

This WACV 2023 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

Learning 3D Human Pose Estimation from Dozens of Datasets using a Geometry-Aware Autoencoder to Bridge Between Skeleton Formats

István Sárándi Alexander Hermans Bastian Leibe RWTH Aachen University, Germany

{sarandi,hermans,leibe}@vision.rwth-aachen.de

Abstract

Deep learning-based 3D human pose estimation performs best when trained on large amounts of labeled data, making combined learning from many datasets an important research direction. One obstacle to this endeavor are the different skeleton formats provided by different datasets, i.e., they do not label the same set of anatomical landmarks. There is little prior research on how to best supervise one model with such discrepant labels. We show that simply using separate output heads for different skeletons results in inconsistent depth estimates and insufficient information sharing across skeletons. As a remedy, we propose a novel affine-combining autoencoder (ACAE) method to perform dimensionality reduction on the number of landmarks. The discovered latent 3D points capture the redundancy among skeletons, enabling enhanced information sharing when used for consistency regularization. Our approach scales to an extreme multi-dataset regime, where we use 28 3D human pose datasets to supervise one model, which outperforms prior work on a range of benchmarks, including the challenging 3D Poses in the Wild (3DPW) dataset. Our code and models are available for research purposes.¹

1. Introduction

Research on 3D human pose estimation has gone through enormous progress in recent years [12, 17, 37, 39, 43, 46, 57, 69]. While semi-supervised and self-supervised approaches are on the rise [41, 79], best results are still achieved when using as much labeled training data as possible. However, individual 3D pose datasets tend to be rather small and lacking in diversity, as they are often recorded in a single studio with few subjects. Therefore, to provide the best possible models for downstream applications (*e.g.*, action recognition, sports analysis, medical rehabilitation, collaborative robotics), it becomes important to use many datasets in the training process. Thanks to sustained efforts by the research community, numerous publicly re-



Figure 1: Different 3D human pose datasets (*e.g.*, CMU-Panoptic and Human3.6M) provide annotations for different sets of body landmarks (*left*). To best leverage such discrepant labels for multi-dataset 3D pose estimation, we discover a smaller set of latent 3D keypoints (*right*), from which the dataset-specific points can be reconstructed. This allows us to capture the redundancy among the different skeleton formats and enhance information sharing between datasets, ultimately leading to improved pose accuracy.

leased, labeled datasets exist. However, as prior published works only train on at most a handful of them, it remains unknown what performance could be achieved by combining more than a decade of dataset collection efforts into a single model. Unfortunately, this is not a trivial undertaking, since different datasets do not use the same skeleton format for their labels (see Fig. 1), *e.g.*, the hip keypoints are at different heights, some body parts are only labeled in some datasets, some provide surface markers while others provide keypoints inside the body, etc. Prior work has rectified such

¹https://vision.rwth-aachen.de/wacv23sarandi



Figure 2: We train models to jointly estimate 3D human pose according to multiple different skeleton formats so that we can train on many datasets at once. a) Simply using separate prediction heads on a shared backbone is an insufficient solution to this multi-skeleton learning problem, as we obtain inconsistent outputs along the depth axis. b) We propose a method to capture and exploit the redundancy among the different skeletons using a novel affine-combining autoencoder-based (ACAE) regularization. This leads to a clear improvement in skeleton consistency.

differences through a handful of individually defined rules (*e.g.*, shrink the hip-pelvis distance by a certain factor [65]), but this does not scale to many keypoints and datasets – we need a more systematic and automatic method. The question we tackle in this work is therefore: *How can we automatically merge dozens of 3D pose datasets into one training process, given the label discrepancies?* We refer to this task as *multi-skeleton 3D human pose estimation.*²

If we ignored the discrepancies altogether and proceeded as if keypoints with the same name represented the same body landmark, the model would be supervised with inconsistently labeled examples and would learn to output a skeleton format that is some kind of average of the true ones, leading to subpar benchmark performance. Alternatively, we may consider this as a multi-task learning problem, and predict the skeletons on separate output heads on a shared backbone, without assuming any skeleton correspondences. But as we will see, this is not ideal either, as there is insufficient information sharing between skeletons, which is most apparent in inconsistencies between the depth predictions of such a model, as shown in Fig. 2a.

To strike the right balance between those two extremes, we aim to establish *some* connections between the skeleton formats without assuming them to be the same. To learn such geometric relations between skeletons, we introduce a novel autoencoder-based dimensionality reduction technique to compress a larger set of 3D keypoints (the joints from all datasets) into a lower-cardinality representation (a smaller latent keypoint set). The encoder and decoder compute affine combinations of their input points, and are thus equivariant to rotation and translation. We further induce chirality equivariance (left-right symmetry) via weight sharing [77]. We call this model an Affine-Combining Autoencoder (ACAE). We employ the ACAE in pose estimation training as an output regularizer, to encourage consistent predictions. This improves prediction results both qualitatively and quantitatively. As an alternative to the regularization approach, we can also directly predict the latent keypoints of the ACAE with a 3D pose estimator. This latter variant avoids the need for the underlying pose estimator to estimate a large number of joints, which may be costly for some methods. In both cases, the final predictions become consistent, showing the value of our approach in tackling multi-dataset 3D pose estimation.

Through an extensive literature review, we have identified 28 datasets with high-quality 3D human pose labels. By systematically pre-processing these datasets and discarding redundant poses, we constructed a meta-dataset of 13 million examples, spanning more than a thousand people. This is almost two orders of magnitude more data than in typical research papers (*e.g.*, Human3.6M has ~165k examples after redundancy filtering). We show that using more data indeed helps, and that our approach scales to 28 datasets providing a total of 555 joints in their skeleton formats, summarized in Tab. 1. Our final models show excellent in-thewild performance, outperforming currently available models, making them highly useful for downstream research.

In summary, we make the following **contributions** in this paper. (1) We assemble the largest scale meta-dataset for 3D human pose estimation to date, consisting of 28 individual datasets, and release scripts for reproducing the process. We call special attention to the problem of dis-

²For simplicity, we call any set of landmarks provided in a particular dataset a "skeleton", and use "landmark" and "joint" synonymously.

Table 1: We study the extreme multi-dataset setting of 3D human pose estimation, using all datasets below in one training process. We define three dataset combinations (indicated by \diamond , \blacklozenge , and, \heartsuit) to study the effect of training data amount. (* marks custom dataset-specific skeletons.)

	ample	ب م	ould'	points et	
Dataset name	#EX.c.	*Reice	*ter	· 51105-	Skeleton
Real imag	ges with i	marker	less N	1oCap	
MuCo-3DHP [56]	677k	8	28	♦♣♡	3DHP
CMU-Panoptic [36]	2.81M	>60	19	\$♡	COCO
AIST-Dance++ [42, 72]	1.86M	30	19	\$♡	COCO
HUMBI [78]	1.26M	772	19	♦♡	COCO
MPI-INF-3DHP [55]	627k	8	28	\$♡	3DHP
RICH [29]	96k	15	42	¢¢	SMPL-X [62]
BEHAVE [6]	42k	7	43	\heartsuit	SMPL, COCO
ASPset [58]	124k	15	17	\heartsuit	*
3DOH50K [80]	50k	<10	14	\heartsuit	LSP [35]
IKEA ASM [5]	23k	48	17	\heartsuit	*
Real image	es with m	arker-l	pased	МоСар	
Human3.6M [9, 30]	165k	5	25	♦♣♡	H3.6M
TotalCapture [71]	130k	5	21	¢	*
BML-MoVi [21]	553k	13	87	\heartsuit	*
Berkeley-MHAD [60]	526k	12	43	\heartsuit	*
UMPM [1]	164k	30	15	\heartsuit	*
Fit3D [20]	147k	8	25	\heartsuit	H3.6M
GPA [75]	109k	13	34	\heartsuit	*
HumanSC3D [19]	72k	4	25	\heartsuit	H3.6M
CHI3D [18]	46k	6	25	\heartsuit	H3.6M
Human4D [10]	40k	4	32	\heartsuit	*
MADS [81]	33k	5	15	\heartsuit	*
	Syntheti	c image	es		
SURREAL [73]	1.9M		24	♦♣♡	SMPL [48]
3DPeople [64]	946k		29		*
JTA [16]	562k		22		*
HSPACE [4]	195k		35		GHUM [76]
SAIL-VOS [28]	101k		26	¢	*
AGORA [61]	79k		66		SMPL[-X]
SPEC [40]	59k		24	\heartsuit	SMPL
Real images with	2D anno	otations	(wea	k superv	vision)
COCO [45]	47k		17		
MPII [2]	27k		16		
PoseTrack [3]	40k		15		
JRDB [54]	59k		17		
Total	s (for 3D	-labele	d dat	a)	
Small (3 datasets)	2 8M	13	77	, 	
Medium (14 datasets)	10.8M	>900	277	•	
GRAND TOTAL (28 ds.)	13.4M	>1k	555	÷ v	
5.1.1.1.D 101112 (20 d3.)	10.401	/ IK	555	*	

parate skeleton annotation formats in these datasets, which has rarely been addressed in the literature so far. (3) We propose affine-combining autoencoders (ACAE), a novel linear dimensionality reduction technique applicable to keypointbased representations such as poses. (4) We apply the ACAE to regularize model predictions to become more consistent, leading to qualitative and quantitative improvements, and we show that the latent points can be predicted directly as well. (5) We release high-quality 3D pose estimation models with excellent and consistent in-the-wild performance due to diverse supervision and our regularization tying together different skeleton formats.

2. Related Work

3D Human Pose Estimation. For an overview on the current trends in 3D pose estimator design, we refer the reader to excellent current surveys [11, 15, 34, 47, 79]. We emphasize that our approach is independent of the internals of the pose estimation method.

Handling Discrepancy in Skeleton Formats. In 2D-to-3D pose lifting, Rapczyński et al. [65] combine pairs of datasets in training by concatenating the training data sets, and harmonize the joints through hand-crafted rules. In 2D pose tracking, Guo et al. [24] train dataset-specific output heads and combine their results via hand-crafted rules. To unify pose representations, some prior works on (imageindependent) MoCap data standardized the height and bone length of skeletons [26, 52]. The AMASS dataset [51] addresses the problem of discrepancy in MoCap data representations by mapping them to the SMPL [48] representation, but the dataset does not provide corresponding images and cannot be used for image-based pose estimation. Furthermore, the underlying MoSH++ algorithm relies on a complex, multi-stage procedure requiring temporal sequence data and a pre-existing body mesh model. In contrast, our method has different goals and is much simpler in comparison. We do not aim to generate a definitive, universal ground truth representation for all datasets, instead our latent keypoint set is only used as an intermediate representation for the pose estimator, but the losses and evaluations are still computed in the original skeleton formats, after decoding the latent points into full skeletons.

Keypoint Discovery. Discovering a good set of landmarks to describe objects has been investigated in other contexts in computer vision. 2D keypoint discovery has been used to disentangle pose and shape in 2D human pose estimation [31, 32]. In 3D, Jakab *et al.* [33] discover control points for deforming 3D shapes. Rhodin *et al.* [66] learn a 3D human representation that consists of a set of 3D points, which encode both pose and appearance, optimizing for the unsupervised auxiliary task of novel view synthesis. Loper *et al.* [49] optimize the placement of sparse markers on the body to best capture both human shape and pose.

Linear Subspace Learning. Linear dimensionality reduction has a long history, with principal component analysis being the best known representative [63]. Its relation to autoencoders was discovered by Bourlard and Kamp [8], and a recent paper by the same first author reviews the developments since [7]. Linear autoencoders have been employed in robust and sparse [23] variants, a detailed overview is presented in [14]. Our proposed affine-combining autoencoders are related, but have different constraints, tailored to our use case, *i.e.*, that the weights sum to unity, and there is no requirement of orthogonality, unlike in PCA.



Figure 3: Our complete training workflow. We train an initial model on all skeletons of multiple datasets without enforcing consistency. Using this model, we create pseudo-ground truth, needed to train an autoencoder that learns a latent keypoint space. In turn, we use this frozen autoencoder to regularize the initial model during fine-tuning, encouraging consistent predictions.

3. Method

Our goal is to obtain a strong, monocular RGB-based 3D human pose estimation model by integrating numerous datasets into one mixed training process, even when the different datasets provide annotations according to different skeleton formats. Suppose we have D skeleton formats, with $\{J_d\}_{d=1}^D$ joints in each, for a total of $J = \sum_{d=1}^D J_d$ joints overall. Further, we have a merged dataset with Ntraining examples, each consisting of an image of a person and annotations for a subset of the J body joints in 3D. Our proposed workflow consists of three main steps. First, we train an initial model that predicts the different skeletons on separate prediction heads, branching out from a common backbone network (Fig. 3a). With the resulting model, we can run inference and produce a pseudo-ground truth "parallel corpus" of many poses given in every skeleton format. From this, the geometric relations between skeleton formats can be captured. We accomplish this in the second step, by training an undercomplete geometry-aware autoencoder, which discovers a latent 3D body landmark set that best captures human pose variations in the pseudo-GT data (Fig. 3b). Finally, equipped with the trained autoencoder, we rely on its learned latent space to make the model output consistent across skeleton formats through output regularization (Fig. 3c). We also experiment with direct latent point prediction, and a hybrid variant for the last step.

3.1. Initial Model Training

The first step of our workflow is to train an initial pose estimator to predict all *J* joints separately (Fig. 3a). This means that no correspondences or relations across different skeletons are assumed, *i.e.*, without specifying or enforcing that the left shoulder joint of one skeleton should be predicted near the left shoulder of another skeleton. This is akin to multi-task architectures that use different taskspecific heads on one backbone. The pose loss we minimize is $\mathcal{L}_{pose} = \mathcal{L}_{meanrel} + \lambda_{proj} \mathcal{L}_{proj} + \lambda_{abs} \mathcal{L}_{abs}$, where $\mathcal{L}_{meanrel}$ is an ℓ_1 loss computed after aligning the prediction and ground truth at the mean, \mathcal{L}_{proj} is an ℓ_1 loss on the 2D coordinates after projection onto the image, and \mathcal{L}_{abs} is an ℓ_1 loss on the absolute pose (in camera coordinates). Since each training example is annotated only with a subset of the *J* joints, we ignore any unlabeled joints when averaging the loss.

When visualizing the different skeleton outputs of this trained model, we see inconsistencies among them along the challenging depth axis (see Fig. 2). This is understandable, since we have not employed any training mechanism that would ensure any relations between the output skeletons (except that they are predicted from shared backbone features). On the other hand, when projected onto the image plane, the predictions appear sufficiently consistent.

3.2. Pseudo-Ground Truth Generation

To characterize how the joints of the different skeletons relate to one another, we need pose labels according to all skeleton formats for the same examples, to function as a "Rosetta Stone". Since no such ground truth is available (datasets only provide one type of skeletons, rarely two), we generate pseudo-ground truth using the initial separate-head model. It is important to use images that the model can handle well in this step, hence we choose a relatively clean, clutter-free subset of the *training* data for this purpose (H36M and MoVi). This yields a set of K pseudo-ground truth poses, with all J joints: $\{P_k \in \mathbb{R}^{J \times 3}\}_{k=1}^K$.

3.3. Affine-Combining Autoencoder

To capture the redundancy among the full set of J joints, and ultimately to improve the consistency in estimating them, we introduce a simple but effective dimensionality reduction technique. Since the pseudo-GT is more reliable in 2D (the X and Y axes) than in the depth dimension, the transformation to and from the latent representation should be viewpoint-independent, in other words the representation should be equivariant to rotation and translation. This equivariance in turn requires the latent representation to be geometric, *i.e.* to consist of a list of L latent 3D points $Q_k \in \mathbb{R}^{L \times 3}$ (L < J). This makes intuitive sense: the way the different skeletons relate to each other is only dependent on how joints are defined on the human body, not on the camera angle. The latent points are then responsible for spanning the overall structure of a pose. Specific skeleton formats can then be computed in relation to these latents. Further, the latent points should only have sparse influence on the joints, *e.g.*, some latent points should be responsible for the positioning of the left arm and these should have no influence on the right leg's pose. We find that these requirements can be fulfilled effectively by adopting a novel constrained undercomplete linear autoencoder structure, which we call *affine-combining autoencoder* (ACAE). Instead of operating on general *n*-dimensional vectors, an ACAE's **encoder** takes as input a list of *J* points $\mathbf{p}_j \in \mathbb{R}^3$ and encodes them into *L* latent points $\mathbf{q}_l \in \mathbb{R}^3$ by computing affine combinations according to

$$\mathbf{q}_{l} = \sum_{j=1}^{J} w_{l,j}^{\text{enc}} \mathbf{p}_{j}, \quad \sum_{j=1}^{J} w_{l,j}^{\text{enc}} = 1, \quad \forall l = 1, \dots, L.$$
(1)

Similarly, the **decoder**'s goal is to reproduce the original points from the latents, again through affine combinations:

$$\hat{\mathbf{p}}_{j} = \sum_{l=1}^{L} w_{j,l}^{\text{dec}} \mathbf{q}_{l}, \quad \sum_{l=1}^{L} w_{j,l}^{\text{dec}} = 1, \quad \forall j = 1, \dots, J.$$
(2)

Since affine combinations are equivariant to any affine transformation, our encoder and decoder are guaranteed to be rotation and translation equivariant. (Note that the same weighting is used for the X, Y and Z coordinates.)

The learnable parameters of the ACAE are the affine combination weights $w_{l,j}^{enc}$ and $w_{j,l}^{dec}$, which can also be understood as (potentially negative) generalized barycentric coordinates [27] for the latents w.r.t. the full joint set and vice versa. Allowing negative coordinates is necessary, as this allows the latents to spread outwards from the body, similar to a cage used in graphics [59]. Restricting the encoder and decoder to convex combinations would severely limit its expressiveness. To achieve sparsity in the weights (*i.e.*, spatially localized influence), we use ℓ_1 regularization, and this also reduces the amount of negative weights, preferring nearly convex combinations. We further adopt the ℓ_1 reconstruction loss, as it is robust to outliers which may be present due to noise in the pseudo-GT.

Problem Statement. We can now formally state our proposed ACAE problem in matrix notation for the weights. Given *K* training poses with *J* joints $\{P_k \in \mathbb{R}^{J \times 3}\}_{k=1}^{K}$,

$$\begin{aligned} \underset{W_{\text{enc}} \in \mathbb{R}^{L \times J}, W_{\text{dec}} \in \mathbb{R}^{J \times L}}{\text{minimize}} \mathcal{L}_{\text{reconstr}} + \lambda_{\text{sparse}} \mathcal{L}_{\text{sparse}} \\ \mathcal{L}_{\text{reconstr}} &= \frac{1}{K} \sum_{k=1}^{K} \|P_k - W_{\text{dec}} W_{\text{enc}} P_k\|_1 \quad (3) \\ \mathcal{L}_{\text{sparse}} &= \|W_{\text{enc}}\|_1 + \|W_{\text{dec}}\|_1 \\ \text{s. t. } W_{\text{enc}} \mathbf{1}_J = \mathbf{1}_L, \quad W_{\text{dec}} \mathbf{1}_L = \mathbf{1}_J, \end{aligned}$$

where $\mathbf{1}_a$ is a vector of dimension a filled with ones and

 λ_{sparse} controls the strength of the sparsity regularization. The sum-to-one (partitioning of unity) constraints ensure that the weights express affine combinations, which is necessary for translation equivariance.

Reconstruction Loss on 2D Projection. As discussed above, the pseudo-GT is more reliable in its 2D projection than along the depth axis. We therefore adapt the above general problem formulation to take this into account by defining the reconstruction loss on 2D projections:

$$\mathcal{L}_{\text{reconstr}}^{\text{proj}} = \frac{1}{K} \sum_{k=1}^{K} \|\Pi(P_k) - \Pi(W_{\text{dec}} W_{\text{enc}} P_k)\|_1, \quad (4)$$

where $\Pi(\cdot)$ denotes camera projection.

Our key insight here is that it is sufficient to observe the high-quality 2D image-plane projections of this model's outputs to characterize how the joints of different skeleton formats geometrically interrelate, because these relations are viewpoint-independent. As a simplified example, if we observe on many poses, that a certain joint tends to be halfway in between two other joints in 2D, then this will also have to hold along the depth axis.

Chirality Equivariance. As humans have bilateral symmetry, it is natural to expect the autoencoder to be chiralityequivariant, *i.e.* to process the left and right sides the same way [77]. To this end, we partition the latent keypoints into three disjoint sets: left, right and central latents, following the same proportions as in the full joint set. Assume, without loss of generality, that the points are sorted and grouped into left-side, right-side and center blocks. We then impose the following weight-sharing block structure on both the encoder and decoder weight matrices:

$$W = \begin{bmatrix} W_1 & W_2 & W_3 \\ W_2 & W_1 & W_3 \\ W_4 & W_4 & W_5 \end{bmatrix}.$$
 (5)

This structure indeed ensures chirality equivariance, since the matrix remains the same if we permute both its rows and columns by swapping the first two sections, *i.e.*, swapping the left and right points in the inputs and the outputs.

Head Keypoint Weighting. Based on the intuition that smaller motions of head and facial keypoints can be more semantically relevant, we weight these joints higher (by a factor of 10) in the loss, ensuring that the latents sufficiently cover the head as well. (We later found that this is not strictly necessary and the method also works without this as well.)

Training. We train the autoencoder using the Adam optimizer [38] with batch size 32. To enforce the sum-to-one constraints, we normalize the weight matrices within the computational graph.



(b) Hybrid of Fig. 3c and 4a with a further student-teacher loss.

Figure 4: Alternative model structures for the fine-tuning phase, to be used instead of Fig. 3c in our training workflow.

3.4. Consistency Fine-Tuning

Once our affine-combining autoencoder is trained on pseudo-ground truth, we freeze its weights and use it to enhance the consistency of 3D pose estimation outputs, with one of three alternative methods.

Output Regularization. In this case (Fig. 3c), we estimate all J joints $\hat{P} \in \mathbb{R}^{J \times 3}$ with the underlying pose estimator, but we feed this output through the autoencoder, and apply an additional loss term that measures the consistency of the prediction with the latent space, through an ℓ_1 loss, as

$$\mathcal{L}_{\text{cons}} = \|\hat{P} - W_{\text{dec}} W_{\text{enc}} \hat{P}\|_1.$$
(6)

This encourages that the separately predicted skeletons can be projected to latent keypoints and back without information loss, thereby discouraging inconsistencies between them. The pose loss \mathcal{L}_{pose} (*c.f.* Sec. 3.1) applied on \hat{P} .

Direct Latent Prediction. To avoid having to predict a large number of J joints in the base pose estimator, we define an alternative approach where the latents $\hat{Q} \in \mathbb{R}^{L \times 3}$ are directly predicted and then fed to the frozen decoder (Fig. 4a). The last layer is reinitialized from scratch, as the number of predicted joints changes from J to L. The pose loss $\mathcal{L}_{\text{pose}}$ is applied on $W_{\text{dec}}\hat{Q}$.

Hybrid Student-Teacher. In a hybrid of the above two variants, we keep the full prediction head and add a newly initialized one to predict the latents \hat{Q} directly (Fig. 4b). To distill the knowledge of the full prediction head to the latent head, we add a student-teacher-like ℓ_1 loss

$$\mathcal{L}_{\text{teach}} = \|\hat{Q} - \text{stop}_{\text{gradient}}(W_{\text{enc}}\hat{P})\|_{1}, \quad (7)$$

which is only backpropagated to the latent predictor (the student). During inference, we use $W_{\text{dec}}\hat{Q}$ as the output, to be as lightweight as direct latent prediction.

4. Experimental Setup

Base Model. We adopt the recent state-of-the-art Me-TRAbs [67] 3D human pose estimator as the platform for our experiments, but we note that our method is agnostic to the specifics of the underlying pose estimator. Unless mentioned otherwise, the backbone is EfficientNetV2-S [70].

Training Details. We perform 400k training steps with AdamW [50] and batch size 128, with every dataset represented with a fixed number of examples per batch. Batch composition and learning rate schedule are specified in the supplementary. We use ghost BatchNorm [25, 68] with size 16, as this improved convergence in multi-dataset training, together with switching the BatchNorm layers to inference mode for the last 1000 updates. The final fine-tuning phase has 40k iterations with smaller learning rate on the backbone than the heads. The autoencoder weights are trained on pseudo-GT obtained with EffV2-L. We crop a 256×256 px square around the person, apply perspective undistortion with camera intrinsics and perform augmentation as in [67].

Datasets. See Tab. 1 for an overview of all used datasets, which employ a variety of skeleton formats. In some cases, e.g., when annotations are derived through triangulating COCO-like predictions (of *e.g.*, OpenPose), or through fitting a body model (e.g., SMPL), we can assume that multiple datasets use the same convention (indicated in the last column). For other datasets, we assume the skeleton is a custom one, yielding 555 distinct keypoints in total. As most 3D human datasets contain videos, rather than isolated images, the number of sufficiently different poses is smaller than the total number of annotated frames. We hence discard examples where all joints remain within 100 mm of the last stored example. Our overall processing ensures that each training example has a person-centered image crop, camera intrinsics, 3D coordinates for some subset of the joints, a bounding box and a segmentation mask.

Evaluation Metrics. We evaluate on four datasets: MuPoTS [56], 3DPW [53], 3DHP [55] and H36M [9, 30]. Over the years, different evaluation metrics and protocols have become customary on different datasets, whose details can be very arcane. Especially in a multi-dataset setting, we find it important to use consistent metrics. For our main experiments, we therefore adopt the following four metrics everywhere: MPJPE: mean Euclidean distance between predicted and ground truth joints after alignment at the root joint. PMPJPE: mean Euclidean distance after Procrustes alignment. PCK@100mm: percentage of joints predicted within 100 mm of the ground truth after root alignment. CPS@200mm: percentage of poses where all joints are within 200 mm distance of the ground truth after root alignment [74]. More details on the experimental setup are in the supplementary.

Table 2: Results using different amounts of datasets when training a separate-head model. Tab. 1 defines which datasets belong in which combination size. Using more datasets improves results on the 3DPW, 3DHP and MuPoTS benchmarks. On Human3.6M the small dataset combination gives better results, but this studio benchmark it less suited for studying real-world generalization capacity, as opposed to in-the-wild and outdoor benchmarks such as 3DPW and MuPoTS.

	MuPoTS-3D			3DPW			MPI-INF-3DHP				Human3.6M					
	MPJPE↓	PMPJPE↓	$PCK_{100}↑$	CPS_{200} †	MPJPE↓	PMPJPE↓	PCK ₁₀₀ ↑	CPS_{200} †	MPJPE↓	PMPJPE↓	$PCK_{100}↑$	CPS_{200}^{\uparrow}	MPJPE↓	PMPJPE↓	$PCK_{100}↑$	$\text{CPS}_{200} \uparrow$
Single dataset	91.3	62.9	65.3	53.5	-	_	_	-	66.5	46.7	83.0	78.4	48.3	33.2	92.1	89.8
Small (\$)	88.4	61.3	67.3	59.9	81.0	54.5	72.9	35.8	64.8	45.9	83.7	78.9	42.1	33.8	94.6	90.3
Medium (+)	86.1	59.4	69.0	67.9	64.3	45.6	82.5	70.0	61.7	44.6	85.6	80.2	43.2	34.5	94.3	90.3
Full (♡)	84.6	59.0	70.1	66.0	61.8	43.4	83.8	71.1	59.6	44.1	86.6	81.8	44.7	34.3	94.3	90.1

Table 3: Evaluation of different strategies for handling different skeleton annotation formats during training.

		MuPoTS-3D				3DPW			MPI-INF-3DHP				Human3.6M				
		MPJPE↓	PMPJPE↓	PCK ₁₀₀ ↑	$\mathrm{CPS}_{200}^{}\uparrow$	MPJPE↓	PMPJPE↓	PCK_{100}^{\dagger}	CPS_{200}^{\uparrow}	MPJPE↓	PMPJPE↓	PCK ₁₀₀ ↑	$\text{CPS}_{200}^{\uparrow}$	MPJPE↓	PMPJPE↓	PCK ₁₀₀ ↑	$\text{CPS}_{200}^{}\uparrow$
	Merged joints	91.9	67.3	63.2	69.9	72.5	48.3	79.5	69.7	69.8	51.6	80.7	79.4	44.6	34.2	93.9	89.8
2-S	Separate joints (F. 3a)	84.6	59.0	70.1	66.0	61.8	43.4	83.8	71.1	59.6	44.1	86.6	81.8	44.7	34.3	94.3	90.1
etV	Consistency regul. (F. 3c)	81.8	57.8	72.5	72.9	61.5	43.0	84.0	71.9	59.2	43.6	86.6	82.7	45.2	33.3	94.4	90.1
E.	Latent pred. (F. 4a)	83.0	58.9	71.4	71.2	62.0	43.6	84.0	71.7	60.2	44.7	86.1	80.2	46.5	34.4	93.9	89.5
щ	Hybrid (F. 4b)	82.7	58.5	71.6	72.1	61.8	43.3	84.0	71.8	60.4	44.8	85.9	80.9	46.1	34.2	94.1	89.4
2-L	Separate joints	82.9	57.7	71.0	70.9	60.9	42.1	84.4	73.4	59.1	42.2	88.0	85.3	41.6	32.0	95.1	92.1
letV	Consistency regul.	81.0	57.4	72.8	74.8	60.6	41.7	84.7	74.3	57.9	41.8	88.2	84.7	40.6	30.7	95.7	92.6
Eff	Hybrid	81.3	57.9	72.4	73.9	61.1	42.0	84.6	74.3	59.2	42.8	87.2	84.3	41.8	31.4	95.6	92.6

5. Results

5.1. Benefit of Training on Many Datasets

Since one contribution of our paper is the study of the large-scale multi-dataset training regime, an important question is whether this brings improvements or whether performance saturates with just a few large-scale datasets. As a simple baseline, we train models on individual datasets and evaluate on the corresponding test splits. (With MuPoTS, we use MuCo-3DHP for training). We then train on three dataset combinations, as shown in Tab. 2. There is a clear trend showing performance improvement when training with more datasets, and the small dataset combination also outperforms single-dataset baselines. We note that H36M scores sometimes suffer from additional data. H36M uses the same studio environment in the training and evaluation split, therefore the model works better when a large part of the training batches are filled with H36M examples, allowing it to specialize on images from this room, but this does not reflect true generalization ability. The model trained on the large dataset combination achieves very strong scores across the board, confirming that using many datasets makes a difference.

Despite the good benchmark scores, we qualitatively observe (Fig. 2) that the different skeleton outputs can still be inconsistent among themselves.

5.2. Consistent Multi-Skeleton Prediction

A first naive baseline for achieving consistent predictions is to merge joints from different skeletons (*e.g.*, we predict only one "left shoulder" joint), reducing the joint count from 555 to 163. This leads to weaker results than predicting all joints separately (see Tab. 3), since joints with similar names may represent somewhat different keypoints.

H36M is again an outlier, as the prediction with merged joints works well for it. Since the model can easily recognize that a test image comes from the H36M studio, it can adapt its prediction to match the H36M skeleton format. This is not possible on *e.g.*, 3DPW, since the model cannot know in advance what skeleton format will be used for the reference poses of these images, since they come from diverse in-the-wild scenes.

When using our proposed ACAE-based regularization (c.f. Fig. 3c), we can see consistent improvements for almost all metrics. However, the improvement in the qualitative performance of the model is even more striking. As seen in Fig. 2, the regularized model creates significantly more consistent skeleton predictions. Especially the depth-consistency is improved, but some errors in the frontal view are also corrected. More qualitative results in the supplementary show that this observation holds broadly.

Overall, the model that estimates latent keypoints (Fig. 4a) has slightly lower performance than the separatehead baseline, likely because latent keypoints may be placed at less characteristic locations on the body and can thus be harder to localize. Further, the latent keypoint head's weights are initialized from scratch, whereas the regularization-based method fine-tunes a pre-trained head. The hybrid combination from Fig. 4a performs slightly worse than the model that is only regularized, but in many cases still outperforms the baseline. This shows that a direct

Table 4:Comparison to recent state-of-the-art works.(*twice as long training, 384 px resolution)

	MuPoTS		3DPW		3	3DHP			
	PCK ₁₅₀ ↑	MPJPE↓	PMPJPE	↓ PCK ₅₀ ↑	MPJPH	E↓ PCK ₁₅₀ ↑	MPJPE↓		
ROMP [69]	-	80.1	56.8	36.5	_	_	_		
Lin [44]	_	74.7	45.6	_	_	_	51.2		
PoseAug [22]	_	_	_	_	71.1	89.2	50.2		
Cheng [13]	89.6	-	-	-	-	-	49.3		
Ours RN50	92.2	65.5	47.2	49.0	64.2	93.3	45.8		
Ours EffV2S	93.7	61.5	43.0	51.8	60.0	95.3	45.2		
Ours EffV2L	94.1	60.6	41.7	52.1	59.2	95.8	40.6		
Ours EffV2L*	95.4	58.9	39.5	53.9	55.4	97.1	36.5		

estimation of the discovered latent keypoints is also a viable option. By design, this approach also produces consistent results, since we compute a single latent set of keypoints from which we decode all skeletons.

We also train the regularization and hybrid variants with EffNetV2-Large (lower part of Tab. 3). Overall, the results follow the same order, and they are better across the board. Regularization improves results and also leads to consistent predictions, and the hybrid approach is somewhat better than the initial model trained to predict separate joints.

This means that our autoencoder-based regularization is effective at improving results both quantitatively and qualitatively, and the discovered latent keypoints can be predicted directly. This opens up interesting future research directions, as the latent keypoints can be seen as a model agnostic interface, potentially allowing us to incorporate new skeleton formats by expanding the decoder, without a need for model specific fine-tuning or probing.

5.3. Comparison to Prior Works

In Tab. 4, we compare our final results to recent stateof-the-art published works (using standard protocols) and observe much better accuracy than SOTA models. We emphasize that this comparison is not "fair" w.r.t. the amount of training data. However, our goal in this paper is to show the value in large-scale multi-dataset training, and to investigate how to best supervise models in that setting.

5.4. Ablations

Chirality Equivariance Constraints. In Tab. 5, we analyze the effect of enforcing chirality-equivariance on the ACAE. In the quantitative metrics, we see approximately no change or a slight positive effect on both evaluated models. Given that symmetry makes sense as an inductive bias, we use chirality equivariance in our default setting.

Latent Keypoint Count. Fig. 5 shows that once a minimum number of latent points is reached, the reconstruction error only decreases slowly (evaluated on a held-out pseudo-GT validation set). We evaluate several latent sizes for fine-tuning in Tab. 6. 48 points work well in practice, and our regularization method is robust w.r.t. this hyperparameter. When directly predicting latent keypoints, using

Table 5: Effect of enforcing chirality equivariance constraints on the autoencoder weight matrices.

	ality	MuPoTS				3DPW	T	3DHP			
	Chir	MPJPE↓	РСК ₁₀₀	CPS ₂₀₀	MPJPE↓	PCK100	CPS ₂₀₀	MPJPE↓	PCK ₁₀₀	CPS ₂₀₀	
Cons. regul.		81.8	72.4	73.1	61.6	83.9	71.9	59.2	86.6	82.1	
Cons. regul.	\checkmark	81.8	72.5	72.9	61.5	84.0	71.9	59.2	86.6	82.7	
Hybrid		83.2	71.2	71.7	61.7	84.0	72.0	60.3	85.9	80.7	
Hybrid	\checkmark	82.7	71.6	72.1	61.8	84.0	71.8	60.4	85.9	80.9	

Table 6: Effect of the number of latent points on final performance. We use 48 as our default setting.

MuPoTS				3DPW			3DHP	
MPJPE↓	PCK ₁₀₀ ↑	CPS_{200}^{\uparrow}	MPJPE↓	PCK ₁₀₀ ↑	CPS ₂₀₀ ↑	MPJPE↓	PCK ₁₀₀ ↑	CPS ₂₀₀ ↑
81.6	72.4	73.1	62.0	83.9	71.7	58.9	86.5	81.7
82.2	72.1	72.5	61.8	83.9	71.8	59.2	86.5	81.9
81.8	72.5	72.9	61.5	84.0	71.9	59.2	86.6	82.7
82.3	72.0	73.0	61.8	83.8	71.8	59.2	86.6	82.1
86.0	69.4	65.0	67.2	81.2	64.6	70.2	80.0	62.2
82.7	71.4	72.1	62.3	83.9	71.8	60.2	86.2	81.4
82.7	71.6	72.1	61.8	84.0	71.8	60.4	85.9	80.9
84.1	70.6	67.0	62.2	83.7	71.3	62.0	84.5	80.2
$30 \\ 25 \\ 20 \\ 15 \\ 10 \\ 5 \\ 0$	32 64	128		256	~~~~~		51	2 555
	I MPJPE↓ 81.6 82.2 81.8 82.3 86.0 82.7 82.7 82.7 82.7 82.7 82.7 82.7 82.7 82.7 82.7 82.7 82.7 80 25 20 15 0	MuPoTS MPJPE4 PCK100 [†] 81.6 72.4 82.2 72.1 81.8 72.5 82.3 72.0 86.0 69.4 82.7 71.4 82.7 71.6 84.1 70.6 30 15 0 32 64	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	$\begin{tabular}{ c c c c c c c c c c c c c c c c c c c$	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	$\begin{array}{c c c c c c c c c c c c c c c c c c c $	$\begin{array}{c c c c c c c c c c c c c c c c c c c $

Figure 5: Intrinsic dimension analysis of our pseudo-ground truth with 555 joints. The residual error curve shows a characteristic elbow shape.

too few or too many latent keypoints has a negative effect, but the differences are small beyond 32.

6. Conclusion

We have proposed a principled, automatic approach to the problem of large-scale multi-skeleton training of 3D human pose estimation. Despite its practical relevance in exploiting a large number of 3D pose datasets in one training, this problem has been largely overlooked in the literature.

Our approach relies on a novel formulation of dimensionality reduction of sets of keypoints, via an affinecombining autoencoder with guaranteed built-in equivariances to common transformations. By regularizing a 3D human pose estimator's output to stay close to the learned latent space discovered by the autoencoder, we can more effectively share information between the different datasets, resulting in an overall more accurate and consistent pose estimator. We release code for data processing and training, as well as trained models to serve as high-quality off-the-shelf methods for downstream research.

Acknowledgments. This work was supported by the ERC Consolidator Grant project "DeeViSe" (ERC-CoG-2017-773161) and by Robert Bosch GmbH under the project "Context Understanding for Autonomous Systems".

References

- [1] N. P. van der Aa, X. Luo, G. J. Giezeman, R. T. Tan, and R. C. Veltkamp. Utrecht multi-person motion (UMPM) benchmark: a multi-person dataset with synchronized video and motion capture data for evaluation of articulated human motion and interaction. In *ICCV Workshops*, 2011.
- [2] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2D human pose estimation: New benchmark and state of the art analysis. In *CVPR*, 2014.
- [3] M. Andriluka, U. Iqbal, E. Insafutdinov, L. Pishchulin, A. Milan, J. Gall, and B. Schiele. PoseTrack: A benchmark for human pose estimation and tracking. In *CVPR*, 2018.
- [4] E. G. Bazavan, A. Zanfir, M. Zanfir, W. T. Freeman, R. Sukthankar, and C. Sminchisescu. HSPACE: Synthetic parametric humans animated in complex environments. arXiv preprint arXiv:2112.12867, 2021.
- [5] Y. Ben-Shabat, X. Yu, F. Saleh, D. Campbell, C. Rodriguez-Opazo, H. Li, and S. Gould. The IKEA ASM dataset: Understanding people assembling furniture through actions, objects and pose. In WACV, 2021.
- [6] B. L. Bhatnagar, X. Xie, I. Petrov, C. Sminchisescu, C. Theobalt, and G. Pons-Moll. BEHAVE: Dataset and method for tracking human object interactions. In *CVPR*, 2022.
- [7] H. Bourlard and S. H. Kabil. Autoencoders reloaded. *Biol. Cybern.*, 116(4):1–18, 2022.
- [8] H. Bourlard and Y. Kamp. Auto-association by multilayer perceptrons and singular value decomposition. *Biol. Cybern.*, 59(4):291–294, 1988.
- [9] C. S. Catalin Ionescu, Fuxin Li. Latent structured models for human pose estimation. In *ICCV*, 2011.
- [10] A. Chatzitofis, L. Saroglou, P. Boutis, P. Drakoulis, N. Zioulis, S. Subramanyam, B. Kevelham, C. Charbonnier, P. Cesar, D. Zarpalas, *et al.* HUMAN4D: A human-centric multimodal dataset for motions and immersive media. *IEEE Access*, 8:176241–176262, 2020.
- [11] Y. Chen, Y. Tian, and M. He. Monocular human pose estimation: A survey of deep learning-based methods. *CVIU*, 192:102897, 2020.
- [12] Y. Cheng, B. Wang, B. Yang, and R. T. Tan. Monocular 3D multi-person pose estimation by integrating top-down and bottom-up networks. In *CVPR*, 2021.
- [13] Y. Cheng, B. Wang, and R. Tan. Dual networks based 3D multi-person pose estimation from monocular video. *PAMI*, 2022. doi: 10.1109/TPAMI.2022.3170353. Early access.
- [14] J. P. Cunningham and Z. Ghahramani. Linear dimensionality reduction: Survey, insights, and generalizations. *JMLR*, 16 (1):2859–2900, 2015.
- [15] Q. Dang, J. Yin, B. Wang, and W. Zheng. Deep learning based 2D human pose estimation: A survey. *Tsinghua Science and Technology*, 24(6):663–676, 2019.
- [16] M. Fabbri, F. Lanzi, S. Calderara, A. Palazzi, R. Vezzani, and R. Cucchiara. Learning to detect and track visible and occluded body joints in a virtual world. In *ECCV*, 2018.
- [17] M. Fabbri, F. Lanzi, S. Calderara, S. Alletto, and R. Cucchiara. Compressed volumetric heatmaps for multi-person 3D pose estimation. In *CVPR*, 2020.

- [18] M. Fieraru, M. Zanfir, E. Oneata, A.-I. Popa, V. Olaru, and C. Sminchisescu. Three-dimensional reconstruction of human interactions. In *CVPR*, 2020.
- [19] M. Fieraru, M. Zanfir, E. Oneata, A.-I. Popa, V. Olaru, and C. Sminchisescu. Learning complex 3D human self-contact. In AAAI, 2021.
- [20] M. Fieraru, M. Zanfir, S.-C. Pirlea, V. Olaru, and C. Sminchisescu. AIFit: Automatic 3D human-interpretable feedback models for fitness training. In CVPR, 2021.
- [21] S. Ghorbani, K. Mahdaviani, A. Thaler, K. Kording, D. J. Cook, G. Blohm, and N. F. Troje. Movi: A large multipurpose human motion and video dataset. *PLOS One*, 16(6): e0253157, 2021.
- [22] K. Gong, J. Zhang, and J. Feng. PoseAug: A Differentiable Pose Augmentation Framework for 3D Human Pose Estimation. In *CVPR*, 2021.
- [23] R. Guerra-Urzola, K. Van Deun, J. C. Vera, and K. Sijtsma. A guide for sparse pca: model comparison and applications. *Psychometrika*, 86(4):893–919, 2021.
- [24] H. Guo, T. Tang, G. Luo, R. Chen, Y. Lu, and L. Wen. Multidomain pose network for multi-person pose estimation and tracking. In *ECCV Workshops*, 2018.
- [25] E. Hoffer, I. Hubara, and D. Soudry. Train longer, generalize better: Closing the generalization gap in large batch training of neural networks. In *NIPS*, 2017.
- [26] D. Holden, J. Saito, and T. Komura. A deep learning framework for character motion synthesis and editing. ACM Trans. Graphics, 35(4):1–11, 2016.
- [27] K. Hormann and N. Sukumar. Generalized barycentric coordinates in computer graphics and computational mechanics. CRC Press, 2017.
- [28] Y.-T. Hu, H.-S. Chen, K. Hui, J.-B. Huang, and A. G. Schwing. SAIL-VOS: Semantic amodal instance level video object segmentation a synthetic dataset and baselines. In *CVPR*, 2019.
- [29] C.-H. P. Huang, H. Yi, M. Höschle, M. Safroshkin, T. Alexiadis, S. Polikovsky, D. Scharstein, and M. J. Black. Capturing and inferring dense full-body human-scene contact. In *CVPR*, 2022.
- [30] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6M: Large scale datasets and predictive methods for 3D human sensing in natural environments. *PAMI*, 36(7): 1325–1339, 2014.
- [31] T. Jakab, A. Gupta, H. Bilen, and A. Vedaldi. Unsupervised learning of object landmarks through conditional image generation. In *NeurIPS*, 2018.
- [32] T. Jakab, A. Gupta, H. Bilen, and A. Vedaldi. Self-supervised learning of interpretable keypoints from unlabelled videos. In *CVPR*, 2020.
- [33] T. Jakab, R. Tucker, A. Makadia, J. Wu, N. Snavely, and A. Kanazawa. KeypointDeformer: Unsupervised 3D keypoint discovery for shape control. In *CVPR*, 2021.
- [34] X. Ji, Q. Fang, J. Dong, Q. Shuai, W. Jiang, and X. Zhou. A survey on monocular 3D human pose estimation. *Virtual Reality & Intelligent Hardware*, 2(6):471–500, 2020.
- [35] S. Johnson and M. Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *BMVC*, 2010.

- [36] H. Joo, T. Simon, X. Li, H. Liu, L. Tan, L. Gui, S. Banerjee, T. Godisart, B. Nabbe, I. Matthews, T. Kanade, S. Nobuhara, and Y. Sheikh. Panoptic Studio: A massively multiview system for social interaction capture. *PAMI*, 41(1):190–204, 2019.
- [37] H. Joo, N. Neverova, and A. Vedaldi. Exemplar fine-tuning for 3D human model fitting towards in-the-wild 3D human pose estimation. In *3DV*, 2021.
- [38] D. P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In *ICLR*, 2015.
- [39] M. Kocabas, N. Athanasiou, and M. J. Black. Vibe: Video inference for human body pose and shape estimation. In *CVPR*, 2020.
- [40] M. Kocabas, C.-H. P. Huang, J. Tesch, L. Müller, O. Hilliges, and M. J. Black. SPEC: Seeing people in the wild with an estimated camera. In *ICCV*, 2021.
- [41] J. N. Kundu, S. Seth, P. YM, V. Jampani, A. Chakraborty, and R. V. Babu. Uncertainty-aware adaptation for self-supervised 3D human pose estimation. In *CVPR*, 2022.
- [42] R. Li, S. Yang, D. A. Ross, and A. Kanazawa. AI choreographer: Music conditioned 3D dance generation with AIST++. In *ICCV*, 2021.
- [43] K. Lin, L. Wang, and Z. Liu. End-to-end human pose and mesh reconstruction with transformers. In CVPR, 2021.
- [44] K. Lin, L. Wang, and Z. Liu. Mesh graphormer. In *ICCV*, 2021.
- [45] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft COCO: Common objects in context. In *ECCV*, 2014.
- [46] R. Liu, J. Shen, H. Wang, C. Chen, S.-c. Cheung, and V. Asari. Attention mechanism exploits temporal contexts: Real-time 3D human pose reconstruction. In *CVPR*, 2020.
- [47] W. Liu and T. Mei. Recent advances of monocular 2D and 3D human pose estimation: A deep learning perspective. ACM Computing Surveys, 2022. doi: 10.1145/3524497. Early access.
- [48] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black. SMPL: A skinned multi-person linear model. ACM Trans. Graphics (Proc. SIGGRAPH Asia), 34(6):248:1– 248:16, 2015.
- [49] M. M. Loper, N. Mahmood, and M. J. Black. MoSh: Motion and shape capture from sparse markers. ACM Trans. Graphics (Proc. SIGGRAPH Asia), 33(6):220:1–220:13, 2014.
- [50] I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.
- [51] N. Mahmood, N. Ghorbani, N. F. Troje, G. Pons-Moll, and M. J. Black. AMASS: Archive of motion capture as surface shapes. In *ICCV*, 2019.
- [52] C. Mandery, Ö. Terlemez, M. Do, N. Vahrenkamp, and T. Asfour. Unifying representations and large-scale wholebody motion databases for studying human motion. *IEEE Trans. Robotics*, 32(4):796–809, 2016.
- [53] T. von Marcard, R. Henschel, M. Black, B. Rosenhahn, and G. Pons-Moll. Recovering accurate 3D human pose in the wild using IMUs and a moving camera. In *ECCV*, 2018.
- [54] R. Martin-Martin, M. Patel, H. Rezatofighi, A. Shenoi, J. Gwak, E. Frankel, A. Sadeghian, and S. Savarese. JRDB:

A dataset and benchmark of egocentric robot visual perception of humans in built environments. *PAMI*, 2021. doi: 10.1109/TPAMI.2021.3070543. Early access.

- [55] D. Mehta, H. Rhodin, D. Casas, P. Fua, O. Sotnychenko, W. Xu, and C. Theobalt. Monocular 3D human pose estimation in the wild using improved CNN supervision. In *3DV*, 2017.
- [56] D. Mehta, O. Sotnychenko, F. Mueller, W. Xu, S. Sridhar, G. Pons-Moll, and C. Theobalt. Single-shot multi-person 3D pose estimation from monocular RGB. In *3DV*, 2018.
- [57] G. Moon, J. Chang, and K. M. Lee. Camera distanceaware top-down approach for 3D multi-person pose estimation from a single rgb image. In *ICCV*, 2019.
- [58] A. Nibali, J. Millward, Z. He, and S. Morgan. ASPset: An outdoor sports pose video dataset with 3D keypoint annotations. *Image and Vision Computing*, 111:104196, 2021.
- [59] A. Nieto, Jesús R.and Susín. *Deformation Models: Tracking, Animation and Applications*, chapter Cage Based Deformations: A Survey, pages 75–99. Springer Netherlands, 2013. doi: 10.1007/978-94-007-5446-1_.
- [60] F. Ofli, R. Chaudhry, G. Kurillo, R. Vidal, and R. Bajcsy. Berkeley MHAD: A comprehensive multimodal human action database. In WACV, 2013.
- [61] P. Patel, C.-H. P. Huang, J. Tesch, D. T. Hoffmann, S. Tripathi, and M. J. Black. AGORA: Avatars in geography optimized for regression analysis. In *CVPR*, 2021.
- [62] G. Pavlakos, V. Choutas, N. Ghorbani, T. Bolkart, A. A. A. Osman, D. Tzionas, and M. J. Black. Expressive body capture: 3D hands, face, and body from a single image. In *CVPR*, 2019.
- [63] K. Pearson. On lines and planes of closest fit to systems of points in space. Lond. Edinb. Dublin philos. mag. j. sci., 2 (11):559–572, 1901.
- [64] A. Pumarola, J. Sanchez, G. Choi, A. Sanfeliu, and F. Moreno-Noguer. 3DPeople: Modeling the geometry of dressed humans. In *ICCV*, 2019.
- [65] M. Rapczyński, P. Werner, S. Handrich, and A. Al-Hamadi. A baseline for cross-database 3D human pose estimation. *Sensors*, 21(11):3769, 2021.
- [66] H. Rhodin, M. Salzmann, and P. Fua. Unsupervised geometry-aware representation for 3D human pose estimation. In *ECCV*, 2018.
- [67] I. Sárándi, T. Linder, K. O. Arras, and B. Leibe. MeTRAbs: metric-scale truncation-robust heatmaps for absolute 3D human pose estimation. *IEEE Trans. Biometr. Behav. Ident. Sci.* (*T-BIOM*), 3(1):16–30, 2021.
- [68] C. Summers and M. J. Dinneen. Four things everyone should know to improve batch normalization. In *ICLR*, 2019.
- [69] Y. Sun, Q. Bao, W. Liu, Y. Fu, B. Michael J., and T. Mei. Monocular, One-stage, Regression of Multiple 3D People. In *ICCV*, 2021.
- [70] M. Tan and Q. Le. EfficientNetV2: Smaller Models and Faster Training. In *ICML*, 2021.
- [71] M. Trumble, A. Gilbert, C. Malleson, A. Hilton, and J. Collomosse. Total capture: 3D human pose estimation fusing video and inertial sensors. In *BMVC*, 2017.
- [72] S. Tsuchida, S. Fukayama, M. Hamasaki, and M. Goto. AIST dance video database: Multi-genre, multi-dancer, and

multi-camera database for dance information processing. In *Proc. Int. Soc. Music Inf. Retr. Conf. (ISMIR)*, 2019.

- [73] G. Varol, J. Romero, X. Martin, N. Mahmood, M. J. Black, I. Laptev, and C. Schmid. Learning from synthetic humans. In *CVPR*, 2017.
- [74] B. Wandt, M. Rudolph, P. Zell, H. Rhodin, and B. Rosenhahn. Canonpose: Self-supervised monocular 3D human pose estimation in the wild. In *CVPR*, 2021.
- [75] Z. Wang, L. Chen, S. Rathore, D. Shin, and C. Fowlkes. Geometric pose affordance: 3D human pose with scene constraints. arXiv preprint arXiv:1905.07718, 2019.
- [76] H. Xu, E. G. Bazavan, A. Zanfir, W. T. Freeman, R. Sukthankar, and C. Sminchisescu. GHUM & GHUML: Generative 3D human shape and articulated pose models. In *CVPR*, 2020.
- [77] R. Yeh, Y.-T. Hu, and A. Schwing. Chirality nets for human pose regression. In *NeurIPS*, 2019.
- [78] Z. Yu, J. S. Yoon, I. K. Lee, P. Venkatesh, J. Park, J. Yu, and H. S. Park. HUMBI: A large multiview dataset of human body expressions. In *CVPR*, 2020.
- [79] D. Zhang, Y. Wu, M. Guo, and Y. Chen. Deep learning methods for 3D human pose estimation under different supervision paradigms: A survey. *Electronics*, 10(18):2267, 2021.
- [80] T. Zhang, B. Huang, and Y. Wang. Object-occluded human shape and pose estimation from a single color image. In *CVPR*, 2020.
- [81] W. Zhang, Z. Liu, L. Zhou, H. Leung, and A. B. Chan. Martial arts, dancing and sports dataset: a challenging stereo and multi-view dataset for 3D human pose estimation. *Image and Vision Computing*, 61:22–39, 2017.