

HybrIK: A Hybrid Analytical-Neural Inverse Kinematics Solution for 3D Human Pose and Shape Estimation

Jiefeng Li, Chao Xu, Zhicun Chen, Siyuan Bian, Lixin Yang, Cewu Lu[†]
Shanghai Jiao Tong University, China

{ljf_likit, xuchao.19962007, zhicun_chen, biansiyuan, siriusyang, lucewu}@sjtu.edu.cn

Abstract

Model-based 3D pose and shape estimation methods reconstruct a full 3D mesh for the human body by estimating several parameters. However, learning the abstract parameters is a highly non-linear process and suffers from image-model misalignment, leading to mediocre model performance. In contrast, 3D keypoint estimation methods combine deep CNN network with the volumetric representation to achieve pixel-level localization accuracy but may predict unrealistic body structure. In this paper, we address the above issues by bridging the gap between body mesh estimation and 3D keypoint estimation. We propose a novel hybrid inverse kinematics solution (HybrIK). HybrIK directly transforms accurate 3D joints to relative body-part rotations for 3D body mesh reconstruction, via the twist-and-swing decomposition. The swing rotation is analytically solved with 3D joints, and the twist rotation is derived from the visual cues through the neural network. We show that HybrIK preserves both the accuracy of 3D pose and the realistic body structure of the parametric human model, leading to a pixel-aligned 3D body mesh and a more accurate 3D pose than the pure 3D keypoint estimation methods. Without bells and whistles, the proposed method surpasses the state-of-the-art methods by a large margin on various 3D human pose and shape benchmarks. As an illustrative example, HybrIK outperforms all the previous methods by **13.2 mm MPJPE** and **21.9 mm PVE** on 3DPW dataset. Our code is available at <https://github.com/Jeff-sjtu/HybrIK>.

1. Introduction

Recovering the 3D surface from a monocular RGB image is a fundamentally ill-posed problem. It has a wide range of application scenarios [52, 31, 27, 28, 8]. With the development of the parametric statistical human body shape models [2, 29, 47], a realistic and controllable 3D mesh of

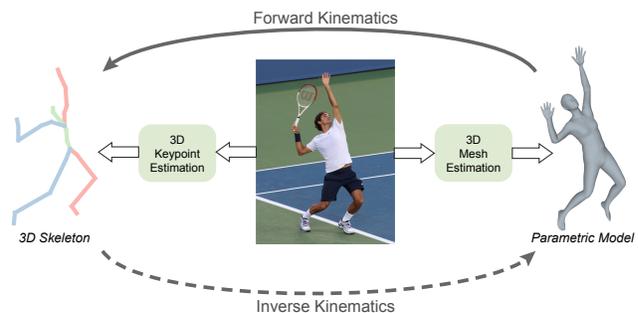


Figure 1. **Closing the loop between the 3D skeleton and the parametric model via HybrIK.** A 3D skeleton predicted by the neural network can be transformed into a parametric body mesh by inverse kinematics without loss of accuracy. The parametric body mesh can generate structural realistic 3D skeleton by forward kinematics.

human body can be generated from only a few parameters, e.g. shape parameters and relative rotations of body parts. Recent studies develop the model-based methods [7, 18, 23] to obtain these parameters from the monocular RGB input and produce 3D pose and shape of human bodies.

Most of the model-based methods can be categorized into two classes: optimization-based approach and learning-based approach. Optimization-based approaches [13, 7, 47] estimate the body pose and shape by an iterative fitting process. The parameters of the statistical model are tuned to reduce the error between its 2D projection and 2D observations, e.g. 2D joint locations and silhouette. However, the optimization problem is non-convex and takes a long time to solve. The results are sensitive to the initialization. These issues shift the spotlight towards the learning-based approaches. With a parametric body model, learning-based approaches use neural networks to regress the model parameters directly [18, 23, 20]. But the parameter space in the statistical model is abstract, making it difficult for the networks to learn the mapping function.

This challenge prompts us to look into the field of 3D keypoint estimation. Instead of the direct regression, previous methods [48, 60] adopt volumetric heatmap as the target representation to learn 3D joint locations and have achieved

[†]Cewu Lu is the corresponding author. He is the member of Qing Yuan Research Institute and MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, China.

impressive performance. This inspires us to build a collaboration between the 3D joints and the body mesh (Fig. 1). On the one hand, the accurate 3D joints facilitate the 3D body mesh estimation. On the other hand, the shape prior in parametric body model in turn fixes the unrealistic body structure issue of the 3D keypoint estimation methods. Since the current 3D keypoint estimation methods lack explicit modelling of the distribution of body bone length, it may predict unrealistic body structures like left-right asymmetry and abnormal proportions of limbs. By leveraging the parametric body model, the presented human shape better conforms to the actual human body.

In this work, we propose a hybrid analytical-neural inverse kinematics solution (HybrIK) to bridge the gap between 3D keypoint estimation and body mesh estimation. Inverse kinematics (IK) process is the mathematical process of finding the relative rotations to produce the desired locations of body joints. It is an ill-posed problem because there is no unique solution. The core of our approach is to propose an innovative IK solution via twist-and-swing decomposition. The relative rotation of a skeleton part is decomposed into *twist* and *swing*, *i.e.* a longitudinal rotation and an in-plane rotation. In HybrIK, we composite the entire rotation recursively along the kinematic tree by analytically calculating *swing* rotation and predicting *twist* rotation. A critical characteristic of our approach is that the relative rotation estimated by HybrIK is naturally aligned with the 3D skeleton, without the need for additional optimization procedures in the previous approaches [7, 47, 23]. All operations in HybrIK are differentiable, which allows us to simultaneously train 3D joints and human body mesh in an end-to-end manner. Besides, experiments indicate that HybrIK raises the performance of body mesh estimation to the same level as 3D keypoint estimation and takes a step forward. The proposed approach is benchmarked in various 3D human pose and shape datasets, and it significantly outperforms state-of-the-art approaches [23, 38] by **21.9** mm PVE on 3DPW [67], **6.6** mm PA-MPJPE on Human3.6M [16] and **10.8** AUC on MPI-INF-3DHP [33].

The contributions of our approach can be summarized as follows:

- We propose HybrIK, a hybrid analytical-neural IK solution that converts the accurate 3D joint locations to full 3D human body mesh. HybrIK is differentiable and allows end-to-end training.
- Our approach closes the loop between the 3D skeleton and the parametric model. It fixes the alignment issue of current model-based body mesh estimation methods and the unrealistic body structure problem of 3D keypoint estimation methods at the same time.
- Our approach achieves state-of-the-art performance across various 3D human pose and shape benchmarks.

2. Related Work

3D Keypoint Estimation. Many works formulate 3D human pose estimation as the problem of locating the 3D joints of the human body. Previous studies can be divided into two categories: single-stage and two-stage approaches. Single-stage approaches [49, 54, 34, 74, 36, 61, 37, 69] directly estimate the 3D joint locations from the input image. Various representations are developed, including 3D heatmap [49], location-map [34] and 2D heatmap + z regression [74]. Two-stage approaches first estimate 2D pose and then lift them to 3D joint locations by a learned dictionary of 3D skeleton [1, 53, 64, 56, 76, 77] or regression [46, 71, 40, 11, 32, 59]. Two-stage approaches highly rely on the accurate 2D pose estimators, which have achieved impressive performance by the combination of powerful backbone network [58, 15, 42, 44, 45] and the 2D heatmap.

These privileged forms of supervision contribute to the recent performance leaps of 3D keypoint estimation. However, the human structural information is modelled implicitly by the neural network, which can not ensure the output 3D skeletons to be realistic. Our approach combines the advantages of both the 3D skeleton and parametric model to predict accurate and realistic human pose and shape.

Model-based 3D Pose and Shape Estimation. Pioneer works on the model-based 3D pose and shape estimation methods use parametric human body model [2, 29, 47] as the output target because they capture the statistics prior of body shape. Compared with the model-free methods [65, 24, 38], the model-based methods directly predict controllable body mesh, which can facilitate many downstream tasks for both computer graphics and computer vision. Bogo *et al.* [7] propose SMPLify, a fully automatic approach, without manual user intervention [57, 13]. This optimization paradigm was further extended with silhouette cues [26], volumetric grid [65], multiple people [73] and whole-body parametric model [47].

With the advances of the deep learning networks, there are increasing studies that focus on the learning-based methods, using a deep network to estimate the pose and shape parameters. Since the mapping from RGB image to shape space and relative body-part rotation is hard to learn, many works use some form of intermediate representation to alleviate this problem, such as keypoints and silhouettes [50], semantic part segmentation [43] and 2D heatmap input [63]. Kanazawa *et al.* [18] use an adversarial prior and an iterative error feedback (IEF) loop to reduce the difficulty of regression. Arnab *et al.* [4] and Kocabas *et al.* [20] exploit temporal context, while Guler *et al.* [14] use a part-voting expression and test-time post-processing to improve the regression network. Kolotouros *et al.* [23] leverage the optimization paradigm to provide extra 3D su-

pervision from unlabeled images.

In this work, we address this challenging learning problem by a transformation from the pixel-aligned 3D joints to the relative body-part rotations.

Body-part Rotation in Pose Estimation. The core of our approach is to calculate the relative rotation of human body parts through a hybrid IK process. There are several works that estimate the relative rotations in the 3D pose estimation literature. Zhou *et al.* [75] use the network to predict the rotation angle of each body joint, followed by an FK layer to generate the 3D joint coordinates. Pavllo *et al.* [51] switch to quaternions, while Yoshiyasu *et al.* [72] directly predict the 3×3 rotation matrices. Mehta *et al.* [35] first estimate the 3D joints and then use a fitting procedure to find the rotation Euler angles. Previous approaches are either limited to a hard-to-learn problem or require an additional fitting procedure. Our approach recovers the body-part rotation from 3D joint locations in a direct, accurate and feed-forward manner.

Inverse Kinematics Process. The inverse kinematics (IK) problem has been extensively studied during recent decades. Numerical solutions [6, 70, 12, 19, 68, 9] are simple ways to implement the IK process, but they suffer from time-consuming iterative optimization. Heuristic methods are efficient solutions to the IK problem. For example, CDC[30], FABRIK[3] and IK-FA[55] have a low computational cost for each heuristic iteration. In some special cases, there exist analytical solutions to the IK problem. Tolani *et al.* [62] propose a reliable algorithm by the combination of analytical and numerical methods. Kallmann *et al.* [17] solve the IK for arm linkage, *i.e.* a three-joint system. Recently, researchers have been interested in using neural networks to solve the IK problem in robotic control [10], motion retargeting [66] and hand pose estimation [41, 22].

In this work, we combine the interpretable characteristic of analytical solution and the flexibility of the neural network, introducing a feed-forward hybrid IK algorithm with twist-and-swing decomposition. Twist-and-swing decomposition is introduced by Baerlocher *et al.* [5]. The twist angles are limited based on the particular body joint. In our works, the twist angles are estimated by neural networks, which is more flexible and can be generalized to all body joints. Compared with previous analytical solutions [17] designed for specific joint linkage, our algorithm can be applied to the entire body skeleton in a direct and differentiable manner.

3. Method

In this section, we present our hybrid analytical-neural inverse kinematics solution that boosts 3D human pose and shape estimation (Fig. 2). First, in §3.1, we briefly describe

the forward kinematics process, the inverse kinematics process and the SMPL model. In §3.2, we introduce the proposed inverse kinematics solution, HybriIK. Then, in §3.3, we present the overall learning framework to estimate the pixel-aligned body mesh and realistic 3D skeleton. Finally, we provide the necessary implementation details in §3.4.

3.1. Preliminary

Forward Kinematics. Forward kinematics (FK) for human pose usually refers to the process of computing the reconstructed pose $\mathbf{Q} = \{q_k\}_{k=1}^K$, with the rest pose template $\mathbf{T} = \{t_k\}_{k=1}^K$ and the relative rotations $\mathbf{R} = \{R_{\text{pa}(k),k}\}_{k=1}^K$ as input:

$$\mathbf{Q} = \text{FK}(\mathbf{R}, \mathbf{T}), \quad (1)$$

where K is the number the body joints, $q_k \in \mathbb{R}^3$ denotes the reconstructed 3D location of the k -th joint, $t_k \in \mathbb{R}^3$ denotes the k -th joint location of the rest pose template, $\text{pa}(k)$ return the parent’s index of the k -th joint, and $R_{\text{pa}(k),k}$ is the relative rotation of k -th joint with respect to its parent joint. FK can be performed by recursively rotating the template body part from the root joint to the leaf joints:

$$q_k = R_k(t_k - t_{\text{pa}(k)}) + q_{\text{pa}(k)}, \quad (2)$$

where $R_k \in \mathbb{SO}(3)$ is the global rotation of the k -th joint with respect to the canonical rest pose space. The global rotation can be calculated recursively:

$$R_k = R_{\text{pa}(k)} R_{\text{pa}(k),k}. \quad (3)$$

For the root joint that has no parent, we have $q_0 = t_0$.

Inverse Kinematics. Inverse kinematics (IK) is the reverse process of FK, computing relative rotations \mathbf{R} that can generate the desired locations of input body joints $\mathbf{P} = \{p_k\}_{k=1}^K$. This process can be formulated as:

$$\mathbf{R} = \text{IK}(\mathbf{P}, \mathbf{T}), \quad (4)$$

where p_k denotes the k -th joint of the input pose. Ideally, the resulting rotations should satisfy the following condition:

$$p_k - p_{\text{pa}(k)} = R_k(t_k - t_{\text{pa}(k)}) \quad \forall 1 \leq k \leq K. \quad (5)$$

Similar to the FK process, we have $p_0 = t_0$ for the root joint that has no parent. While the FK problem is well-posed, the IK problem is ill-posed because there is either no solution or because there are many solutions to fulfill the target joint locations.

SMPL Model. In this work, we employ the SMPL [29] parametric model for human body representation. SMPL allows us to use shape parameters and pose parameters

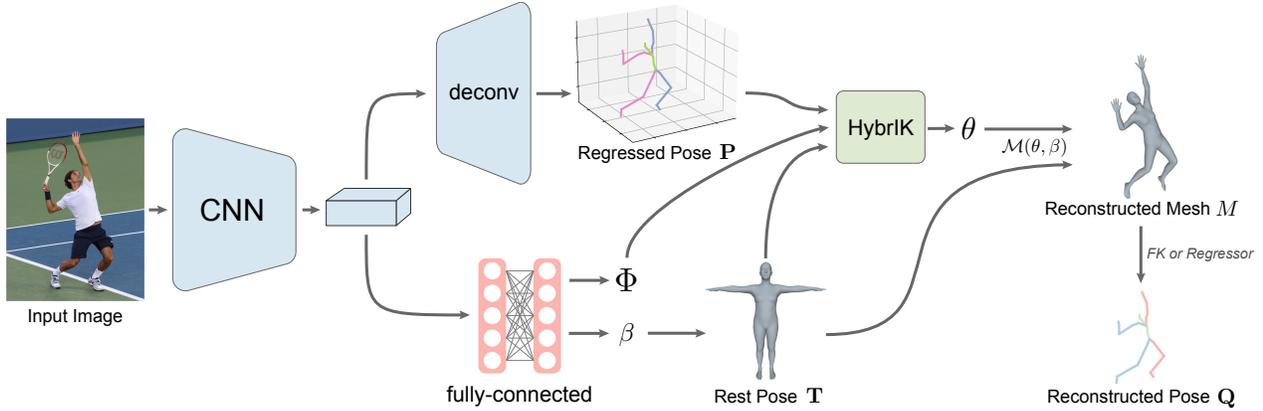


Figure 2. **Overview of the proposed framework.** A 3D heatmap is generated by the deconvolution layers and used to regress the 3D joints \mathbf{P} . The shape parameters β and the *twist* angle Φ are learned from the visual cues through the fully-connected layers. These results are then sent to the HybrIK process to solve the relative rotation, *i.e.* the pose parameters θ . Finally, with the pose and shape parameters, we can obtain the reconstructed body mesh M , and the reconstructed pose \mathbf{Q} via a further FK process or linear regression.

to control the full human body mesh. The shape parameters $\beta \in \mathbb{R}^{10}$ are parameterized by the first 10 principal components of the shape space. The pose parameters θ are modelled by relative 3D rotation of $K = 23$ joints, $\theta = (\theta_1, \theta_2, \dots, \theta_K)$. SMPL provides a differentiable function $\mathcal{M}(\theta, \beta)$ that takes the pose parameters θ and the shape parameters β as input and outputs a triangulated mesh $M \in \mathbb{R}^{N \times 3}$ with $N = 6980$ vertices. Conveniently, the reconstructed 3D joints \mathbf{Q}_{smpl} can be obtained by an FK process, *i.e.* $\mathbf{Q}_{smpl} = \text{FK}(\mathbf{R}, \mathbf{T})$. Also, the joints of Human3.6M [16] can be obtained by a linear combination of the mesh vertices through a linear regressor W , *i.e.* $\mathbf{Q}_{h36m} = WM$.

3.2. Hybrid Analytical-Neural Inverse Kinematics

Estimating the human body mesh by direct regression of the relative rotations is too difficult [18, 23, 20]. Here, we propose a hybrid analytical-neural inverse kinematics solution (HybrIK) to leverage 3D keypoints estimation to boost 3D body mesh estimation. Since 3D joints cannot uniquely determine the relative rotation, we decompose the original rotation into *twist* and *swing*. The 3D joints are utilized to calculate the *swing* rotation analytically, and we exploit the visual cues by a neural network to estimate the 1-DoF *twist* rotation. In HybrIK, the relative rotations are solved recursively along the kinematic tree. We conduct error analysis and further develop an adaptive solution to reduce the reconstruction error.

Twist-and-Swing Decomposition. In the analytical IK formulation, some body joints are usually assigned lower degree-of-freedom (DoFs) to simplify the problem, *e.g.* 1 or 2 DoFs [21, 62, 17]. In this work, we consider a general case where each body joint is assumed to have full 3 DoFs. As illustrated in Fig. 3, a rotation $R \in \text{SO}(3)$ can be decomposed into a *twist* rotation R^{tw} and a *swing* rotation

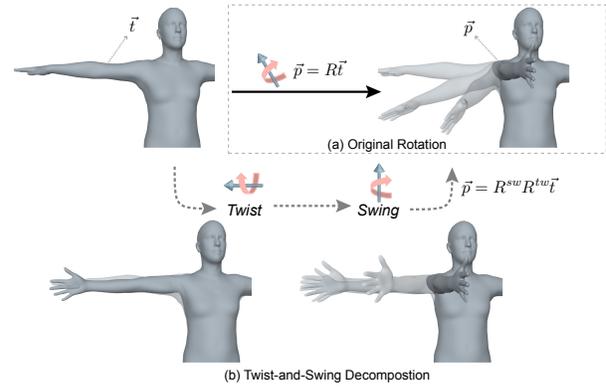


Figure 3. **Illustration of the twist-and-swing decomposition.** (a) The original rotation turns the right palm-down hand to the front and the palm to the left in one step. (b) With twist-and-swing decomposition, the rotation can be divided into two steps: First, turn the palm 90° , and then move the entire hand to the front.

R^{sw} . Given the start template body-part vector \vec{t} and the target vector \vec{p} , the solution process of R can be formulated as:

$$R = \mathcal{D}(\vec{p}, \vec{t}, \phi) = \mathcal{D}^{sw}(\vec{p}, \vec{t}) \mathcal{D}^{tw}(\vec{t}, \phi) = R^{sw} R^{tw}, \quad (6)$$

where ϕ is the *twist* angle that estimated by a neural network, $\mathcal{D}^{sw}(\cdot)$ is a closed-form solution of the *swing* rotation, and $\mathcal{D}^{tw}(\cdot)$ transforms ϕ to the *twist* rotation. Here, R should satisfy the condition in Eq. 5, *i.e.* $\vec{p} = R\vec{t}$.

- **Swing:** The *swing* rotation has the axis \vec{n} that is perpendicular to \vec{t} and \vec{p} . Therefore, it can be formulated as:

$$\vec{n} = \frac{\vec{t} \times \vec{p}}{\|\vec{t} \times \vec{p}\|}, \quad (7)$$

and the *swing* angle α satisfies:

$$\cos \alpha = \frac{\vec{t} \cdot \vec{p}}{\|\vec{t}\| \|\vec{p}\|}, \quad \sin \alpha = \frac{\|\vec{t} \times \vec{p}\|}{\|\vec{t}\| \|\vec{p}\|}. \quad (8)$$

Hence, the closed-form solution of the *swing* rotation R^{sw} can be derived by the *Rodrigues formula*:

$$R^{sw} = \mathcal{D}^{sw}(\vec{p}, \vec{t}) = \mathcal{I} + \sin \alpha [\vec{n}]_{\times} + (1 - \cos \alpha) [\vec{n}]_{\times}^2, \quad (9)$$

where $[\vec{n}]_{\times}$ is the skew symmetric matrix of \vec{n} and \mathcal{I} is the 3×3 identity matrix.

- **Twist**: The *twist* rotation is rotating around \vec{t} itself. Thus, with \vec{t} itself the axis and ϕ the angle, we can determine *twist* rotation R^{tw} :

$$R^{tw} = \mathcal{D}^{tw}(\vec{t}, \phi) = \mathcal{I} + \frac{\sin \phi}{\|\vec{t}\|} [\vec{t}]_{\times} + \frac{(1 - \cos \phi)}{\|\vec{t}\|^2} [\vec{t}]_{\times}^2, \quad (10)$$

where $[\vec{t}]_{\times}$ is the skew symmetric matrix of \vec{t} .

Note that function \mathcal{D}^{sw} and \mathcal{D}^{tw} are fully differentiable, which allows us to integrate the twist-and-swing decomposition into the training process. Although we need a neural network to learn the *twist* angle, the difficulty of learning is significantly reduced. Compared with the 3-DoF rotation that is directly regressed in previous work [18, 23, 20], the *twist* angle is only a 1-DoF variable. Moreover, due to the physical limitation of the human body, the *twist* angle has a small range of variation. Therefore, it is much easier for the networks to learn the mapping function. We further analyze its variation in §4.2.

Naive HybriK. The IK process can be performed recursively along the kinematic tree like the FK process. First of all, we need to determine the global root rotation R_0 , which has a closed-form solution using the locations of spine, left hip, right hip and Singular Value Decomposition (SVD). Detailed mathematical proof is provided in the supplemental document. Then, in each step, *e.g.* the k -th step, we assume the rotation of the parent joint $R_{pa(k)}$ is known. Hence, we can reformulate Eq. 5 with Eq. 3 as:

$$R_{pa(k)}^{-1}(p_k - p_{pa(k)}) = R_{pa(k),k}(t_k - t_{pa(k)}). \quad (11)$$

Let $\vec{p}_k = R_{pa(k)}^{-1}(p_k - p_{pa(k)})$ and $\vec{t}_k = (t_k - t_{pa(k)})$, we can solve the relative rotation via Eq. 6:

$$R_{pa(k),k} = \mathcal{D}(\vec{p}_k, \vec{t}_k, \phi_k), \quad (12)$$

where ϕ_k is the network predicting *twist* angle for the k -th joint. The set of *twist* angle is denoted as $\Phi = \{\phi_k\}_{k=1}^K$. Since the rotation matrices are orthogonal, their inverse equals to their transpose, *i.e.* $R_{pa(k)}^{-1} = R_{pa(k)}^T$, which keeps the solving process differentiable.

The whole process is named Naive HybriK and summarized in Alg. 1. Note that we solve the relative rotation $R_{pa(k),k}$ instead of the global rotation R_k . The reason is that if we directly decompose the global rotation, the resulting *twist* angle will depend on all ancestors' rotations along the kinematic tree, which increases the variation of the distal limb joints and the difficulty for the network to learn.

Algorithm 1: Naive HybriK

Input: $\mathbf{P}, \mathbf{T}, \Phi$

Output: \mathbf{R}

- 1 Determine R_0 ;
 - 2 **for** k along the kinematic tree **do**
 - 3 $\vec{p}_k \leftarrow R_{pa(k)}^{-1}(p_k - p_{pa(k)});$
 - 4 $\vec{t}_k \leftarrow (t_k - t_{pa(k)});$
 - 5 $R_{pa(k),k}^{sw} \leftarrow \mathcal{D}^{sw}(\vec{p}_k, \vec{t}_k);$
 - 6 $R_{pa(k),k}^{tw} \leftarrow \mathcal{D}^{tw}(\vec{t}_k, \phi_k);$
 - 7 $R_{pa(k),k} \leftarrow R_{pa(k),k}^{sw} R_{pa(k),k}^{tw};$
-

Adaptive HybriK. Although the Naive HybriK process seems effective, it follows an unstated hypothesis: $\|p_k - p_{pa(k)}\| = \|t_k - t_{pa(k)}\|$. Otherwise, there is no solution for Eq. 5. Unfortunately, in our case, the body-parts predicted by the 3D keypoint estimation method are not always consistent with the rest pose template. In Naive HybriK, Eq. 6 can still be solved because the condition is turned into:

$$p_k - p_{pa(k)} = R_k(t_k - t_{pa(k)}) + \vec{\epsilon}_k, \quad (13)$$

where $\vec{\epsilon}_k$ denotes the error in the k -th step, which has the same direction of $p_k - p_{pa(k)}$ and $\|\vec{\epsilon}_k\| = \|\|p_k - p_{pa(k)}\| - \|t_k - t_{pa(k)}\|\|$. To analyze the reconstruction error, we compare the difference between the input pose \mathbf{P} and the reconstructed pose \mathbf{Q} :

$$\|\mathbf{P} - \mathbf{Q}\| \Leftrightarrow \sum_{k=1}^K \|p_k - q_k\|, \quad (14)$$

where $\mathbf{Q} = \text{FK}(\mathbf{R}, \mathbf{T}) = \text{FK}(\text{IK}(\mathbf{P}, \mathbf{T}), \mathbf{T})$. Combining Eq. 2 and Eq. 13, we have:

$$\begin{aligned} p_k - q_k &= p_{pa(k)} - q_{pa(k)} + \vec{\epsilon}_k \\ &= p_{pa^2(k)} - q_{pa^2(k)} + \vec{\epsilon}_{pa(k)} + \vec{\epsilon}_k \\ &= \dots = \sum_{i \in A(k)} \vec{\epsilon}_i, \end{aligned} \quad (15)$$

where $pa^2(k)$ denotes the parent index of the $pa(k)$ -th joint, and $A(k)$ denotes the set of ancestors of the k -th joint. That means the difference between the input joint p_k and the reconstructed joint q_k will accumulate along the kinematic tree, which brings more uncertainty to the distal joint.

To address this error accumulation problem, we further propose Adaptive HybriK. In Adaptive HybriK, the target vector is adaptively updated by the newly reconstructed parent joints. Let $\vec{p}_k = R_{pa(k)}^{-1}(p_k - q_{pa(k)})$ and \vec{t}_k the same as the one in the naive solution. In this way, the condition in Adaptive HybriK can be formulated as:

$$p_k - q_{pa(k)} = R_k(t_k - t_{pa(k)}) + \vec{\epsilon}_k. \quad (16)$$

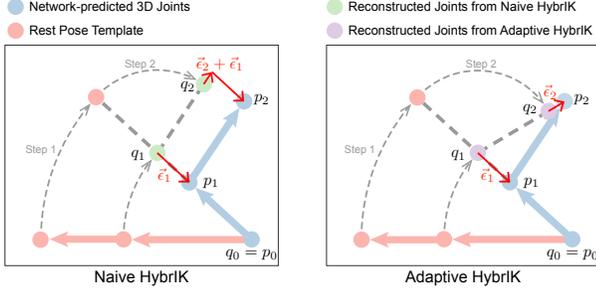


Figure 4. **Example of the reconstruction error.** The rest pose is rotated to q_1 and q_2 by two steps. In the first step, due to the bone-length inconsistency, the reconstruction error is $\vec{\epsilon}_1$. In the second step, Naive HybrIK takes $p_2 - p_1$ as the target direction, resulting in the accumulation of error $\vec{\epsilon}_1 + \vec{\epsilon}_2$. Instead, Adaptive HybrIK selects the reconstructed joint q_1 to form the target direction $p_2 - q_1$, which reduces the error to only $\vec{\epsilon}_2$.

Therefore, we have:

$$\begin{aligned} p_k - q_{\text{pa}(k)} &= q_k - q_{\text{pa}(k)} + \vec{\epsilon}_k \\ \Rightarrow p_k - q_k &= \vec{\epsilon}_k. \end{aligned} \quad (17)$$

Compared to the naive solution (Eq. 15), the reconstructed error of the adaptive solution only depends on the current joint and will not accumulate from its ancestors. As illustrated in Fig. 4, in Naive HybrIK, once the parent joint is out of position, its children will continue this mistake. Instead, in Adaptive HybrIK, the solved relative rotation is always pointing towards the target joint and reduce the error. We conduct empirical experiments in §4.2 to validate its robustness. Note that an iterative global optimization process can further reduce the error, but it is non-differentiable and does not allow end-to-end training. Adaptive HybrIK is robust enough and remains differentiable. The whole process is summarized in Alg. 2.

3.3. Learning Framework

The overall framework of our approach is illustrated in Fig. 2. Firstly, a neural network is utilized to predict 3D joints \mathbf{P} , the *twist* angle Φ and the shape parameters β . Secondly, the shape parameters are used to obtain the rest pose \mathbf{T} by the SMPL model. Then, by combining \mathbf{P} , \mathbf{T} and Φ , we perform HybrIK to solve the relative rotations \mathbf{R} of the 3D pose, *i.e.* the pose parameters θ . Finally, with the function $\mathcal{M}(\theta, \beta)$ provided by the SMPL model, the body mesh M is obtained. The reconstructed pose \mathbf{Q} can be obtained from M by FK or a regressor, which is guaranteed to be realistic. Since the HybrIK process is differentiable, the whole framework is trained in an end-to-end manner.

3D Keypoint Estimation. We adopt a simple yet effective architecture to estimate the 3D body keypoints. Following [60], we use ResNet as our backbone and 3 deconvolution layers followed by a 1×1 convolution to generate the 3D

Algorithm 2: Adaptive HybrIK

Input: $\mathbf{P}, \mathbf{T}, \Phi$
Output: \mathbf{R}

- 1 Determine R_0 ;
- 2 **for** k along the kinematic tree **do**
- 3 $q_{\text{pa}(k)} \leftarrow R_{\text{pa}(k)}(t_{\text{pa}(k)} - t_{\text{pa}^2(k)}) + q_{\text{pa}^2(k)}$;
- 4 $\vec{p}_k \leftarrow R_{\text{pa}(k)}^{-1}(p_k - q_{\text{pa}(k)})$;
- 5 $\vec{t}_k \leftarrow (t_k - t_{\text{pa}(k)})$;
- 6 $R_{\text{pa}(k),k}^{\text{sw}} \leftarrow \mathcal{D}^{\text{sw}}(\vec{p}_k, \vec{t}_k)$;
- 7 $R_{\text{pa}(k),k}^{\text{tw}} \leftarrow \mathcal{D}^{\text{tw}}(\vec{t}_k, \phi_k)$;
- 8 $R_{\text{pa}(k),k} \leftarrow R_{\text{pa}(k),k}^{\text{sw}} R_{\text{pa}(k),k}^{\text{tw}}$;

heatmaps. The soft-argmax operation is used to obtain 3D pose from the heatmap in a differentiable manner. We supervise the predicted pose coordinates with ℓ_1 loss:

$$\mathcal{L}_{\text{pose}} = \frac{1}{K} \sum_{k=1}^K \|p_k - \hat{p}_k\|_1, \quad (18)$$

where \hat{p}_k denotes the ground-truth joint.

Twist Angle Estimation. Instead of the direct regression of scalar value ϕ_k , we choose to learn a 2-dimensional vector (c_{ϕ_k}, s_{ϕ_k}) that represents $\cos \phi_k$ and $\sin \phi_k$ to avoid the discontinuity problem. The ℓ_2 loss is applied:

$$\mathcal{L}_{\text{tw}} = \frac{1}{K} \sum_{k=1}^K \|(c_{\phi_k}, s_{\phi_k}) - (\cos \hat{\phi}_k, \sin \hat{\phi}_k)\|_2, \quad (19)$$

where $\hat{\phi}_k$ denotes the ground-truth *twist* angle.

Collaboration with SMPL. The SMPL model allows us to obtain the rest pose skeleton with the additive offsets according to the shape parameters β :

$$\mathbf{T} = W(\bar{M}_{\mathbf{T}} + B_S(\beta)), \quad (20)$$

where $\bar{M}_{\mathbf{T}}$ is the mesh vertices of mean rest pose, and $B_S(\beta)$ is the blend shapes function provided by SMPL. Then the pose parameters θ are calculated by HybrIK in a differentiable manner. In the training phase, we supervise the shape parameters β :

$$\mathcal{L}_{\text{shape}} = \|\beta - \hat{\beta}\|_2, \quad (21)$$

and the rotation parameters θ :

$$\mathcal{L}_{\text{rot}} = \|\theta - \hat{\theta}\|_2. \quad (22)$$

The overall loss of the learning framework is formulated as:

$$\mathcal{L} = \mathcal{L}_{\text{pose}} + \mu_1 \mathcal{L}_{\text{shape}} + \mu_2 \mathcal{L}_{\text{rot}} + \mu_3 \mathcal{L}_{\text{tw}}, \quad (23)$$

where μ_1, μ_2 and μ_3 are weights of the loss items.

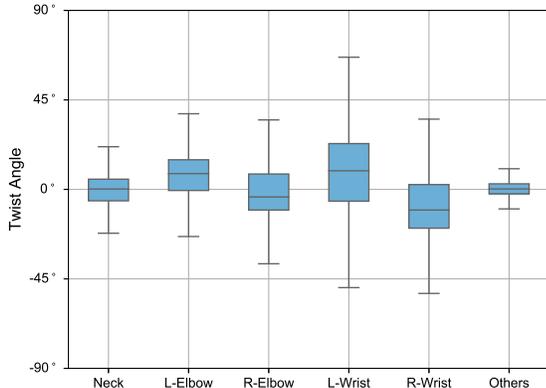


Figure 5. **Distribution of the *twist* angle.** Only a few joints have a range over 30° . Other joints have a limited range of *twist* angle.

3.4. Implementation Details

Here we elaborate more implementation details. We use ResNet-34 [15] as the network backbone, initialized with ImageNet pre-trained weights. The ResNet output is divided into two branches. The first branch is to generate 3D heatmaps. The second branch consists of an average pooling, two fully-connected layers with 1024 neurons (each with a dropout layer in between) and a final layer of 56 neurons (10 for β , 46 for Φ). The input image is resized to 256×192 . The learning rate is set to 1×10^{-3} at first and reduced by a factor of 10 at the 90th and 120th epoch. We use the Adam solver and train for 140 epochs, with a mini-batch size of 32 per GPU and 8 GPUs in total. In all experiments, $\mu_1 = 1$ and $\mu_2 = \mu_3 = 1 \times 10^{-2}$. In the testing phase, the absolute depth of the root joint is obtained from RootNet [37]. Implementation is in PyTorch.

4. Empirical Evaluation

In this section, we first describe the datasets employed for training and quantitative evaluation. Next, ablation experiments are conducted to evaluate the proposed HybrIK. Finally, we report our results and compare the proposed method with state-of-the-art approaches.

4.1. Datasets

3DPW: It is a challenging outdoor benchmark for 3D pose and shape estimation. We use this dataset only for evaluation on its test set.

MPI-INF-3DHP: It consists of both constrained indoor and complex outdoor scenes. Following [18, 23], we use its train set for training and evaluate on its test set.

Human3.6M: It is an indoor benchmark for 3D pose estimation. Following [18, 23], we use 5 subjects (S1, S5, S6, S7, S8) for training and 2 subjects (S9, S11) for evaluation.

MSCOCO: It is a large-scale in-the-wild 2D human pose datasets. We incorporate its train set for training.

	Random Twist			Estimated Twist			Zero Twist		
	24 <i>jts</i>	14 <i>jts</i>	Vert.	24 <i>jts</i>	14 <i>jts</i>	Vert.	24 <i>jts</i>	14 <i>jts</i>	Vert.
Error	0.1	40.0	67.3	0.1	6.1	10.0	0.1	6.8	12.1

Table 1. **Reconstruction error with different *twist* angle.** The accurate *twist* angles significantly reduce the reconstruction error.

4.2. Ablation Study

In this study, we evaluate the effectiveness of the twist-and-swing decomposition and the HybrIK algorithm. All evaluation is conducted on the 3DPW test set as it contains challenging in-the-wild scenes to demonstrate the strength of our model. More experimental results are provided in the supplemental document.

Analysis of the twist rotation. To demonstrate the effectiveness of twist-and-swing decomposition, we first count the distribution of the *twist* angle in the 3DPW test set. The distribution is illustrated in Fig. 5. As expected, due to the physical limitation, only neck, elbow and wrist have a wide range of variations. All other joints have a limited range of *twist* angle (less than 30°). It indicates that the *twist* angle can be reliably estimated.

Besides, we develop an experiment to see how the *twist* angles affect the reconstructed pose and shape. We take the ground-truth 24 SMPL joints and shape parameters as the input of the HybrIK process. As for the *twist* angle, we compare random values in $[-\pi, \pi]$ and the values estimated by the network. We evaluate the mean error of the reconstructed 24 SMPL joints, the 14 LSP joints, the body mesh and the *twist* angle. Here, following previous works [7, 18, 23], the 14 LSP joints are regressed from the body mesh by a pretrained regressor. Quantitative results are reported in Tab. 1. It shows that the accurate regressed *twist* angles significantly reduce the error on the mesh vertices and the LSP joints that regressed from the mesh. Since most of the twist angles are close to zeros, the zero twist angles produce acceptable performance. Notice that the wrong *twist* angles do not affect the reconstructed SMPL joints. Only the *swing* rotations change the joint locations.

Robustness of HybrIK. To demonstrate the superiority of Adaptive HybrIK over Naive HybrIK, we compare the reconstructed joints error of these two algorithms. First, we feed the ground-truth joints, *twist* angle and shape parameters to the two IK algorithms to see whether they will introduce extra error. Then we add jitters to the input to observe the performance of the HybrIK algorithm with the noisy joints. As shown in Tab. 3, when the input joints are correct, both HybrIK algorithms introduce negligible errors. For noisy joints input, the Naive HybrIK algorithm accumulates errors along the kinematic tree, while the Adaptive HybrIK algorithm is more robust to the noise.

Method	3DPW			Human3.6M		MPI-INF-3DHP		
	PA-MPJPE ↓	MPJPE ↓	PVE ↓	PA-MPJPE ↓	MPJPE ↓	PCK ↑	AUC ↑	MPJPE ↓
SMPLify [7]	-	-	-	82.3	-	-	-	-
HMR [18]	81.3	130.0	-	56.8	88.0	72.9	36.5	124.2
Kolotouros <i>et al.</i> [25]	70.2	-	-	50.1	-	-	-	-
Pavlakos <i>et al.</i> [50]	-	-	-	75.9	-	-	-	-
Arnab <i>et al.</i> [4]	72.2	-	-	54.3	77.8	-	-	-
SPIN [23]	59.2	96.9	116.4	41.1	-	76.4	37.1	105.2
Moon <i>et al.</i> [39]*	58.6	93.2	-	41.7	55.7	-	-	-
Ours (Naive HybrIK)	49.0	80.2	94.6	35.3	55.8	85.9	41.7	91.5
Ours (Adaptive HybrIK)	48.8	80.0	94.5	34.5	54.4	86.2	42.2	91.0

Table 2. **Benchmark of state-of-the-art models on 3DPW, Human3.6M and MPI-INF-3DHP datasets.** “*” denotes the method is trained on different datasets. “-” shows the results that are not available.

	GT Joints	±10 mm	±20 mm	±30 mm
Naive HybrIK	0.1	16.2	34.0	53.4
Adaptive HybrIK	0.1	9.8	20.2	31.2

Table 3. **Naive vs. Adaptive** with different input joints. MPJPE of 24 joints is reported. Adaptive HybrIK is more robust to the jitters.

Error correction capability of HybrIK. In this experiment, we examine the error correction capability of the HybrIK algorithm. The HybrIK algorithm is fed with the 3D joints, *twist* angles and shape parameters that predicted by the neural network. Additionally, we apply the SMPLify [7] algorithm on the predicted pose and compare it to our method. As shown in Tab. 4, the error of reconstructed joints after HybrIK is reduced the error to 79.2mm, while SMPLify raises the error to 114.3 mm. The error correction capability of HybrIK comes from the fact that the network may predict unrealistic body pose, *e.g.* left-right asymmetry and abnormal limbs proportions. In contrast, the rest pose is generated by the parametric statistical body model, which guarantees that the reconstructed pose is consistent with the realistic body shape distribution. Since our proposed framework is agnostic to the way we obtain 3D joints, we can improve the performance of any single-stage 3D keypoint estimation methods.

4.3. Comparison with the State-of-the-art

To make a fair comparison with previous 3D human pose and shape estimation methods, we use a regressor to obtain the 14 LSP joints from the body mesh for the evaluation on 3DPW and Human3.6M datasets and 17 joints for MPI-INF-3DHP dataset. Procrustes aligned mean per joint position error (PA-MPJPE), mean per joint position error (MPJPE), Percentage of Correct Keypoints (PCK) and Area Under Curve (AUC) are reported to evaluate the 3D pose results. We also report Per Vertex Error (PVE) to evaluate the entire estimated body mesh.

In Tab. 2, we compare our method with previous 3D human pose and shape estimation methods, including both

	Predicted Pose	HybrIK	SMPLify [7]
MPJPE (24 <i>jts</i>) ↓	88.2 mm	79.2 mm	114.3 mm

Table 4. **Error correction capability of HybrIK.** HybrIK improves the results predicted by the 3D keypoint estimation method.

model-based and model-free methods, on 3DPW, Human3.6M and MPI-INF-3DHP datasets. Without bells and whistles, our method surpasses all previous state-of-the-art methods by a large margin on all three datasets. It is worth noting that our method improve 21.9 mm PVE on 3DPW dataset, which shows that it is accurate and reliable to recover body mesh through inverse kinematics.

5. Conclusion

In this paper, we bridge the gap between 3D keypoint estimation and body mesh estimation via a novel hybrid analytical-neural inverse kinematics solution, HybrIK. It transforms the 3D joint locations to a pixel-aligned accurate human body mesh, and then obtains a more accurate and realistic 3D skeleton from the reconstructed 3D mesh, closing the loop between the 3D skeleton and the parametric body model. Our method is fully differentiable and allows simultaneously training of 3D joints and human body mesh in an end-to-end manner. We demonstrate the effectiveness of our method on various 3D pose and shape datasets. The proposed method surpasses state-of-the-art methods by a large margin. Besides, comprehensive analyses demonstrate that HybrIK is robust and has error correction capability. We hope HybrIK can serve as a solid baseline and provide a new perspective for the 3D human pose and shape estimation task.

Acknowledgements This work is supported in part by the National Key R&D Program of China, No. 2017YFA0700800, National Natural Science Foundation of China under Grants 61772332 and Shanghai Qi Zhi Institute, SHEITC (018-RGZN-02046).

References

- [1] Ijaz Akhter and Michael J Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *CVPR*, 2015. [2](#)
- [2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. In *SIGGRAPH*. 2005. [1](#), [2](#)
- [3] Andreas Aristidou and Joan Lasenby. Fabrik: A fast, iterative solver for the inverse kinematics problem. *Graphical Models*, 2011. [3](#)
- [4] Anurag Arnab, Carl Doersch, and Andrew Zisserman. Exploiting temporal context for 3d human pose estimation in the wild. In *CVPR*, 2019. [2](#), [8](#)
- [5] Paolo Baerlocher and Ronan Boulic. Parametrization and range of motion of the ball-and-socket joint. In *Deformable avatars*. Springer, 2001. [3](#)
- [6] A Balestrino, Giuseppe De Maria, and L Sciavicco. Robust control of robotic manipulators. *IFAC Proceedings Volumes*, 1984. [3](#)
- [7] Federica Bogo, Angjoo Kanazawa, Christoph Lassner, Peter Gehler, Javier Romero, and Michael J Black. Keep it smpl: Automatic estimation of 3d human pose and shape from a single image. In *ECCV*, 2016. [1](#), [2](#), [7](#), [8](#)
- [8] Ronan Boulic, Pascal Bécheiraz, Luc Emering, and Daniel Thalmann. Integration of motion control techniques for virtual human and avatar real-time animation, 1997. [1](#)
- [9] Samuel R Buss and Jin-Su Kim. Selectively damped least squares for inverse kinematics. *Journal of Graphics tools*, 2005. [3](#)
- [10] Akos Csizsar, Jan Eilers, and Alexander Verl. On solving the inverse kinematics problem using neural networks. In *M2VIP*, 2017. [3](#)
- [11] Haoshu Fang, Yuanlu Xu, Wenguan Wang, Xiaobai Liu, and Song-Chun Zhu. Learning pose grammar to encode human body configuration for 3d pose estimation. *AAAI*, 2017. [2](#)
- [12] Michael Girard and Anthony A Maciejewski. Computational modeling for the computer animation of legged figures. In *SIGGRAPH*, 1985. [3](#)
- [13] Peng Guan, Alexander Weiss, Alexandru O Balan, and Michael J Black. Estimating human shape and pose from a single image. In *ICCV*, 2009. [1](#), [2](#)
- [14] Riza Alp Guler and Iasonas Kokkinos. Holopose: Holistic 3d human reconstruction in-the-wild. In *CVPR*, 2019. [2](#)
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. [2](#), [7](#)
- [16] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3D human sensing in natural environments. *TPAMI*, 2014. [2](#), [4](#)
- [17] Marcelo Kallmann. Analytical inverse kinematics with body posture control. *Computer animation and virtual worlds*, 2008. [3](#), [4](#)
- [18] Angjoo Kanazawa, Michael J Black, David W Jacobs, and Jitendra Malik. End-to-end recovery of human shape and pose. In *CVPR*, 2018. [1](#), [2](#), [4](#), [5](#), [7](#), [8](#)
- [19] Charles A Klein and Ching-Hsiang Huang. Review of pseudoinverse control for use with kinematically redundant manipulators. *IEEE Transactions on Systems, Man, and Cybernetics*, 1983. [3](#)
- [20] Muhammed Kocabas, Nikos Athanasiou, and Michael J Black. Vibe: Video inference for human body pose and shape estimation. In *CVPR*, 2020. [1](#), [2](#), [4](#), [5](#)
- [21] Nikos Kofinas, Emmanouil Orfanoudakis, and Michail G Lagoudakis. Complete analytical inverse kinematics for nao. In *ICARSC*, 2013. [4](#)
- [22] Mia Kokic, Danica Kragic, and Jeannette Bohg. Learning to estimate pose and shape of hand-held objects from rgb images. In *IROS*, 2019. [3](#)
- [23] Nikos Kolotouros, Georgios Pavlakos, Michael J Black, and Kostas Daniilidis. Learning to reconstruct 3d human pose and shape via model-fitting in the loop. In *ICCV*, 2019. [1](#), [2](#), [4](#), [5](#), [7](#), [8](#)
- [24] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, 2019. [2](#)
- [25] Nikos Kolotouros, Georgios Pavlakos, and Kostas Daniilidis. Convolutional mesh regression for single-image human shape reconstruction. In *CVPR*, 2019. [8](#)
- [26] Christoph Lassner, Javier Romero, Martin Kiefel, Federica Bogo, Michael J Black, and Peter V Gehler. Unite the people: Closing the loop between 3d and 2d human representations. In *CVPR*, 2017. [2](#)
- [27] Yong-Lu Li, Xinpeng Liu, Han Lu, Shiyi Wang, Junqi Liu, Jiefeng Li, and Cewu Lu. Detailed 2d-3d joint representation for human-object interaction. In *CVPR*, 2020. [1](#)
- [28] Yong-Lu Li, Liang Xu, Xinpeng Liu, Xijie Huang, Yue Xu, Shiyi Wang, Hao-Shu Fang, Ze Ma, Mingyang Chen, and Cewu Lu. Pastanet: Toward human activity knowledge engine. In *CVPR*, 2020. [1](#)
- [29] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multi-person linear model. *TOG*. [1](#), [2](#), [3](#)
- [30] David G Luenberger, Yinyu Ye, et al. *Linear and nonlinear programming*. 1984. [3](#)
- [31] Diogo Carbonera Luvizon, Hedi Tabia, and David Picard. Learning features combination for human action recognition from skeleton sequences. *Pattern Recognition*, 2017. [1](#)
- [32] Julieta Martinez, Rayat Hossain, Javier Romero, and James J Little. A simple yet effective baseline for 3d human pose estimation. In *ICCV*, 2017. [2](#)
- [33] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3D human pose estimation in the wild using improved cnn supervision. In *3DV*, 2017. [2](#)
- [34] Dushyant Mehta, Helge Rhodin, Dan Casas, Pascal Fua, Oleksandr Sotnychenko, Weipeng Xu, and Christian Theobalt. Monocular 3d human pose estimation in the wild using improved cnn supervision. In *3DV*, 2017. [2](#)
- [35] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Mohamed Elgharib, Pascal Fua, Hans-Peter Seidel, Helge Rhodin, Gerard Pons-Moll, and Christian Theobalt. Xnect: Real-time multi-person 3d

- human pose estimation with a single rgb camera. *arXiv preprint arXiv:1907.00837*, 2019. 3
- [36] Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Srinath Sridhar, Gerard Pons-Moll, and Christian Theobalt. Single-shot multi-person 3d pose estimation from monocular rgb. In *3DV*, 2018. 2
- [37] Gyeongsik Moon, Ju Yong Chang, and Kyoung Mu Lee. Camera distance-aware top-down approach for 3d multi-person pose estimation from a single rgb image. In *ICCV*, 2019. 2, 7
- [38] Gyeongsik Moon and Kyoung Mu Lee. I2l-meshnet: Image-to-lixel prediction network for accurate 3d human pose and mesh estimation from a single rgb image. In *ECCV*, 2020. 2
- [39] Gyeongsik Moon and Kyoung Mu Lee. I2L-meshnet: Image-to-lixel prediction network for accurate 3d human pose and mesh estimation from a single rgb image. In *ECCV*, 2020. 8
- [40] Francesc Moreno-Noguer. 3d human pose estimation from a single image via distance matrix regression. In *CVPR*, 2017. 2
- [41] Franziska Mueller, Dushyant Mehta, Oleksandr Sotnychenko, Srinath Sridhar, Dan Casas, and Christian Theobalt. Real-time hand tracking under occlusion from an egocentric rgb-d sensor. In *ICCVW*, 2017. 3
- [42] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016. 2
- [43] Mohamed Omran, Christoph Lassner, Gerard Pons-Moll, Peter Gehler, and Bernt Schiele. Neural body fitting: Unifying deep learning and model based human pose and shape estimation. In *3DV*, 2018. 2
- [44] Bo Pang, Kaiwen Zha, Hanwen Cao, Chen Shi, and Cewu Lu. Deep rnn framework for visual sequential applications. In *CVPR*, 2019. 2
- [45] Bo Pang, Kaiwen Zha, Hanwen Cao, Jiajun Tang, Minghui Yu, and Cewu Lu. Complex sequential understanding through the awareness of spatial and temporal concepts. *Nature Machine Intelligence*, 2020. 2
- [46] Sunghoon Park, Jihye Hwang, and Nojun Kwak. 3d human pose estimation using convolutional neural networks with 2d pose information. In *ECCV*, 2016. 2
- [47] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed AA Osman, Dimitrios Tzionas, and Michael J Black. Expressive body capture: 3d hands, face, and body from a single image. In *CVPR*, 2019. 1, 2
- [48] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3D human pose. In *CVPR*, 2017. 1
- [49] Georgios Pavlakos, Xiaowei Zhou, Konstantinos G Derpanis, and Kostas Daniilidis. Coarse-to-fine volumetric prediction for single-image 3d human pose. In *CVPR*, 2017. 2
- [50] Georgios Pavlakos, Luyang Zhu, Xiaowei Zhou, and Kostas Daniilidis. Learning to estimate 3d human pose and shape from a single color image. In *CVPR*, 2018. 2, 8
- [51] Dario Pavlo, David Grangier, and Michael Auli. Quaternet: A quaternion-based recurrent model for human motion. In *BMVC*, 2018. 3
- [52] Liliana Lo Presti and Marco La Cascia. 3d skeleton-based human action classification: A survey. *Pattern Recognition*, 2016. 1
- [53] Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Reconstructing 3d human pose from 2d image landmarks. In *ECCV*, 2012. 2
- [54] Gregory Rogez, Philippe Weinzaepfel, and Cordelia Schmid. Lcr-net: Localization-classification-regression for human pose. In *CVPR*, 2017. 2
- [55] Nizar Rokbani, Alicia Casals, and Adel M Alimi. Ik-fa, a new heuristic inverse kinematics solver using firefly algorithm. In *Computational Intelligence Applications in Modeling and Control*. 2015. 3
- [56] Marta Sanzari, Valsamis Ntouskos, and Fiora Pirri. Bayesian image based 3d pose estimation. In *ECCV*, 2016. 2
- [57] Leonid Sigal, Alexandru Balan, and Michael J Black. Combined discriminative and generative articulated pose and non-rigid shape estimation. In *NeurIPS*, 2008. 2
- [58] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014. 2
- [59] Xiao Sun, Jiayang Shang, Shuang Liang, and Yichen Wei. Compositional human pose regression. In *ICCV*, 2017. 2
- [60] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *ECCV*, 2018. 1, 6
- [61] Xiao Sun, Bin Xiao, Fangyin Wei, Shuang Liang, and Yichen Wei. Integral human pose regression. In *ECCV*, 2018. 2
- [62] Deepak Tolani, Ambarish Goswami, and Norman I Badler. Real-time inverse kinematics techniques for anthropomorphic limbs. *Graphical models*, 2000. 3, 4
- [63] Hsiao-Yu Tung, Hsiao-Wei Tung, Ersin Yumer, and Katerina Fragkiadaki. Self-supervised learning of motion capture. In *NeurIPS*, 2017. 2
- [64] Hsiao-Yu Fish Tung, Adam W Harley, William Seto, and Katerina Fragkiadaki. Adversarial inverse graphics networks: Learning 2d-to-3d lifting and image-to-image translation from unpaired supervision. In *ICCV*, 2017. 2
- [65] Gul Varol, Duygu Ceylan, Bryan Russell, Jimei Yang, Ersin Yumer, Ivan Laptev, and Cordelia Schmid. Bodynet: Volumetric inference of 3d human body shapes. In *ECCV*, 2018. 2
- [66] Ruben Villegas, Jimei Yang, Duygu Ceylan, and Honglak Lee. Neural kinematic networks for unsupervised motion retargetting. In *CVPR*, 2018. 3
- [67] Timo von Marcard, Roberto Henschel, Michael Black, Bodo Rosenhahn, and Gerard Pons-Moll. Recovering accurate 3d human pose in the wild using imus and a moving camera. In *ECCV*, 2018. 2
- [68] Charles W Wampler. Manipulator inverse kinematic solutions based on vector formulations and damped least-squares methods. *IEEE Transactions on Systems, Man, and Cybernetics*, 1986. 3
- [69] Can Wang, Jiefeng Li, Wentao Liu, Chen Qian, and Cewu Lu. Hmor: Hierarchical multi-person ordinal relations for monocular multi-person 3d pose estimation. In *ECCV*, 2020. 2

- [70] William A Wolovich and H Elliott. A computational technique for inverse kinematics. In *CDC*, 1984. 3
- [71] Hashim Yasin, Umar Iqbal, Bjorn Kruger, Andreas Weber, and Juergen Gall. A dual-source approach for 3d pose estimation from a single image. In *CVPR*, 2016. 2
- [72] Yusuke Yoshiyasu, Ryusuke Sagawa, Ko Ayusawa, and Akihiko Murai. Skeleton transformer networks: 3d human pose and skinned mesh from single rgb image. In *ACCV*, 2018. 3
- [73] Andrei Zanfir, Elisabeta Marinoiu, and Cristian Sminchisescu. Monocular 3d pose and shape estimation of multiple people in natural scenes-the importance of multiple scene constraints. In *CVPR*, 2018. 2
- [74] Xingyi Zhou, Qixing Huang, Xiao Sun, Xiangyang Xue, and Yichen Wei. Towards 3d human pose estimation in the wild: a weakly-supervised approach. In *ICCV*, 2017. 2
- [75] Xingyi Zhou, Xiao Sun, Wei Zhang, Shuang Liang, and Yichen Wei. Deep kinematic pose regression. In *ECCV*, 2016. 3
- [76] Xiaowei Zhou, Menglong Zhu, Spyridon Leonardos, and Kostas Daniilidis. Sparse representation for 3d shape estimation: A convex relaxation approach. *TPAMI*, 2016. 2
- [77] Xiaowei Zhou, Menglong Zhu, Spyridon Leonardos, Konstantinos G Derpanis, and Kostas Daniilidis. Sparseness meets deepness: 3d human pose estimation from monocular video. In *CVPR*, 2016. 2