

Human Pose Estimation in Monocular Omnidirectional Top-View Images

Jingrui Yu* Tobias Scheck Roman Seidel Yukti Adya Dipankar Nandi Gangolf Hirtz
 Chemnitz University of Technology, Germany

*jingrui.yu@etit.tu-chemnitz.de

Abstract

Human pose estimation (HPE) with convolutional neural networks (CNNs) for indoor monitoring is one of the major challenges in computer vision. In contrast to HPE in perspective views, an indoor monitoring system can consist of an omnidirectional camera with a field of view of 180° to detect the pose of a person with only one sensor per room. To recognize human pose, the detection of keypoints is an essential upstream step. In our work we propose a new dataset for training and evaluation of CNNs for the task of keypoint detection in omnidirectional images. The training dataset, THEODORE+, consists of 50,000 images and is created by a 3D rendering engine, where humans are randomly walking through an indoor environment. In a dynamically created 3D scene, persons move randomly with simultaneously moving omnidirectional camera to generate synthetic RGB images and 2D and 3D ground truth. For evaluation purposes, the real-world PoseFES dataset with two scenarios and 701 frames with up to eight persons per scene was captured and annotated. We propose four training paradigms to finetune or re-train two top-down models in MMPose and two bottom-up models in CenterNet on THEODORE+. Beside a qualitative evaluation we report quantitative results. Compared to a COCO pretrained baseline, we achieve significant improvements especially for top-view scenes on the PoseFES dataset. Our datasets can be found at https://www.tu-chemnitz.de/etit/dst/forschung/comp_vision/datasets/index.php.en.

1. Introduction

With the growing need for indoor monitoring, human pose estimation (HPE) has become one of the main research areas of modern computer vision research [23, 27]. However, most of this research was conducted on perspective views with a limited field of view which is not appropriate for indoor scenarios when an entire room is to be covered with only one sensor. For this problem, cameras with an omnidirectional camera model and wide-angle fisheye lenses are best suited. Nevertheless, fish-eye lenses have

the characteristic of distortion which makes it necessary to model these distortions implicitly in the training data. In this work, we propose a fuss free way to use convolutional neural networks (CNNs) for human pose estimation in omnidirectional images. The application field of this work is Ambient Assisted Living (AAL). AAL stands for concepts, products and services that introduce new technologies into daily living in order to improve the quality of life for people in all phases of life, especially in old age. These systems, here a camera-based smart sensor that is mounted in the ceiling of a room, monitor elderlies' activities to provide behaviour analysis results for care attendants, relatives and physicians. Besides a daily activity protocol, the system is capable to recognize emergency situations such as unpredictable human falls [30]. With the help of infrared lighting, video-based sleep quality estimation is employed to be able to draw conclusions about the current behaviour of an affected person as well as anomalies of the behaviour. The contribution of our work is threefold:

- we provide a top-view omnidirectional synthetic dataset with keypoint annotations, namely THEODORE+.
- For evaluation purposes, the PoseFES dataset, a new real-world top-view omnidirectional dataset with keypoint annotations was created.
- The adaption of state-of-the-art approaches for keypoint detection to a monocular top-view omnidirectional camera.

2. Related work

2.1. Human pose estimation in perspective and omnidirectional images

HPE is a popular topic in computer vision because of its wide range of applications, such as action and activity recognition, augmented reality (AR) and rehabilitation feedback systems. Early implementations have used hand-crafted features as well as pre-defined human models [11]. In recent years, researchers have shifted to deep learning-based methods. Xiao *et al.* [40] provided a simple base-

line by using ResNet-50 [13] to generate heatmaps for each keypoint. HRNet [33] further improved the results by introducing a high resolution backbone network. These are top-down methods, which use a separate person detector to isolate image areas with the person, and then perform keypoint estimation on them. In contrast, bottom-up methods are able to detect keypoints for multiple persons in a single inference pipeline. OpenPose [2] uses part affinity fields to associate the estimated keypoints with the individual persons in the image. CenterNet [45] estimates the offset of each joint to the person center point and use it to associate them with the persons.

Though HPE is well researched for perspective images, it is only starting to get attention for omnidirectional images. There has been research that focus on using head-mounted fisheye cameras for pose estimation, such as *Ego-Cap* [26], *Mo²Cap²* [41], *xR-EgoPose* [35]/*SelfPose* [34], *EgoGlass* [44] and Wang *et al.* [38]. However, the application of CNNs for HPE in overhead omnidirectional images has not been thoroughly investigated. This is largely due to the lack of training data in this domain.

Georgakopoulos *et al.* [5,10] employ a 3D human model to create a dataset of binary silhouettes, which are rendered through the calibration of the fisheye camera. The CNN is trained to differentiate between the pre-set poses, rather than estimate the joint positions. Haque *et al.* [12] train CNN and LSTM [15] to achieve view-point invariant 3d pose estimation on a singular depth image. Denecke and Jauch [6] use the 3D point cloud calculated by the smart sensor and prior knowledge of the human body to estimate the joint positions. The results of this method are restricted by factors such as the mounting position of the camera and differences between each individual body. The inference speed is limited by the speed of the smart sensor. Heindl *et al.* [14] generate multiple rectilinear views from a fish-eye image and perform 2D keypoint estimation using OpenPose without finetuning the model. The 3D skeleton is reconstructed by a stereo vision setup. This method requires that the calibration parameters of the cameras are known. The temporal performance is heavily restricted by the overhead of generating multiple views and inferencing on all of them. Garau *et al.* achieve viewpoint-invariant 3D HPE with a capsule auto-encoder named DECA [8] on depth and RGB images, namely ITOP and PanopTOP31K datasets.

2.2. Top-view HPE datasets

Haque *et al.* [12] introduced the Invariant-Top View Dataset (ITOP), which consists of 100K real-world depth images. It contains no RGB-images, therefore limiting the use of popular CNN-based HPE models. Garau *et al.* introduced the PanopTOP framework in [9] for generating semi-synthetic top-view human images of normal perspective camera with 2D- and 3D-pose groundtruth from the

multi-view dataset Panoptic [19]. With this framework they create the PanopTOP31K dataset. It contains top-view and front-view persons as well as the corresponding depth maps, point clouds and 3D meshes. This is the first semi realistic HPE dataset of RGB images from the top-view. However, there are a few shortcomings in this dataset. Firstly, the image resolution is very low at 256×256 pixels, while the persons in the images effectively occupy no more than 100×100 pixels. Secondly, there are a lot of artifacts in the synthesized images, the most severe of which is hand position ghosting, where there are multiple instances of each arm/hand in the images. Finally, the subjects are positioned in front of a white background and the camera position is fixed, thus the variations are low across the dataset.

2.3. Synthetic data generation for CNN training

There exist many datasets for the task HPE, such as MPII [1], Human3.6M [18] and COCO keypoints [21]. However, capturing real-world HPE data with accurate groundtruth annotations either require a specific motion capture system [18] or a large-scale manual annotation process [21]. Therefore, it is often expensive and time-consuming. Synthetic data generation, on the other hand, can create pixel precise annotations without additional steps. Additionally, it does not raise privacy concerns.

Song *et al.* [32] introduce 3D computer graphics to create interior scenes with realistic textures and furniture. The authors in [39] extended this concept with agent functionality to navigate freely in the 3D environment. McCormac *et al.* [22] describe methods for creating physical and photorealistic interior renderings. A similar approach is taken by Li *et al.* [20], but with a vast amount of professional interior designs and object assets while creating a synthetic dataset. In [3] the authors simulate different human bodies with various assigned textures to train CNNs to recognize poses. Hoffman *et al.* [16] compare synthetically generated humans to investigate the influence of synthetic data compared to real but augmented data. With SURREAL [37], the authors introduce a dataset with photorealistic computer-generated images containing annotations for body parts and action sequences. In a multi-agent simulation concerning interaction in household scenarios, Puig *et al.* [25] present VirtualHome. Since VirtualHome provides interactions of humans with objects, this simulation serves very well in an AAL context. With ElderSIM [17], another AAL-focused eldercare simulation is introduced. Here various modalities, including RGB videos, 2D and 3D skeletons and different camera angles, are provided. Concerning omnidirectional imagery, none of the above methods provide such content, which is mandatory if CNNs are to be applied in this domain. A simulation using computer graphics for omnidirectional ego-pose imagery is introduced in [26,34,35].

To obtain omnidirectional top-view images to train

CNNs, the work of Scheck *et al.* [28] introduces a procedure for generating such images, resulting in a dataset for object detection and semantic segmentation. This approach was extended by Seuffert *et al.* [31] with stereo image pairs and the corresponding depth maps. Seidel *et al.* [29] present an omnidirectional dataset for the optical flow, focusing on household activities. However, skeletal keypoints for simulated humans are still missing in these datasets for the application of HPE.

3. Datasets

3.1. THEODORE+ dataset

Building upon the work of Scheck *et al.* [28], a large-scale synthetic dataset of indoor omnidirectional scenes is generated. The implementation details closely follow the original work, utilizing six indoor settings and domain randomization to create a diverse dataset. The object textures and the human model parameters (height and weight) are randomized, and at the same time the camera position changes constantly to ensure varied perspectives. During the simulation, each character has eight pre-defined animations that belong to four categories: *sitting*, *lying*, *falling*, and *walking*. An animation is selected from this animation set and executed depending on the character’s action. For instance, when a character interacts with a chair, the sitting animation is triggered. Furthermore, a fall may be activated on the way to the selected chair in the virtual environment, which then executes the respective animation. The corresponding action of a virtual character is also part of the exported data and is usable for activity recognition. Our simulation uses Unity’s universal render pipeline (UDP) for rendering the images with a distortionless virtual fisheye camera. We did not use the high-definition render pipeline (HDRP) because of restrictions of the texture package we used, and it is debatable whether using resource intensive photorealistic synthetic images for training CNNs is advantageous [36].

Beside bounding boxes and segmentation masks, our dataset features full body 2D and 3D pose information and action information of the human model. The implementation based on the Unity engine is extended to map the human body skeleton points from the engine’s internal format to the COCO format [21] with 13 keypoints. No data is available for the simulated characters for the keypoints *left_eye*, *right_eye*, *left_ear* and *right_ear*. Therefore, these points receive the coordinates [0,0] during the export. The new dataset consists of 50,000 images with a resolution of 2048×2048 pixel and $\sim 160,000$ character instances. Each file is saved in PNG format to exclude compression artifacts in the exported images. The 2D keypoints are converted to pixel coordinate space during export, while the 3D keypoints remain unchanged. An overview of the dataset

modalities is shown in Fig. 1. The dataset contains a scene’s RGB image (1a), 2D Keypoints in pixel space (1b), and non-normalized 3D keypoints (1c) of each person. Furthermore, the keypoints’ occluded attribute is true if superimposed by another object and not visible to the camera (keypoints with a red border in 1b).

3.2. PoseFES dataset

We created the real-world dataset PoseFES for evaluation purposes by extending the FES Dataset [28] with one scenario and pose annotations. It consists of two sequences, which have been recorded in a laboratory apartment with an omnidirectional fisheye camera. The image resolution is 1680×1680 pixels. The first sequence, *Scenario 1* (Sc1), contains 400 frames (Record_00000.png – Record_00399.png), in which three persons walk through the apartment performing daily activities. Overlapping of persons is kept very seldom for this sequence. *Scenario 2* (Sc2) contains 301 frames (Record_00600.png – Record_00900.png), in which a maximum of eight persons appear at the same time. Heavy overlapping is present in most frames of this sequence. There are 735 and 2161 instances in Sc1 and Sc2, respectively.

Axis-aligned bounding boxes and keypoints are annotated for the dataset. The bounding boxes are generated using OmniPD [43] and then adjusted manually. 17 keypoints are annotated for all persons, which conform to the keypoints provided by COCO [21]. Two extra keypoints, *shoulder_center* and *hip_center* are extrapolated by averaging shoulder and hip keypoints, respectively. Annotations are available in CVAT and COCO format. Sample images as well as annotations are shown in Fig. 2. More images are available in Sec. 5.4 where we present some qualitative evaluation.

4. Training CNNs for pose estimation in omnidirectional images

4.1. Model choices

We choose three models to train on our synthetic dataset. **SimpleBaseline2D** [40] adds three de-convolution layers with batch normalization and ReLU activation to the commonly used backbone network ResNet [13]. A 1×1 convolutional layer at the end generates k heatmaps for k keypoints of a person object. The loss is calculated by the Mean Squared Error (MSE) between the predicted heatmaps and the groundtruth heatmaps. Using 256×192 pixels as input resolution and ResNet-50 backbone, the authors are able to reach state-of-the-art performance on the COCO Keypoints dataset with this simple structure.

HRNet [33] introduces a new network structure for HPE. It contains four parallel multi-resolution subnetworks. The first stem net reduces the input resolution by 4. The fol-




Figure 1. Example scene from THEODORE+ dataset with an RGB image, 2D and 3D keypoints visualized. Annotations are grouped per instance (person) by colour.




Figure 2. Example images with annotations from PoseFES dataset. Annotations are grouped per instance (person) by colour.

lowing subnetworks reduce the resolution to the half of the former one and simultaneously double the depth of the feature maps. Information is exchanged between each stage of the subnetworks by exchange units. At the end of the network, the heatmaps are predicted from the stem net which has the highest resolution.

Both SimpleBaseline2D and HRNet are top-down methods. We use **MMPose** [4] to train both networks. MMPose is a pose estimation toolbox under the OpenMMLab project. It supports many state-of-the-art methods as well as popular datasets. This enables us to conveniently deal with dataset manipulation and model evaluation. We use the COCO-pretrained models provided by MMPose as our baseline, namely *res50_coco_256x192-ec54d7f3_20200709.pth* for SimpleBaseline2D and *hrnet_w48_coco_384x288-314c8528_20200708.pth* for HRNet, both of which can be found in MMPose model zoo.

CenterNet [45] employs a completely different approach for object detection and pose estimation. Object detection is performed by estimating the heatmaps of the center points of objects and regressing the object size in x and y directions. Keypoints are regarded as properties of the center point, and thus regressed as offset values to the center point. In this way, CenterNet is able to perform object (in this case, person) detection and keypoint estimation for multiple objects at the same time. It is preferable for our AAL application because of its low calculation cost, which enables it to be implemented in an embedded platform. We use the original implementation by the authors¹. The baseline models are *multi_pose_hg_1x.pth* and *multi_pose_dla_1x.pth* from its model zoo.

¹<https://github.com/xingyizhou/CenterNet>

4.2. Training

The training paradigm is the same for MMPose and CenterNet:

- The afore mentioned pretrained models are finetuned on the THEODORE+ dataset. For MMPose, learning rate (LR) is reduced to 1/10 of the original LR for training on COCO. For CenterNet, LR is slightly reduced from $1.25e-4$ to $1e-4$. No LR decay is used for all models.
- All networks are trained from scratch on THEODORE+. For MMPose, LR starts at $5e-4$ and is decayed until $5e-5$. For CenterNet, the constant LR of $7e-4$ and $5e-4$ are used for the network with hourglass and DLA backbones respectively.
- All networks are trained from scratch on the combined dataset from COCO keypoints dataset and THEODORE+ (C&T+). LR settings are kept the same as in **b**).
- Pretrained models are finetuned on the combined dataset C&T+. Settings are kept the same as in **a**).

The pretrained models are first tested on the PoseFES dataset and the baseline performance is noted. During finetuning or training, the models are validated directly on PoseFES after every epoch. If the model outperforms the one from the last epoch, it is saved as the best performing model. We finetune for 30 epochs and train for 60 epochs and choose the best performing models to fully test on PoseFES.

5. Evaluation results on PoseFES

5.1. Evaluation metrics

We use the object keypoint similarity (OKS)-based average precision (AP) and average recall (AR) as defined in COCO [21] to evaluate the trained models. COCO API evaluates 17 keypoints by default. However, THEODORE+ dataset only has 13 keypoints. Due to catastrophic forgetting [7], the model is unable to estimate the positions of eyes and ears after training or finetuning solely with THEODORE+ dataset. Therefore, we evaluate all the models on the 13 keypoints, excluding eyes and ears. In this case, the COCO API is adapted by deleting the sigma values for eyes and ears. Sigma values for other keypoints are kept unchanged. The models that are trained or finetuned with the combined dataset (training routines **b**) and **c**) are evaluated additionally on all 17 keypoints. Unfortunately, there exists no commonly used large scale dataset for evaluating HPE in top-down view of a fisheye camera. The PanopTOP31K dataset [9] is not suitable due to the shortcomings mentioned in Sec. 2.2 and the fact that the images

are not generated with a fisheye camera model. Therefore, the performance of all models are evaluated based on their performance on PoseFES dataset.

For SimpleBaseline2D and HRNet, the person bounding boxes are required for evaluation, since they are top-down methods. We provide the models with ground truth bounding boxes, which means the person detection accuracy is 100%. To estimate the influence of the person detection accuracy, we tested the models finetuned on C&T+ using bounding boxes that are inferred by OmniPD [43], whose accuracy on PoseFES is 85.6%. CenterNet performs person detection and keypoints estimation at the same time, therefore it does not need bounding box input.

5.2. Evaluation results of MMPose models




Figure 3. Comparison of MMPose models before and after training, evaluated on PoseFES for 13 KPs. Top diagram shows evaluation AP on Sc1 and bottom diagram shows evaluation AP on Sc2. C corresponds to COCO-pretrained models, T+ ft, T+, C&T+ and C&T+ ft correspond to training routines **a**) - **d**), respectively. ft stands for finetuning. AP is given at OKS = 0.5 : 0.05 : 0.95.

We first evaluate the COCO pretrained models and our trained models on the PoseFES Dataset for 13 keypoints on scenario 1 and 2 separately. Fig. 3 visualizes the evaluation APs of MMPose models on PoseFES as bar charts, so that the trends can be easily spotted. Commonly for MMPose models *SimpleBaseline2D* and *HRNet*, finetuning pretrained models on THEODORE+ alone results in heavy performance degradation. For both scenarios, AP is reduced by 10% for SimpleBaseline2D and over 20% for HRNet. Predictably, training from scratch on THEODORE+ alone

yields very low performance. Finetuning on the combined dataset results in performance improvements of 1%–2% for both models. The two models behave differently when trained from scratch on the combined dataset. For HRNet, the trained model is slightly worse than the COCO pretrained model, while for SimpleBaseline2D the results are slightly better than the finetuned model for scenario 1 but worse for scenario 2. The complete evaluation results are shown in Tab. 1. The evaluation results using OmniPD bounding boxes are denoted C&T+ ft w/PD.

Table 1. Evaluation results of MMPose models for 13 KPs on PoseFES Dataset. The best results are marked bold.

Model		SimpleBaseline2D		HRNet	
Sc	Dataset	AP	AR	AP	AR
1	COCO	72.2	74.8	81.0	83.0
	T+ ft	61.5	63.6	60.9	63.1
	T+	19.9	23.6	44.8	47.6
	C&T+	74.9	77.6	79.2	81.6
	C&T+ ft	74.6	76.8	81.8	83.7
	C&T+ ft w/ PD	60.1	61.5	65.4	66.7
2	COCO	68.1	74.8	77.0	79.2
	T+ ft	57.8	61.5	50.9	62.2
	T+	9.1	13.8	21.2	27.8
	C&T+	67.3	70.6	75.7	77.7
	C&T+ ft	70.0	72.6	78.6	80.7
	C&T+ ft w/ PD	59.0	60.5	63.8	64.9

Scenario 2 is more difficult for both models. Overall, the highest AP and AR values are reached by finetuning pretrained models on the combined dataset. For HRNet, the improvement is 1.6% for AP and 1.5% for AR. It is less consistent for SimpleBaseline2D.

Table 2. Evaluation results of MMPose models for 17 KPs on PoseFES Dataset. The best results are marked bold.

Model		SimpleBaseline2D		HRNet	
Sc	Training	AP	AR	AP	AR
1	COCO	67.4	69.9	75.7	78.2
	C&T+	66.9	69.7	72.2	75.0
	C&T+ ft	67.4	69.7	76.1	78.3
	C&T+ ft w/ PD	53.7	55.3	60.4	62.1
2	COCO	63.2	67.0	71.9	74.6
	C&T+	60.0	63.4	68.4	71.7
	C&T+ ft	62.6	65.8	72.9	75.3
	C&T+ ft w/ PD	52.3	54.3	58.9	60.2

Tab. 2 lists evaluation results for 17 keypoints. For HRNet the results stay in line with those for 13 keypoints. Finetuning pretrained model on combined dataset yields the best performance, albeit less than 1% over COCO pretrained

model. For SimpleBaseline2D, training or finetuning on the combined dataset provides no benefit over COCO pretrained model.

5.3. Evaluation results of CenterNet models




Figure 4. Comparison of CenterNet models before and after training, evaluated on PoseFES for 13 KPs. *CN-DLA* stands for CenterNet with DLA [42] backbone, and *CN-HG* stands for CenterNet with Hourglass [24] backbone.

Table 3. Evaluation results of CenterNet models for 13 KPs on PoseFES Dataset. The best results are marked bold.

Model		CN-DLA		CN-HG	
Sc	Dataset	AP	AR	AP	AR
1	COCO	31.1	35.0	34.8	38.4
	T+ ft	29.2	34.8	51.0	55.5
	T+	38.9	44.4	44.0	48.8
	C&T+	41.2	44.4	48.4	52.6
	C&T+ ft	45.9	50.7	53.5	56.6
2	COCO	23.0	27.6	24.9	30.4
	T+ ft	10.4	17.9	34.3	42.0
	T+	17.6	27.4	21.5	31.1
	C&T+	26.7	34.9	39.5	49.1
C&T+ ft	34.5	44.0	43.9	51.7	

CenterNet shows very different behaviours than MMPose models. In Fig. 4, both variants benefit from joint training of object detection and keypoints estimation when evaluated on PoseFES. Both models significantly outperform the COCO pretrained models except for CN-DLA

Table 4. Evaluation results of CenterNet models for **person detection** on PoseFES Dataset.

Sc	Model	CN-DLA		CN-HG	
	Dataset	AP	AR	AP	AR
1	COCO	44.2	52.7	47.9	55.2
	C&T+ ft	48.4	56.0	53.8	62.0
2	COCO	56.9	64.3	54.5	61.9
	C&T+ ft	65.9	71.7	69.8	74.4

when finetuned only on THEODORE+. To isolate the influence of the improvement in person detection on HPE, we compare the evaluation results of keypoint estimation in Tab. 3 to the evaluation results of bounding boxes in Tab. 4. To keep it simple, we compare only the models of finetuning with the combined dataset with the COCO pretrained model. CN-DLA’s improvement in person detection AP is 4.2% for Sc1 and 9% for Sc2. Its improvement in keypoint estimation AP is 14.8% for Sc1 and 11.5% for Sc2. For CN-HG, the person detection AP improvement lies at 5.9% for Sc1 and 15.3% for Sc2. Its improvement in keypoint estimation AP is even more significant at 18.7% for Sc1 and 19%. We can conclude that the relationship between improvement in person detection and keypoint estimation is not proportional. Keypoint estimation benefits much more from the training than person detection.

Evaluation results for 17 keypoints are listed in Tab. 5. Similarly to 13 keypoints, both models perform best when COCO pretrained model is finetuned on combined COCO and THEODORE+ dataset.

Table 5. Evaluation results of CenterNet models for 17 KPs on PoseFES Dataset. The best results are marked bold.

Sc	Model	CN-DLA		CN-HG	
	Training	AP	AR	AP	AR
1	COCO	28.1	31.7	32.4	35.7
	C&T+	29.6	32.7	34.4	38.4
	C&T+ ft	35.1	40.3	38.1	42.4
2	COCO	19.9	24.7	21.7	27.0
	C&T+	15.1	22.4	24.1	34.4
	C&T+ ft	23.3	33.3	27.7	36.1

5.4. Qualitative evaluation

In this section we evaluate the quality of the estimation results by the two most relevant models, HRNet and CenterNet with Hourglass backbone, both finetuned on C&T+.

We conclude from the quantitative evaluation that the finetuning provides only an improvement of the AP by less than 2% for **HRNet**, which seems not significant. Persons at the edges or in the outer circle of the omnidirectional view




Figure 5. Comparison of estimation examples by **HRNet** finetuned on C&T+. The cropped area are marked with red dashed bounding boxes in the original image.

look similar to those in the perspective image when standing or lying down. These instances are therefore not difficult for a good top-down model like HRNet even without finetuning. However, the appearance of the body changes dramatically when the person bends over or goes to the center of the omnidirectional image as shown in Fig. 5. Our finetuned model shows a clear improvement over the original model in these critical cases. In addition, we observe improvements in estimating joints in occluded body parts either by other body parts of the same person or by another person.

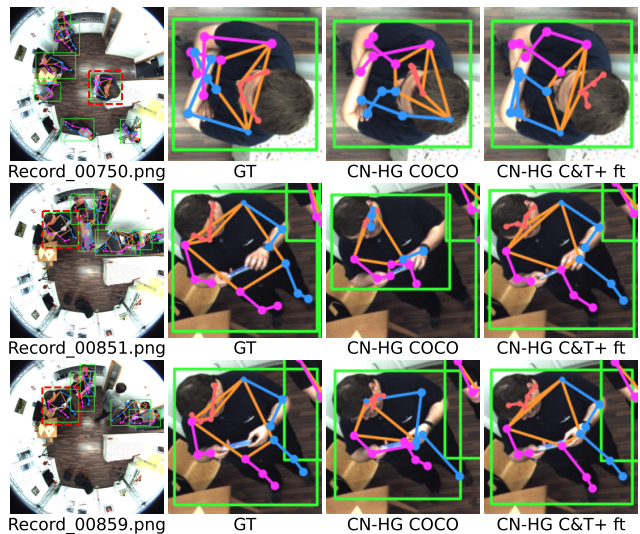


Figure 6. Instances of improved estimation results by CenterNet.




Figure 7. A failed detection and joint estimation by CenterNet.

The improvement is more evident for the **CenterNet** model with Hourglass backbone. At the confidence threshold of 0.4, the finetuned model can reliably detect persons in the scene while keeping false positives at minimum. Smaller instances around the edges and instances near the center of the image are detected more reliably than the COCO pretrained model. The estimated keypoints are more accurate, especially for instances in the center directly under the camera. Fig. 6 shows better estimation results by the finetuned CenterNet model because of more accurate person detection and more precise joint estimation. However, the CenterNet model clearly falls behind the HRNet model. Estimation of occluded joints is hardly improved over COCO pretrained model. Failure cases such as in Fig. 7 is often the result of faulty person detection. Undetected person naturally means no joint estimation, for example the instance in Rechord_00625.png shown in Fig. 5 cannot be detected by the finetuned CenterNet model.

5.5. Temporal performance

Our test bench is a workstation with an Intel Core i9-9960X CPU and 128 GB of DDR4 memory. The GPU used for inference is an Nvidia Titan RTX with 24 GB of GDDR memory.

The MMPose models uses groundtruth bounding boxes, therefore the inference time for person detection is not taken into account. HRNet at the resolution of 384×288 pixels needs on average 214 ms per image for inferencing on Sc1 of PoseFES and 574 ms on Sc2, which makes an average of 369 ms per image on the whole dataset. The inference time depends on the number of instances, as is the same with other top-down methods. The average inference time per instance is 89 ms for HRNet. The corresponding values for SimpleBaseline2D is 159 ms per image for Sc1, 468 ms per image for Sc2, 291 ms per image for the whole dataset and 71 ms per instance at the resolution of 256×192 pixels. Using OmniPD for person detection adds about 20 ~ 30 ms overhead per image to the top-down pipeline.

CenterNet is a bottom-up model and its speed does not depend on the number of instances per image. The measured inference time for CenterNet with DLA backbone is 53 ms per image, and for CenterNet with Hourglass backbone it is 173 ms. Both models process the inputs at the resolution of 512×512 pixels.

6. Conclusion

In this paper, we provide THEODORE+, a new synthetic omnidirectional top-view dataset with 50,000 RGB-images and annotated 2D and 3D keypoints, containing four action categories: *sitting*, *lying*, *falling* and *walking*, rendered through the unity engine. Furthermore, for evaluation purposes, the PoseFES dataset, which contains 701 images of real-world omnidirectional indoor scenes was manually annotated with 2D keypoints. We then trained and finetuned state-of-the-art top-down and bottom-up HPE models for perspective images to successfully perform HPE on PoseFES dataset, namely SimpleBaseline2D, HRNet and CenterNet with two different backbone networks.

The key findings of our work can be summarized as follows. We figured out that (1) the training on synthetic data from a *perfect* omnidirectional camera model significantly improve the results of real-world data, captured by a *non-calibrated* omnidirectional camera. Therefore, the network is able to deal with the distortion from different omnidirectional camera models. (2) The recognition of persons and keypoints close to the center of the image are much more improved than side-viewed persons. This corresponds to the top-view task, which is one of our motivation to generate THEODORE+ for the home monitoring of elderly. Our key finding (3) is already observed in the work of Scheck *et al.* [28], and it could be confirmed through our work that synthetic object detection bounding boxes improve the results for both disciplines, top-view object detection and top-view keypoint estimation.

Optimizing our best trained model to an embedded platform for person monitoring in an AAL-context is our major next step. Further investigations based on THEODORE+ will be activity recognition and the lifting of 2D human poses to 3D poses with and without stereo image data. Thanks to our finetuned models, it is possible to create a large-scale real-world human pose estimation dataset for omnidirectional top-view with limited resources in the future. Our datasets and the synthetic data generation pipeline will be open for download after the publication of this paper.

References

- [1] Mykhaylo Andriluka, Leonid Pishchulin, Peter Gehler, and Bernt Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2014. 2
- [2] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7291–7299, 2017. 2
- [3] Wenzheng Chen, Huan Wang, Yangyan Li, Hao Su, Zhenhua Wang, Changhe Tu, Dani Lischinski, Daniel Cohen-Or, and Baoquan Chen. Synthesizing training images for

- boosting human 3d pose estimation. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 479–488. IEEE, 2016. 2
- [4] MMPose Contributors. Openmmlab pose estimation toolbox and benchmark. <https://github.com/open-mmlab/mmpose>, 2020. 4
- [5] Konstantinos K Delibasis, Spiros V Georgakopoulos, Konstantina Kottari, Vassilis P Plagianakos, and Ilias Maglogiannis. Geodesically-corrected zernike descriptors for pose recognition in omni-directional images. *Integrated Computer-Aided Engineering*, 23(2):185–199, 2016. 2
- [6] Julia Denecke and Christian Jauch. Verification and regularization method for 3d-human body pose estimation based on prior knowledge. *Electronic Imaging*, 33:1–8, 2021. 2
- [7] Robert M French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999. 5
- [8] Nicola Garau, Niccolo Bisagno, Piotr Bródka, and Nicola Conci. Deca: Deep viewpoint-equivariant human pose estimation using capsule autoencoders. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11677–11686, 2021. 2
- [9] Nicola Garau, Giulia Martinelli, Piotr Bródka, Niccolò Bisagno, and Nicola Conci. Panoptop: A framework for generating viewpoint-invariant human pose estimation datasets. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 234–242, October 2021. 2, 5
- [10] Spiros V Georgakopoulos, Konstantina Kottari, Kostas Delibasis, Vassilis P Plagianakos, and Ilias Maglogiannis. Pose recognition using convolutional neural networks on omni-directional images. *Neurocomputing*, 280:23–31, 2018. 2
- [11] Wenjuan Gong, Xuena Zhang, Jordi González, Andrews Sobral, Thierry Bouwmans, Changhe Tu, and El-hadi Zahzah. Human pose estimation from monocular images: A comprehensive survey. *Sensors*, 16(12), 2016. 1
- [12] Albert Haque, Boya Peng, Zelun Luo, Alexandre Alahi, Serena Yeung, and Li Fei-Fei. Towards viewpoint invariant 3d human pose estimation. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part I 14*, pages 160–177. Springer, 2016. 2
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2, 3
- [14] Christoph Heindl, Thomas Pönitz, Andreas Pichler, and Josef Scharinger. Large area 3d human pose detection via stereo reconstruction in panoramic cameras. *arXiv preprint arXiv:1907.00534*, 2019. 2
- [15] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997. 2
- [16] David T. Hoffmann, Dimitrios Tzionas, Michael J. Black, and Siyu Tang. Learning to train with synthetic humans. In *German Conference on Pattern Recognition*, pages 609–623. Springer, 2019. 2
- [17] Hochul Hwang, Cheongjae Jang, Geonwoo Park, Junghyun Cho, and Ig-Jae Kim. Eldersim: A synthetic data generation platform for human action recognition in eldercare applications. *arXiv preprint arXiv:2010.14742*, 2020. 2
- [18] Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(7):1325–1339, jul 2014. 2
- [19] Hanbyul Joo, Hao Liu, Lei Tan, Lin Gui, Bart Nabbe, Iain Matthews, Takeo Kanade, Shohei Nobuhara, and Yaser Sheikh. Panoptic studio: A massively multiview system for social motion capture. In *The IEEE International Conference on Computer Vision (ICCV)*, 2015. 2
- [20] Wenbin Li, Sajad Saeedi, John McCormac, Ronald Clark, Dimos Tzoumanikas, Qing Ye, Yuzhong Huang, Rui Tang, and Stefan Leutenegger. Interiornet: Mega-scale multi-sensor photo-realistic indoor scenes dataset. In *British Machine Vision Conference (BMVC)*, 2018. 2
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 3, 5
- [22] John McCormac, Ankur Handa, Stefan Leutenegger, and Andrew J. Davison. Scenenet rgb-d: Can 5m synthetic images beat generic imagenet pre-training on indoor segmentation? In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2678–2687, 2017. 2
- [23] Lars Meinel, Michel Findeisen, Markus Heß, André Apitzsch, and Gangolf Hirtz. Automated real-time surveillance for ambient assisted living using an omnidirectional camera. In *2014 IEEE International Conference on Consumer Electronics (ICCE)*, pages 396–399, 2014. 1
- [24] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016. 6
- [25] Xavier Puig, Kevin Ra, Marko Boben, Jiaman Li, Tingwu Wang, Sanja Fidler, and Antonio Torralba. Virtualhome: Simulating household activities via programs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8494–8502, 2018. 2
- [26] Helge Rhodin, Christian Richardt, Dan Casas, Eldar Insafutdinov, Mohammad Shafiei, Hans-Peter Seidel, Bernt Schiele, and Christian Theobalt. Egocap: egocentric marker-less motion capture with two fisheye cameras. *ACM Transactions on Graphics (TOG)*, 35(6):1–11, 2016. 2
- [27] Julia Richter, Christian Wiede, and Gangolf Hirtz. Pose estimation and movement detection for mobility assessment of elderly people in an ambient assisted living application. In Ana Fred, Maria De Marsico, and Mário Figueiredo, editors, *Pattern Recognition: Applications and Methods*, pages 172–184. Cham, 2015. Springer International Publishing. 1
- [28] T. Scheck, R. Seidel, and G. Hirtz. Learning from theodore: A synthetic omnidirectional top-view indoor dataset for deep transfer learning. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 932–941, Los

- Alamitos, CA, USA, mar 2020. IEEE Computer Society. 3, 8
- [29] Roman Seidel, André Apitzsch, and Gangolf Hirtz. Omniflow: Human omnidirectional optical flow. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3678–3681, 2021. 3
- [30] Roman Seidel, Tobias Scheck, Ana C. Perez Grassi, Julian B. Seuffert, André Apitzsch, Jingrui Yu, Norbert Nestler, Danny Heinz, Lars Lehmann, Anne Goy, and Gangolf Hirtz. Contactless interactive fall detection and sleep quality estimation for supporting elderly with incipient dementia. *Current Directions in Biomedical Engineering*, 6(3):388–391, 2020. 1
- [31] Julian Seuffert, Ana Pérez Grassi, Tobias Scheck, and Gangolf Hirtz. A study on the influence of omnidirectional distortion on cnn-based stereo vision. In *VISIGRAPP (5: VIS-APP)*, pages 809–816, 2021. 3
- [32] Shuran Song, Fisher Yu, Andy Zeng, Angel X. Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1746–1754, 2017. 2
- [33] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5693–5703, 2019. 2, 3
- [34] Denis Tome, Thiemo Alldieck, Patrick Peluse, Gerard Pons-Moll, Lourdes Agapito, Hernan Badino, and Fernando De la Torre. Selfpose: 3d egocentric pose estimation from a headset mounted camera. *arXiv preprint arXiv:2011.01519*, 2020. 2
- [35] Denis Tome, Patrick Peluse, Lourdes Agapito, and Hernan Badino. xr-egopose: Egocentric 3d human pose from an hmd camera. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7728–7738, 2019. 2
- [36] Jonathan Tremblay, Aayush Prakash, David Acuna, Mark Brophy, Varun Jampani, Cem Anil, Thang To, Eric Cameracci, Shaad Boochoon, and Stan Birchfield. Training deep networks with synthetic data: Bridging the reality gap by domain randomization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2018. 3
- [37] Gul Varol, Javier Romero, Xavier Martin, Naureen Mahmood, Michael J. Black, Ivan Laptev, and Cordelia Schmid. Learning from synthetic humans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 109–117, 2017. 2
- [38] Jian Wang, Lingjie Liu, Weipeng Xu, Kripasindhu Sarkar, and Christian Theobalt. Estimating egocentric 3d human pose in global space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11500–11509, 2021. 2
- [39] Yi Wu, Yuxin Wu, Georgia Gkioxari, and Yuandong Tian. Building generalizable agents with a realistic and rich 3d environment. *arXiv preprint arXiv:1801.02209*, 2018. 2
- [40] Bin Xiao, Haiping Wu, and Yichen Wei. Simple baselines for human pose estimation and tracking. In *European Conference on Computer Vision (ECCV)*, 2018. 1, 3
- [41] Weipeng Xu, Avishek Chatterjee, Michael Zollhoefer, Helge Rhodin, Pascal Fua, Hans-Peter Seidel, and Christian Theobalt. Mo 2 cap 2: Real-time mobile 3d motion capture with a cap-mounted fisheye camera. *IEEE transactions on visualization and computer graphics*, 25(5):2093–2101, 2019. 2
- [42] Fisher Yu, Dequan Wang, Evan Shelhamer, and Trevor Darrell. Deep layer aggregation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2403–2412, 2018. 6
- [43] Jingrui Yu, Roman Seidel, and Gangolf Hirtz. Omnipd: One-step person detection in top-view omnidirectional indoor scenes. *Current Directions in Biomedical Engineering*, 5(1):239–244, 2019. 3, 5
- [44] Dongxu Zhao, Zhen Wei, Jisan Mahmud, and Jan-Michael Frahm. Egoglass: Egocentric-view human pose estimation from an eyeglass frame. In *2021 International Conference on 3D Vision (3DV)*, pages 32–41. IEEE, 2021. 2
- [45] Xingyi Zhou, Dequan Wang, and Philipp Krähenbühl. Objects as points. In *arXiv preprint arXiv:1904.07850*, 2019. 2, 4