal{L}_{scale} + \mathcal{L}_{bio} + \lambda \cdot \mathcal{L}_{pos}, \quad (1)$$

where λ is the weight for \mathcal{L}_{pos} .

Joint angles. The joint angles refer to the coordinates defined in the generic OpenSim skeletal model. The angle loss is calculated differently based on whether the rotation of the corresponding joint is constrained.

The angle loss for free joints, denoted as $\mathcal{L}_{\theta f}$, is calculated as in Eq. (2). This term measures the L1 distance between the predicted and the ground truth when the angles are represented on a unit circle as commonly seen in trigonometry. This representation helps to avoid singularities and angle ambiguities caused by free rotation angles.

$$\mathcal{L}_{\theta^{f}} = \frac{1}{T} \sum_{t=0}^{T-1} \| \hat{a}_{t} - a_{t} \|_{1}, \qquad (2)$$

where $\hat{a}_t = (\cos \hat{\theta}_t^f, \sin \hat{\theta}_t^f)$, $a_t = (\cos \theta_t^f, \sin \theta_t^f)$, $\hat{\theta}_t^f$ and θ_t^f are the predicted and ground truth joint angles of the free joints at time t, respectively. T is the number of frames in a sequence.

The angle loss for constrained joints, denoted as \mathcal{L}_{θ^c} , is calculated as the L1 distance between the predicted and ground truth joint angles:

$$\mathcal{L}_{\boldsymbol{\theta}^c} = \frac{1}{T} \sum_{t=0}^{T-1} \| \hat{\boldsymbol{\theta}}_t^c - \boldsymbol{\theta}_t^c \|_1, \qquad (3)$$

where $\hat{\theta}_t^c$ and θ_t^c are the predicted and ground truth joint angles of the constrained joints at time t, respectively. T is the number of frames in a sequence.

Biomechanical constraints. Predefined constraints are commonly used to ensure biomechanical plausibility by regulating movement. The constraints are further imposed on the network by incorporating $\mathcal{L}_{\theta}^{bio}$ in the loss function to penalize joint angle predictions that violate the constraints

as in [30]. The calculation for $\mathcal{L}_{\theta}^{bio}$ is:

$$\mathcal{L}_{bio} = \frac{1}{T} \sum_{t=0}^{T-1} \| \left(\hat{\boldsymbol{\theta}}_{t}^{c} \geq \boldsymbol{\theta}_{max}^{c} \right) \cdot \left(\hat{\boldsymbol{a}}_{t} - \boldsymbol{a}_{max} \right) \|_{1} + \| \left(\hat{\boldsymbol{\theta}}_{t}^{c} \leq \boldsymbol{\theta}_{min}^{c} \right) \cdot \left(\hat{\boldsymbol{a}}_{t} - \boldsymbol{a}_{min} \right) \|_{1},$$
(4)

where \hat{a}_t , a_{min} , and a_{max} are derived as in Eq. (2). $\hat{\theta}_t^c$ is the predicted joint angles, and $[\theta_{min}^c, \theta_{max}^c]$ is the allowed range for each joint. The term T denotes the number of frames in a sequence.

Body segment scales. The body segment scale loss, denoted as \mathcal{L}_{scale} , is the L1 distance between predicted and ground truth body segment scales. The calculation is:

$$\mathcal{L}_{scale} = \frac{1}{B} \sum_{i=0}^{B-1} \|\hat{s}_i - s_i\|_1,$$
(5)

where \hat{s}_i and s_i represent the predicted and ground truth body segment scales of body segment *i*, respectively. *B* is the total number of body segments.

Keypoints. We define the position of the joints and mass center of the body segments in the OpenSim skeletal model as the keypoints for loss calculation. Supervision of the keypoint positions implicitly considers the body segment scales and the skeleton topology in 3D space.

Given the joint angles at frame t, denoted as $\hat{\theta}_t \in \mathbb{R}^J$, and the body segment scales, denoted as $\hat{s} \in \mathbb{R}^{B \times 3}$, the keypoint $\hat{P}_t \in \mathbb{R}^{K \times 3}$ are derived as $\Phi_{forward}(\hat{\theta}_t, \hat{s})$, where K is the number of keypoints, $\Phi_{forward}$ is the kinematics forward function define in the OpenSim model.

The keypoint position loss \mathcal{L}_{pos} is defined as the L1 distance between the prediction and the ground truth keypoint positions. The derivation is:

$$\mathcal{L}_{pos} = \frac{1}{TK} \sum_{t=0}^{T-1} \sum_{i=0}^{K-1} \|\hat{\boldsymbol{p}}_{i,t} - \boldsymbol{p}_{i,t}\|_{1},$$
(6)

where $\hat{p}_{i,t}$ and $p_{i,t}$ are the predicted and the ground truth keypoint positions, respectively. *T* is the number of frames, and *K* is the number of keypoints. Note that all the positions are relative to the pelvis.

3.4. Synthetic Data Generation

We developed a synthetic data generation pipeline to rig an SMPL-X [24] model against a full-body OpenSim skeletal model [26] and its associated joint angles that produce a variety of human motions (Fig 4). The OpenSim skeletal model and the joint angles serve as the inputs to the pipeline. The subject-specific SMPL-X models were fit to the joint and marker locations extracted from the OpenSim skeletal model. In the case of the AMASS dataset, we defined virtual markers on AMASS mesh vertices to obtain marker trajectories. The joint angle sequences were computed from



Figure 4. An overview of our synthetic data generation pipeline. We first register the OpenSim skeletal model to the SMPL-X mesh; followed by optimizing the body shape and motion parameters of the mesh to fit the subject-specific OpenSim skeletal model and joint angles. Finally, we simulate real-world environments with scene and camera settings to render the synthetic video data.

the marker trajectories by scaling and performing inverse kinematics in OpenSim.

Skeleton and mesh registration. An OpenSim model is composed of rigid bodies (bones) that are connected by joints. Joints connect two reference frames: one on the parent and one on the child body of the joint, which coincide in space at the joint center. The SMPL-X and OpenSim skeletal model do not share the same joint definitions. To remove this discrepancy, we trained a joint regressor for the SMPL-X model that located the joint keypoints corresponding to the OpenSim model joint centers based on the position of the SMPL-X vertices. A variety of poses are included during the training of the joint regressor to avoid overfitting.

Shape and motion acquisition. We first utilized MoSh++ (Motion and Shape capture) [19] to create the initial human mesh. Given marker trajectories and a marker layout on the mesh, MoSh++ generated SMPL-X mesh sequences. The marker layout was defined manually by identifying corresponding SMPL-X vertices as OpenSim markers. From MoSh++, we obtained the initial SMPL-X shape β and pose θ for the given skeletal model and its motion. With additional subject-specific skin or clothing shape for the mesh and a subject-specific OpenSim skeleton, we further optimized β and pose θ to minimize differences between virtual markers on the mesh and the markers on the skeletal model.

To animate the resulting subject-specific SMPL-X mesh, its pose parameters, θ are optimized frame by frame to minimize the distance between joint keypoints and marker positions of the SMPL-X mesh and those on the OpenSim skeletal model driven by the input joint angles.

Scene settings. To augment appearances, we used four types of upper body clothing, ranging from vests to long-sleeved shirts, and four types of lower body clothing, from shorts to trousers. Each type of the clothing is randomly combined with 5 different textures in each trial. We used multiple area light sources evenly distributed on the ceiling. **Cameras settings.** We employed two static cameras for

video rendering that were positioned at a height of 1.1 ± 0.1 meters. One camera captured the frontal view, while the other one captured the sagittal view. To enhance diversity, the positions of the cameras are randomly perturbed within a small range. Both cameras had a fixed focal length of 33 mm. The sensor fit of the cameras was set to horizontal with the sensor width set to 36 mm.

Rendering. The videos were rendered using the Metalaccelerated BLENDER_EEVEE engine in Blender 3.5. The video resolution was set to 1080 by 720 and collected at a framerate of 60 fps. Motion blur effects are disabled. The videos are encoded in the H264 format with a mediumquality configuration.

4. Experiments

4.1. Datasets

ODAH. In the proposed ODAH, there are 56 synthetic subjects generated from the BMLMovi [9] in a subset of the AMASS [19] dataset. The actions include running, jogging, jumping, sideways, scratching head, throwing and catching, hand clapping, walking, checking watch, sitting down, hand waving, crossing arms, stretching, kicking, phone talking, taking photos, pointing, vertical jumping, crawling, crossing legs while sitting, and freestyle. Overall, ODAH has 1132 videos in 60 fps, and each video has a duration of around 10 seconds. We use 42 subjects for training, six subjects for validation, and eight subjects for testing.

To evaluate the performance of the proposed Testing. biomechanics-aware network and the generalization capability in real-world settings, we tested the method on two real datasets, OpenCap [31] and BMLMovi [9], as well as the proposed synthetic ODAH. OpenCap consists of ten subjects performing actions including walking, squatting, rising from a chair, drop jumps, and the asymmetric counterparts. It was recorded using five RGB cameras and a marker-based motion capture system. OpenCap also provides processed marker data and kinematics annotations for a full-body OpenSim skeletal model. BMLMovi involves 90 subjects performing 21 actions, recorded using two cameras and a marker-based motion capture system. Since BMLmovi does not include kinematics annotations, we utilized the provided marker data, the OpenSim Scale tool, and the OpenSim IK tool to acquire accurate joint angles and body segment scales for ground truth. We recognize the root-mean-square of the marker data and the fitted ground truth is between one to three centimeters. We excluded two actions, namely crawling and crossing legs while sitting, from our test set due to the challenges in fitting the underlying OpenSim skeletal model to the bending movement.

4.2. Implementation Details

We first train the stacked hourglass network for 14 epochs with a batch size set to eight, and then integrate it with the spatio-temporal refinement network for another five epochs with a batch size set to two. We use the Adam optimizer [14] with $\beta_1 = 0.5$ and $\beta_2 = 0.999$ is applied, and the learning rate is set to $1 \times 5 \times 10^{-5}$. Input frames are resized to $3 \times 256 \times 256$ before feature extraction. The loss weight λ is set to 100. We train our method with a single A40 GPU.

We set N = 500 3D points for point feature extraction, the point feature length L = 1027, and the reduced one L' = 32. The final frame feature length D = 102. The sequence length T is set to 64. The U-Net consists of three encoder-decoder levels, with downsampling factors r set to two for level two and four for level three. The number of transformer blocks for each level is [2, 4, 6], the number of attention heads is [1, 2, 4], and the number of channels is [48, 96, 384].

In the generic OpenSim skeletal model, the number of the joint angles J is 36. The number of body segments B is 22, and the number of joint keypoints K is 44. Our generic OpenSim skeletal model has nine unconstrained joint angles, controlling the pelvis, and left and right arms. Therefore, only 17 joint angles are restricted by biomechanical constraints.

4.3. Metrics

To make the evaluation focus more on the performance of the 3D joint kinematics, Procrustes alignment [10] is applied as the first step to align the global translation, rotation, and scaling between the predictions and the ground truth.

Mean Absolute Error (MAE) is used to evaluate joint angle error as

$$MAE_{angle} = \frac{1}{T} \sum_{t=1}^{T} \|\hat{\boldsymbol{\theta}}_t - \boldsymbol{\theta}_t\|_1,$$
(7)

where $\hat{\theta}_t$ is the predicted angles, θ_t is the ground truth angles, and T is the number of frames in a sequence.

Mean Per Joint Position Error with Procrustes Alignment (PA-MPJPE) is widely used in 3D human motion estimation to measure the Euclidean distance between the predicted and the ground truth 3D joint keypoint positions. The calculation is

$$PA-MPJPE = \frac{1}{TB} \sum_{t=0}^{T-1} \sum_{i=0}^{B-1} \|\hat{\boldsymbol{p}}_{i,t}^{PA} - \boldsymbol{p}_{i,t}\|_2, \quad (8)$$

where B is the number of joint keypoints, $\hat{p}_{i,t}^{PA}$ is the predicted joint keypoint positions relative to the root joint after Procrustes alignment and $p_{i,t}$ represents the ground truth joint keypoint positions relative to the root joint.

	$MAE_{angle} (deg.) \downarrow$				PA-MPJPE (mm) ↓			
	Pose2Sim	OpenCap	D3KE	Ours	Pose2Sim	OpenCap	D3KE	Ours
OpenCap	9.77	7.37	12.47	9.65	67.69	53.86	95.64	74.85
BMLmovi	10.68	15.27	9.43	9.73	58.55	113.45	49.31	64.68
ODAH	10.99	10.78	12.11	4.81	60.67	64.27	74.29	27.09
Mean	10.48	11.14	11.34	8.06	62.30	77.19	73.08	55.54

Table 1. Comparison between Pose2Sim, OpenCap, D3KE, and our method (columns) with the joint angle error (MAE) and joint keypoint error (PA-MPJPE). The evaluation is performed on OpenCap, BMLMovi, and ODAH (rows) datasets. The last row shows the average error across all datasets.

4.4. Baselines

We compare the performance of the proposed method against state-of-the-art methods of two multi-step methods Pose2Sim [23] and OpenCap [31], and one end-to-end method D3KE [1].

For Pose2Sim and OpenCap, frontal and sagittal views are taken as inputs and the Body25 model in OpenPose [2] is used as the 2D landmark detection backbone. We use Pose2Sim default Butterworth low-pass filter with a cut-off frequency of six Hz as the smoothing filter. For OpenCap, the cut-off frequency is set to half of the framerate. To ensure a fair comparison, we disable the landmark synchronization and the video trimming in OpenCap. For D3KE, it only requires a single-view image as input that we take frontal view. We acknowledge that it is not a totally fair comparison due to different inputs. However, it is natural of the D3KE method with only single-view as the input. We choose the transformer-based temporal model with the sequence length set to 243 frames for the D3KE method. Additionally, we only compare joint angles present in our generic OpenSim skeletal model and exclude arm flexion since D3KE's generic OpenSim skeletal model does not have a joint angle defined for arm flexion for the sake of fair comparison.

4.5. Comparison with State-of-the-Art

We compare the proposed biomechanics-aware network with state-of-the-art methods and show the results in Table 1. We conduct experiments on three datasets Open-Cap, BMLmovi, and our synthetic ODAH as shown in each row, and compare three methods Pose2Sim, OpenCap, and D3KE as shown in columns. We report the errors on both joint angle error (MAE_{angle}) and joint keypoint position error (PA-MPJPE). We also average the performance across three datasets for each method for comparison. From the table, we can see that our method achieves the best performances in terms of averaged joint angle and keypoint position errors. Our method is significantly better than the second-best baseline Pose2Sim with 23% improvement (8.06 degrees vs. 10.48 degrees) on joint angle estimation error and 10% improvement (55.54 mm vs. 62.30 mm) on

	$MAE_{angle} (deg.) \downarrow$			PA-MPJPE (mm) \downarrow		
	small	medium	large	small	medium	large
OpenCap	11.04	10.64	9.65	92.62	81.79	74.85
BMLmovi	10.54	9.95	9.73	82.91	69.53	64.68
ODAH	6.23	5.39	4.81	39.71	33.60	27.09
Mean	9.27	8.66	8.06	71.75	61.64	55.54

Table 2. Ablation study on different data sizes. We configure ODAH into small, medium, and large subsets. The evaluation is performed on OpenCap, BMLMovi, and ODAH (rows) with joint angle error (MAE) and joint position error (PA-MPJPE).

the joint keypoint position error. It clearly shows the advantage of the proposed model over the previous state-of-theart methods with our advanced neural network architecture and spatial-temporal feature refinement.

Note that the OpenCap method was trained on the Open-Cap training dataset, D3KE was trained on the BMLmovi dataset training set, and biomechanics-aware network was trained on ODAH training set. Therefore, each of these methods achieves the best performance on the corresponding test set with their specific domain knowledge. However, our biomechanics-aware network shows strong generalizability by achieving the best averaged performances across the three datasets. Even if we remove results on our synthetic dataset ODAH from the Table 1, our biomechanicsaware network still can achieve the best averaged performances across the rest two real-world datasets in terms of the joint angle estimation task. Since biomechanics-aware network is solely trained on synthetic ODAH without finetuning on any real data, the superior average results show the effectiveness of using synthetic data improving domain generalization. These results also confirm the biomechanical fidelity of the proposed synthetic dataset ODAH.

4.6. Ablation Study

4.6.1 Training Data Size

The benefit of synthetic data generation is its great potential for large amounts of data creation. To verify the effectiveness of the proposed ODAH, we configure the ODAH into three subsets. The small set has 10 subjects and 195 clips,

	MAE _{an}	$_{gle}$ (deg.) \downarrow	PA-MPJ	PE (mm)↓	PA-MPJVE (mm/s)↓		
	Frame	Ours	Frame	Ours	Frame	Ours	
OpenCap	9.91	9.65	76.59	74.85	906.65	471.37	
BMLmovi	9.82	9.73	65.59	64.68	394.71	254.02	
ODAH	4.61	4.81	27.36	27.09	296.45	155.71	
Mean	8.11	8.06	56.51	55.54	532.60	293.70	

Table 3. Ablation study on the effectiveness of temporal information. Frame-based prediction is the proposed network without spatio-temporal refinement, and we compare this baseline to our final model. The evaluation is performed on OpenCap, BMLMovi, and ODAH (rows) with joint angle error (MAE), joint position error (PA-MPJPE), and joint velocity error (MPJVE).

	MAE	C _{angle} (de	eg.)↓	PA-MPJPE (mm) \downarrow		
\mathcal{L}_{angle}	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
\mathcal{L}_{scale}	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark	\checkmark
\mathcal{L}_{bio}		\checkmark	\checkmark		\checkmark	\checkmark
\mathcal{L}_{pos}			\checkmark			\checkmark
OpenCap	10.68	10.20	9.65	90.74	92.64	74.85
BMLmovi	9.78	9.75	9.73	78.51	84.09	64.68
ODAH	4.45	4.68	4.81	36.45	37.24	27.09
Mean	8.30	8.21	8.06	68.57	71.32	55.54

Table 4. Ablation study on the loss function. We incrementally test the effects of biomechanical constraints (+ \mathcal{L}_{angle}) and the keypoint positions (+ \mathcal{L}_{pos}). The evaluation is performed on Open-Cap, BMLMovi, and ODAH (rows) with joint angle error (MAE) and joint position error (PA-MPJPE).

the medium set has 20 subjects and 397 clips, and the large set has 42 subjects and 841 clips. Note the large set is the full ODAH dataset. We train our model on each of these subsets and test on the test set of ODAH. We show results in Table 2 across three datasets as each row and training subsets as each column. From the table we can that the performances of the large subsets achieve the best performances across all datasets, thus, indicating that more data ensures better performances for kinematics estimation in terms of both joint angles and position. Note the large subset training does not cause the overfitting problem on the two real-world datasets OpenCap and BMLmovie. It confirms again the high biomechanical fidelity of the proposed ODAH which benefits the model training for the generalizability. Since it is feasible to generate large synthetic data with the proposed synthetic data generation method, there is great potential for future method development.

4.6.2 Frame-based vs. Sequence-based

In the method design, we have two components including the frame feature encoder for the frame-based prediction and the spatio-temporal feature refinement for the sequencebased prediction. Essentially, the frame-based prediction can already perform the 3D kinematic estimation task. To investigate the effectiveness of feature refinement, we perform a comparison between the frame feature encoder only and the completed model. Specifically, frame-based prediction is implemented by removing the spatio-temporal refinement module in the proposed framework, and directly outputs the joint angles and body segment scales. For experiments, besides the joint angle and position errors, we also calculate the Mean Per Joint Velocity Error (MPJVE) since it indicates the smoothness of the motion. The results are shown in Table 3, which indicates that the completed model achieves better results than the frame-based prediction baseline in general. Moreover, the improvements in the motion coherence, as indicated by Mean Per Joint Velocity Error (MPJVE), are much more significant.

4.6.3 Loss Function

Finally, We further conduct tests to examine the contribution of each loss term. The joint angle loss \mathcal{L}_{angle} and body segment scale loss \mathcal{L}_{scale} are essential for the training that must be included. Therefore, we only examined the effectiveness of biomechanical constraints \mathcal{L}_{bio} and keypoint position loss \mathcal{L}_{pos} . The results are present in Table 4. Although the \mathcal{L}_{bio} does not bring significant improvement numerically, we notice there could be extreme joint angle and position predictions without it. The implicit weights on joint angles introduced by \mathcal{L}_{pos} can improve both the joint angle error and the joint keypoint position error. Combining all loss terms, our model can achieve the optimal performance across the three datasets.

5. Conclusion

We propose a novel end-to-end biomechanics-aware network that is solely trained on self-created synthetic data. The proposed method utilizes direct mapping from two input views to the frame features and is refined with a spatio-temporal refinement module. We create synthetic data by combining the SMPL-X model from the computer vision community and the OpenSim skeletal model from the biomechanics community to provide accurate ground truth of kinematics. By solely trained on synthetic data, our proposed method achieves superior performance with the best generalization across multiple datasets. It demonstrates the effectiveness of the proposed method for improving kinematics estimation and enhancing domain generalization, and also validates the biomechanical fidelity of the proposed dataset generation pipeline.

The limitations include visual quality and motion variations in the synthetic dataset, and the size of the proposed architecture. Future research could focus on improving the visual quality via adversarial training, including variant actions from real humans, and extending the method.

References

- Marian Bittner, Wei-Tse Yang, Xucong Zhang, Ajay Seth, Jan van Gemert, and Frans CT van der Helm. Towards single camera human 3d-kinematics. *Sensors*, 23(1):341, 2022. 2, 3, 7
- [2] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Openpose: Realtime multi-person 2d pose estimation using part affinity fields, 2019. 2, 7
- [3] Tianlang Chen, Chen Fang, Xiaohui Shen, Yiheng Zhu, Zhili Chen, and Jiebo Luo. Anatomy-aware 3d human pose estimation with bone-based pose decomposition. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(1):198–209, 2021. 2, 3
- [4] Mia Chiquier and Carl Vondrick. Muscles in action. In ICCV, 2023. 2
- [5] Enric Corona, Gerard Pons-Moll, Guillem Alenyà, and Francesc Moreno-Noguer. Learned vertex descent: a new direction for 3d human model fitting. In *Computer Vision– ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 146–165. Springer, 2022. 2, 3
- [6] Scott L Delp, Frank C Anderson, Allison S Arnold, Peter Loan, Ayman Habib, Chand T John, Eran Guendelman, and Darryl G Thelen. Opensim: open-source software to create and analyze dynamic simulations of movement. *IEEE transactions on biomedical engineering*, 54(11):1940–1950, 2007. 2, 3
- [7] Sai Kumar Dwivedi, Cordelia Schmid, Hongwei Yi, Michael J Black, and Dimitrios Tzionas. Poco: 3d pose and shape estimation using confidence. In *International conference on 3D vision (3DV)*, 2024. 2
- [8] Glenn S Fleisig, Jonathan S Slowik, Derek Wassom, Yuki Yanagita, Jasper Bishop, and Alek Diffendaffer. Comparison of marker-less and marker-based motion capture for baseball pitching kinematics. *Sports Biomechanics*, pages 1–10, 2022. 1
- [9] Saeed Ghorbani, Kimia Mahdaviani, Anne Thaler, Konrad Kording, Douglas James Cook, Gunnar Blohm, and Nikolaus F Troje. Movi: A large multi-purpose human motion and video dataset. *Plos one*, 16(6):e0253157, 2021. 2, 6
- [10] John C Gower. Generalized procrustes analysis. *Psychometrika*, 40:33–51, 1975. 6
- [11] Glenn Jocher, Chaurasia Ayush, and Jing Qiu. Yolo by ultralytics. 3, 4
- [12] Marilyn Keller, Keenon Werling, Soyong Shin, Scott Delp, Sergi Pujades, C Karen Liu, and Michael J Black. From skin to skeleton: Towards biomechanically accurate 3d digital humans. ACM Transactions on Graphics (TOG), 42(6):1–12, 2023. 2
- [13] Marilyn Keller, Silvia Zuffi, Michael J Black, and Sergi Pujades. Osso: Obtaining skeletal shape from outside. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 20492–20501, 2022. 2
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014. 6

- [15] Wenhao Li, Hong Liu, Runwei Ding, Mengyuan Liu, and Pichao Wang. Lifting transformer for 3d human pose estimation in video. *arXiv preprint arXiv:2103.14304*, 2, 2021.
 3
- [16] Wenhao Li, Hong Liu, Runwei Ding, Mengyuan Liu, Pichao Wang, and Wenming Yang. Exploiting temporal contexts with strided transformer for 3d human pose estimation. *IEEE Transactions on Multimedia*, 2022. 2, 3
- [17] Wenhao Li, Hong Liu, Hao Tang, Pichao Wang, and Luc Van Gool. Mhformer: Multi-hypothesis transformer for 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13147–13156, 2022. 3
- [18] Ruixu Liu, Ju Shen, He Wang, Chen Chen, Sen-ching Cheung, and Vijayan Asari. Attention mechanism exploits temporal contexts: Real-time 3d human pose reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5064–5073, 2020. 3
- [19] 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, 5, 6
- [20] Lars Mündermann, Stefano Corazza, and Thomas P Andriacchi. The evolution of methods for the capture of human movement leading to markerless motion capture for biomechanical applications. *Journal of neuroengineering and rehabilitation*, 3(1):1–11, 2006. 1
- [21] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part VIII 14, pages 483–499. Springer, 2016. 3
- [22] Kien X Nguyen, Liying Zheng, Ashley L Hawke, Robert E Carey, Scott P Breloff, Kang Li, and Xi Peng. Deep learningbased estimation of whole-body kinematics from multi-view images. *Computer Vision and Image Understanding*, page 103780, 2023. 3
- [23] David Pagnon, Mathieu Domalain, and Lionel Reveret. Pose2sim: An open-source python package for multiview markerless kinematics. *Journal of Open Source Software*, 2022. 2, 3, 7
- [24] 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 *Proceedings IEEE Conf.* on Computer Vision and Pattern Recognition (CVPR), pages 10975–10985, 2019. 2, 3, 5
- [25] Dario Pavllo, Christoph Feichtenhofer, David Grangier, and Michael Auli. 3d human pose estimation in video with temporal convolutions and semi-supervised training. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7753–7762, 2019. 3
- [26] Apoorva Rajagopal, Christopher L Dembia, Matthew S De-Mers, Denny D Delp, Jennifer L Hicks, and Scott L Delp. Full-body musculoskeletal model for muscle-driven simulation of human gait. *IEEE transactions on biomedical engineering*, 63(10):2068–2079, 2016. 5

- [27] Ajay Seth, Jennifer L Hicks, Thomas K Uchida, Ayman Habib, Christopher L Dembia, James J Dunne, Carmichael F Ong, Matthew S DeMers, Apoorva Rajagopal, Matthew Millard, et al. Opensim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement. *PLoS computational biology*, 14(7):e1006223, 2018. 2, 3
- [28] 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
- [29] Ke Song, Todd J Hullfish, Rodrigo Scattone Silva, Karin Grävare Silbernagel, and Josh R Baxter. Markerless motion capture estimates of lower extremity kinematics and kinetics are comparable to marker-based across 8 movements. *Journal of Biomechanics*, 157:111751, 2023. 2
- [30] Adrian Spurr, Umar Iqbal, Pavlo Molchanov, Otmar Hilliges, and Jan Kautz. Weakly supervised 3d hand pose estimation via biomechanical constraints. In *Computer Vision– ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVII 16*, pages 211– 228. Springer, 2020. 5
- [31] Scott D Uhlrich, Antoine Falisse, Łukasz Kidziński, Julie Muccini, Michael Ko, Akshay S Chaudhari, Jennifer L Hicks, and Scott L Delp. Opencap: 3d human movement dynamics from smartphone videos. *bioRxiv*, pages 2022–07, 2022. 1, 2, 3, 6, 7
- [32] Eline Van der Kruk and Marco M Reijne. Accuracy of human motion capture systems for sport applications; state-of-theart review. *European journal of sport science*, 18(6):806– 819, 2018. 1
- [33] Size Wu, Sheng Jin, Wentao Liu, Lei Bai, Chen Qian, Dong Liu, and Wanli Ouyang. Graph-based 3d multi-person pose estimation using multi-view images. In *Proceedings of* the IEEE/CVF international conference on computer vision, pages 11148–11157, 2021. 3
- [34] Jingwei Xu, Zhenbo Yu, Bingbing Ni, Jiancheng Yang, Xiaokang Yang, and Wenjun Zhang. Deep kinematics analysis for monocular 3d human pose estimation. In *Proceedings of the IEEE/CVF Conference on computer vision and Pattern recognition*, pages 899–908, 2020. 3
- [35] B.-M. You, P. Siy, W. Anderst, and S. Tashman. In vivo measurement of 3-d skeletal kinematics from sequences of biplane radiographs: Application to knee kinematics. *IEEE Transactions on Medical Imaging*, 20(6):514–525, 2001. 2
- [36] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 5728– 5739, 2022. 2, 3
- [37] Jianfeng Zhang, Yujun Cai, Shuicheng Yan, Jiashi Feng, et al. Direct multi-view multi-person 3d pose estimation. Advances in Neural Information Processing Systems, 34:13153–13164, 2021. 3
- [38] Jinlu Zhang, Zhigang Tu, Jianyu Yang, Yujin Chen, and Junsong Yuan. Mixste: Seq2seq mixed spatio-temporal encoder

for 3d human pose estimation in video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13232–13242, 2022. 2, 3

[39] Ce Zheng, Sijie Zhu, Matias Mendieta, Taojiannan Yang, Chen Chen, and Zhengming Ding. 3d human pose estimation with spatial and temporal transformers. In *Proceedings* of the IEEE/CVF International Conference on Computer Vision, pages 11656–11665, 2021. 3