

Efficient Realistic Data Generation Framework leveraging Deep Learning-based Human Digitization

C. Symeonidis, P. Nousi, P. Tosidis, K. Tsampazis, N. Passalis, A. Tefas, and
N. Nikolaidis

Artificial Intelligence and Information Analysis Lab, Department of Informatics,
Aristotle University of Thessaloniki, Thessaloniki, Greece
{charsyme,passalis,tefas,mnik}@csd.auth.gr

Abstract. The performance of supervised deep learning algorithms depends significantly on the scale, quality and diversity of the data used for their training. Collecting and manually annotating large amount of data can be both time-consuming and costly tasks to perform. In the case of tasks related to visual human-centric perception, the collection and distribution of such data may also face restrictions due to legislation regarding privacy. In addition, the design and testing of complex systems, e.g., robots, which often employ deep learning-based perception models, may face severe difficulties as even state-of-the-art methods trained on real and large-scale datasets cannot always perform adequately as they have not adapted to the visual differences between the virtual and the real world data. As an attempt to tackle and mitigate the effect of these issues, we present a method that automatically generates realistic synthetic data with annotations for a) person detection, b) face recognition, and c) human pose estimation. The proposed method takes as input real background images and populates them with human figures in various poses. Instead of using hand-made 3D human models, we propose the use of models generated through deep learning methods, further reducing the dataset creation costs, while maintaining a high level of realism. In addition, we provide open-source and easy to use tools that implement the proposed pipeline, allowing for generating highly-realistic synthetic datasets for a variety of tasks. A benchmarking and evaluation in the corresponding tasks shows that synthetic data can be effectively used as a supplement to real data.

Keywords: Synthetic data · human-centric visual analysis · person detection · pose estimation · face recognition

1 Introduction

The scale, diversity and quality of data used for training supervised deep learning methods have a major impact on their performance. Algorithms that are intended to be deployed on real-life conditions are usually trained on multiple

datasets in order to improve their generalization abilities and ensure their robustness. COCO [20] and Cityscapes [8] are only a few examples of large and diverse datasets providing annotations for training and evaluation of deep learning algorithms for computer vision tasks, such as object detection, human pose estimation, semantic image segmentation, etc. Collecting and manually annotating such amount of data is usually a challenging and exhausting task, requiring a lot of time and resources. In the case of visual human-centric analysis, the collection and distribution of such data may also face restrictions due to legislation regarding privacy. An alternative approach for collecting training data in a more automated manner is to generate them through a simulator. Indeed, in recent years, the use of synthetic datasets generated in this way has been established on the computer vision domain [10], [37], [35]. On the downside, these datasets often suffer in terms of realism and detail and/or are expensive to generate, requiring artists to carefully design specific models and environments.

Furthermore, apart from using synthetic data to create a new or enrich existing datasets, synthetic data are also often implicitly employed during the validation of complex systems, such as robots [24]. These systems are often designed and tested on simulation environments before being deployed in real life conditions. This allows for minimizing the risk of unwanted behaviors, that often lead to malfunctions that can destroy several hardware components, as well as for reducing development and validation time. However, deep learning methods trained on datasets containing only real data often exhibit an unstable behaviour or fail to perform adequately as they have not adapted to the visual differences between the simulated and the real world data, as we also experimentally demonstrate in this paper.

To overcome these limitations, in this paper we propose an effective and low-cost data generation method for human-centric tasks that:

- generates realistic and diverse data for person detection, pose estimation and face recognition,
- is capable of reusing and augmenting existing datasets, eliminating additional costs that often occur, as well as potential data collection restrictions regarding privacy legislation,
- provides automatically-generated and detailed annotations, and
- bridges the gap between the virtual and real world data, by enabling deep learning methods trained on the synthetic and real data to achieve highly accurate results on both.

At its core, the proposed pipeline for human-centric data generation uses as input real background images and carefully populates them with generated human figures in various poses. Instead of using hand-made 3D human models, we propose using models generated through deep learning methods. In this way, the proposed method eliminates most dataset creation costs, while maintaining a high level of realism. Finally, we provide an open-source implementation of the proposed method at <https://github.com/charsyme/Human-Data-Generation-Framework>, allowing researchers and practitioners easily use the proposed method for generating highly-realistic synthetic datasets for a variety of tasks.

The rest of the paper is structured as follows. The related work is discussed in Section 2. We also provide a brief background on DL-based perception for the three tasks examined in this work in Section 2.3. Then, the proposed method is described in detail in Section 3, while the dataset evaluation and benchmarking is provided in Section 4. Finally, conclusions are drawn in Section 5.

2 Related Work

2.1 Synthetic Data Generation for Computer Vision Methods

The use of synthetic datasets has recently gathered pace in deep learning and such datasets have proven their value either as a replacement or as an augmentation to existing training data [1,29]. Indeed, in [25], it is demonstrated that deep neural networks can achieve state-of-the-art results when trained on synthetic, yet realistic datasets. The domains of computer vision where synthetic datasets are widely used range from semantic image segmentation [31], and object detection [16] to pose estimation [35], and face recognition [23]. In [13], the authors trained scene-specific pedestrian detection models using only synthetic data. Surprisingly, those models managed to outperform models trained on real data. In [7], the authors presented an automatic approach that generates synthetic data for human pose estimation. In [34], a diverse, large and realistic synthetic dataset for human action recognition was generated containing a total of 39,982 videos and 35 action categories. Finally, synthetic datasets are also used for facial recognition. Such an example is [23], where the authors use 3D models of faces to modify existing images in order to generate novel poses and expressions.

Furthermore, an alternative to using fully synthetic datasets is the generation of mixed reality datasets, i.e., datasets that combine synthetic and natural content. In [28], the authors use 3D human models, which are rendered on random backgrounds in order to train a pedestrian detector. In a similar fashion, the authors of [2] render 3D vehicle models into existing captured real world background images. Inspired by these works, our method inserts realistic rendered 3D human models into existing natural background images, while trying to select appropriate scale and insertion locations, using a simple but effective approach. However, in contrast with the aforementioned approaches, the proposed method is capable of generating datasets for a wide variety of tasks, instead of being limited into a single tasks, e.g., pedestrian detection [28].

2.2 3D Human Model Generation from Natural Images

Methods for 3D human model generation from images can be divided based on whether they use a single or multiple (-view) images as input. Most single view methods use parametric models [22], [27] of human bodies and shapes, due to the fundamental depth ambiguity. In [18], the authors have shown that the pose and shape parameters of those models can accurately be estimated,

so that each model fits and aligns perfectly with the figures of the depicted people. Usually, parametric models capture body characteristics and movements of naked human models. However, recent works using parametric models can also capture the shape [3] and texture [19] of clothes, facial details, etc. Finally, the authors in [32], instead of using a parametric human model, use an implicit representation that locally aligns pixels of 2D images with the global context of their corresponding 3D object. The method uses full-body images of people both from single and multiple views as input. In our work, we leverage a state-of-the-art 3D human generation model, i.e., the Pixel-aligned Implicit Function (PIFu) proposed in [32]. This allows us to avoid most costs related to designing hand-crafted 3D human models. Instead, we dynamically create realistic 3D human models from existing datasets, that can be subsequently used for generating realistic and diverse datasets.

2.3 Deep Learning for Human-centric Perception

Person detection constitutes a crucial task related to human-centric perception, where the models aim to detect whether one or more persons exist in a given image, as well as to localize them. Given the proper training annotations, which are usually supplied in the form of bounding box coordinates for each person that appears in an image, generic object detection methods can be trained, such as the Single Shot Detector (SSD) [21]. It is also worth noting that person detection has been studied separately in recent literature, due in part to the single-class nature of this task. This can also be for the development of specialized methods, which can benefit from prior knowledge surrounding the task [30].

Face recognition refers to analyzing images where humans are depicted in order to extract their identities. Three modules are usually needed for a complete face recognition system: a) a face detector to localize faces in a given image, b) a face alignment method, to ensure the proper alignment of the cropped facial image, and c) a feature extractor that is used to extract discriminative features from each detected face. These features are then fed to a matching algorithm to determine the identity of the face, by computing similarity scores against a database, or used by a classifier to directly predict the identity of persons. Most of the recent literature focuses on training deep learning models by using metric learning objectives, such as angular/cosine-margin-based losses, e.g, CosFace [36], and ArcFace [9]. These methods aim to maximize inter-class variance and minimize intra-class variance in the resulting representation space. The annotations used for training face recognition approaches are similar to any classification task, i.e., each face appearing in an image must be annotated by its name (or a pseudo-anonymized id). Typically, these annotations are then used for evaluation by forming pairs of images that correspond either to the same person or to different persons. The most commonly used architectures in recent face recognition systems employ deep models, such as Residual Neural Networks (ResNets) [14] and MobileNets (MobileFaceNet) [6].

Human pose estimation refers to analyzing images or videos of humans to predict their pose, which in essence entails the localization of their joints, allow-

ing for inferring skeleton information. Pose estimation is a key step for enabling machines to understand peoples’ actions and intentions in various robotics applications [40]. To the best of our knowledge, among the most accurate method currently is the Distribution-Aware coordinate Representation of Keypoint (DARK) method [39], which serves as a plugin for other deep learning-based pose estimation methods, improving the results by modifying the coordinate representation of the heatmap and also by changing the way the ground-truth coordinates are converted to heatmaps. Furthermore, for single pose estimation the soft-gated method presented in [5], which introduces gated skip connections with per-channel learnable parameters to control the data flow for each channel, achieved state-of-the-art results This method also employs a complex hybrid network that combines the HourGlass and U-Net architectures to further improve the accuracy of pose estimation.

For all the deep learning architectures used by person detection, face recognition and pose estimation large-scale datasets are employed during the training process. This is required in order to ensure that the models will be trained effectively, avoiding over-fitting phenomena. The emergence of large-scale datasets, such as COCO [20] for person detection and pose estimation and Microsoft Celeb dataset (MSCeleb) [12] for face recognition, enabled training these models by overcoming these issues. However, significant concerns have been raised recently, mainly regarding maintaining the privacy of the human subjects that appear in these datasets [11]. As a results, creating new datasets for these tasks is costly and requires a significant effort, while, in some cases, even the use of existing datasets might be limited due to ethics concerns.

3 Proposed method

The proposed method for realistic synthetic data generation for human-centric tasks consists of two stages. First, 3D human models are generated and skeleton-related information is extracted for each of them. For this task, we employ natural images, avoiding the need for using hand-crafted 3D models. Then, the proposed pipeline proceeds with the second stage, which concerns data generation through careful blending of real background images and 3D human models. The appropriate constraints are enforced during this step to ensure realistic placement of objects. In the remaining of this Section, we analytically describe each of these two steps.

3.1 3D Human Model Generation and Skeleton Extraction

After a careful review of the relevant literature, PIFu [32] was selected as the most suitable deep learning method for realistic 3D human model generation from single-view images. Although the 3D human models generated from PIFu, achieve a high level of detail they are not articulated, which means that they cannot be directly used for pose estimation tasks. To overcome this limitation, we propose approximating the 3D positions of the joints by employing multiple

renderings of the human models from multiple views. Then, an off-the-shelf pose estimator, the OpenPose estimator [26], is employed to estimate the 3D positions of the joints. Increasing the number of views used for estimating the 3D positions of the joints can significantly increase the accuracy of the generated annotations. The corresponding algorithmic procedure is described on Algorithm 1. Some typical examples of images of human models that were generated using the proposed method are depicted in Figure 1.

Algorithm 1: 3D pose estimation through 2D pose estimation using multi-view images

```

1 Input: Initial camera poses  $\mathbf{D} = [\mathbf{d}_1, \dots, \mathbf{d}_N]^T \in \mathbb{R}^{N \times 6}$ 
2 Output: Estimated 3D joint positions  $\mathbf{J}^{3D} = [\mathbf{j}_1^{3D}, \dots, \mathbf{j}_S^{3D}]^T \in \mathbb{R}^{S \times 3}$ 
3 for  $\mathbf{d}_i \in \mathbf{D}$  do
4   Set the camera in pose  $\mathbf{d}_i$ 
5    $\mathbf{I}_i \leftarrow$  Render the image
6    $\mathbf{J}_i^{2D} \leftarrow$  Perform pose estimation on image  $\mathbf{I}_i$ 
7    $[\mathbf{V}^{near}, \mathbf{V}^{far}]^T \leftarrow$  Unproject  $\mathbf{J}_i^{2D}$  in far and near planes
8    $\mathbf{L}_i^{3D} \leftarrow$  For each joint compute a line in 3D space given  $[\mathbf{V}^{near}, \mathbf{V}^{far}]^T$ 
9 for  $s \leftarrow 1$  to  $S$  do
10   $\mathbf{j}_s^{3D} \leftarrow$  Compute the position of a joint in 3D space using Least-Squares
    approximation on lines  $\mathbf{L}^{3D}$ 
11 return  $\mathbf{J}^{3D}$ 

```

3.2 Data generation through synthesis of real background images and 3D human models

The proposed data generation method is implemented in a virtual 3D environment using OpenGL. It receives as input a set of natural background images and a set of 3D human models. Semantic image segmentation ground truth information must also be provided for the corresponding background images. If this information is not available, then the predictions from a semantic image segmentation method, such as [15], can be used instead. As shown in Fig. 2, the proposed method works by projecting different human 3D model into different, yet valid, places of a background image. More specifically, the virtual environment consist of an image projection plane and a virtual camera that is placed and properly calibrated so as the background images are fully visible when projected to the plane.

During the process of data generation, the method selects a background image and projects it to the plane. Then, potential 2D image locations for human figure placement are computed, based on the semantic image segmentation ground truth. For example, in the case of an outdoor environment, pixels that correspond to classes such as “road” or “pavement”, can be selected for human figure placement. Afterwards, randomly selected 3D human models are placed in various poses between the plane and the camera in a way that their projection coincides with the potential 2D image locations for human figure placement that were computed in the previous steps. The projected scale of the models is



Fig. 1. 3D human models in various views generated by PIFu. Full-body images from the CCP dataset were used as input.

modified, by controlling the distance between the camera and the human figure, considering the rules of perspective projection, so that the images captured from the virtual camera look natural and realistic. In the case where the background images have been retrieved from vehicle cameras, such as those in the CityScapes dataset, the higher the figure is projected on the plane, the larger the distance between the camera and the figure is set. Finally, the ground truth annotations required for pose estimation methods is computed through the projection of the 3D locations of the joints of the populated models. Other annotation types, e.g., bounding boxes, can be trivially computed in a similar way.

4 Experimental Evaluation

4.1 Dataset Generation

The proposed dataset generation pipeline was evaluated by generating a large-scale synthetic dataset. To this end, the 3D human models were generated using a pre-trained model of PIFu and full-body images of people from the Clothing Co-Parsing (CCP) [38] dataset. In particular, we were able to generate 133 human




Fig. 2. Visualization of the data generation process. The depicted background image is derived from the Cityscapes [8] dataset.

models, each one corresponding to a different human. As background images for our environment we selected images from Cityscapes dataset [8]. The Cityscapes dataset contains a diverse set of video sequences recorded in street scenes from 50 different cities. We selected to use a subset of 20,000 images, for which coarse annotations for semantic image segmentation were provided. The number of the images that were actually used was smaller, as images already containing humans were dropped. Overall, the generated dataset contains 50,000 images. The training set contains 40,000 images, while the rest form the test set. The generated data were used for three different tasks, i.e., person detection, face recognition and human pose estimation and evaluated using a wide range of different settings and setups.

4.2 Experimental Evaluation

First, we report results on person detection using two different object detection methods and backbones, all pretrained on the COCO dataset [20]. This first set of experiments aims to evaluate the ability of pretrained models to effectively detect persons that appear in the generated data, examining in this way whether the generated data are realistic enough. The results are reported in Table 1, in terms of precision at 0.5 IoU, for the person class of the COCO validation set as well as for the test set of the generated dataset. All detectors exhibit excellent precision on person detection tasks, validating the aforementioned hypothesis. It should be noted that, apart from CenterNet, the detectors generalize better on this dataset compared to the validation set of COCO.

Table 1. Person Detection Evaluation (AP@0.5 is reported) **Table 2.** Human Pose Estimation Evaluation (AP@0.5 is reported)

Model	COCO	Synthetic	Data	COCO	Synthetic
CenterNet (RS-18)	38.2%	39.0%	R	37.7%	52.7%
SSD-512 (RS-50)	41.9%	86.8%	S	1.2%	93.3%
SSD-512 (VGG16)	41.5%	88.1%	R+S	34.0%	88.7%

“RS” refers to ResNet.

“R” refers to using real data for training (COCO dataset), “S” refers to using the proposed synthetic data and “R+S” to using both of them.

Similar results are also reported for pose estimation using the lightweight OpenPose model [26]. The results are reported in Table 2. For these experiments, instead of evaluating the precision of different pretrained models, we evaluate the effect of using only real data for training (“R”), using only synthetic data (“S”) and combining both of them for the training process (“R+S”). The training process ran for 140,000 iterations for all models following the setup proposed in [26]. First, note that the ability of pretrained pose estimators to generalize on the data is validated (“R” row). Then, training only with synthetic data leads to perfect generalization on the test set of the synthetic data, but fails to generalize on real data. On the other hand, when both synthetic and real data are combined, we observe a significant improvement on the precision on the synthetic test set. A small decrease is observed on the COCO test, but this is probably due to using a lightweight model architecture with relatively low learning capacity.

Finally, we evaluated the generated data for the task of face recognition, where we observed several interesting phenomena. The results are reported in Table 3. The used models were a) pretrained on the MSCeleb (trained for 120 epochs), b) trained both on the combined MSCeleb and the proposed dataset (trained for 120 epochs), c) pretrained on the MSCeleb data and fine-tuned on the combined dataset (trained for 5 epochs). An inverted residual model (IR-50) was used as feature extractor [4]. The evaluation datasets consisted of: a) 5,000 positive and 5,000 negative pairs of images taken from the MSCeleb dataset, b) 5,000 positive and 5,000 negative pairs of images of the test set of the generated synthetic data, c) the LFW dataset [17] and d) the CFP_FF dataset [33].

Several interesting conclusions can be drawn from the results reported in Table 3. First, note that the synthetic data are actually “harder” for a network trained on real data (“R” row). This can be explained if we consider the loss of facial detail that often occur in the 3D models. Quite interesting, this is also the case for real face recognition systems that operate *in-the-wild* with low resolution cameras (e.g., footage from CCTV systems), as well as for robotics systems that are validated using simulators. Then, we observe that training using both the combined real and synthetic set leads to tremendous improvements both for the real and synthetic data, without any significant impact on how the model generalizes on other real datasets (e.g., LFW and CFP_FF). Finally, we observed similar positive results even when we only fine-tuned the model trained for 5

Table 3. Face Recognition Evaluation (face verification precision (%) is reported)

Training Setup	MSCeleb	Synthetic Data	LFW	CFP_FF
R	95.01%	90.63%	99.80%	99.67%
R+S	98.44%	99.75%	99.20%	99.05%
R+S (finetuning)	97.70%	99.47%	99.38%	99.27%

epochs using the combined set (last row), demonstrating that we can obtain similar improvements with full training (“R+S”) only spending a fraction of time for training (5 epochs instead of 120).

5 Conclusion

In this paper we presented a method capable of automatically generating realistic synthetic data with annotations, while leveraging existing images both for the creation of 3D models of humans, as well as for employing realistic backgrounds. The proposed method has minimal cost, since it does not require handcrafted 3D models or simulation environments and it can be used for generating data for a variety of tasks, such as person detection, face recognition, and human pose estimation. The conducted experimental evaluation demonstrated that the generated data are suitable for training DL models in most cases. It is worth noting that for some tasks, such as face recognition, further improvements can be obtained. These results highlight the potential of the proposed data generation method both for generating large-scale datasets for data-scarce domains, as well as for minimizing the distribution shift that is experienced in many robotics applications, e.g., when transferring the models from a simulation environment in real deployment and vice versa. indent

Acknowledgment This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 871449 (OpenDR). This publication reflects the authors views only. The European Commission is not responsible for any use that may be made of the information it contains

References

1. Alhajja, H.A., Mustikovela, S.K., Mescheder, L., Geiger, A., Rother, C.: Augmented reality meets deep learning for car instance segmentation in urban scenes. In: Proc. British Machine Vision Conf. (2017)
2. Alhajja, H.A., Mustikovela, S.K., Mescheder, L.M., Geiger, A., Rother, C.: Augmented reality meets computer vision: Efficient data generation for urban driving scenes. *Int. Journal of Computer Vision* **126**, 961–972 (2018)
3. Alldieck, T., Pons-Moll, G., Theobalt, C., Magnor, M.: Tex2shape: Detailed full human body geometry from a single image. In: Proc. Int. Conf. on Computer Vision (2019)
4. Behrmann, J., Grathwohl, W., Chen, R.T.Q., Duvenaud, D., Jacobsen, J.H.: Invertible residual networks (2019)

5. Bulat, A., Kossaifi, J., Tzimiropoulos, G., Pantic, M.: Toward fast and accurate human pose estimation via soft-gated skip connections. arXiv preprint arXiv:2002.11098 (2020)
6. Chen, S., Liu, Y., Gao, X., Han, Z.: Mobilefacenet: Efficient cnns for accurate real-time face verification on mobile devices (2018)
7. Chen, W., Wang, H., Li, Y., Su, H., Wang, Z., Tu, C., Lischinski, D., Cohen-Or, D., Chen, B.: Synthesizing training images for boosting human 3d pose estimation. In: 2016 Fourth Int. Conf. on 3D Vision (3DV). pp. 479–488 (2016)
8. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proc. Conf. on Computer Vision and Pattern Recognition (2016)
9. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition (2019)
10. Gaidon, A., Wang, Q., Cabon, Y., Vig, E.: Virtual worlds as proxy for multi-object tracking analysis. In: Proc. Conf. on Computer Vision and Pattern Recognition (2016)
11. Girasa, R.: Ethics and privacy i: Facial recognition and robotics. In: Artificial Intelligence as a Disruptive Technology, pp. 105–146. Springer (2020)
12. Guo, Y., Zhang, L., Hu, Y., He, X., Gao, J.: Ms-celeb-1m: A dataset and benchmark for large-scale face recognition. pp. 87–102 (10 2016)
13. Hattori, H., Naresh Boddeti, V., Kitani, K.M., Kanade, T.: Learning scene-specific pedestrian detectors without real data. In: Proc. Conf. on Computer Vision and Pattern Recognition (2015)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition (2015)
15. Hengshuang, Z., Jianping, S., Xiaojuan, Q., Xiaogang, W., Jiaya, J.: Pyramid scene parsing network. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition (2017)
16. Hinterstoisser, S., Pauly, O., Heibel, H., Marek, M., Bokeloh, M.: An annotation saved is an annotation earned: Using fully synthetic training for object detection. In: Proc. Int. Conf. on Computer Vision Workshops (2019)
17. Huang, G.B., Ramesh, M., Berg, T., Learned-Miller, E.: Labeled faces in the wild: A database for studying face recognition in unconstrained environments. Tech. Rep. 07-49, University of Massachusetts, Amherst (October 2007)
18. Kanazawa, A., Black, M.J., Jacobs, D.W., Malik, J.: End-to-end recovery of human shape and pose. In: Proc. Conf. on Computer Vision and Pattern Recognition. pp. 7122–7131 (2018)
19. Lazova, V., Insafutdinov, E., Pons-Moll, G.: 360-degree textures of people in clothing from a single image. In: 2019 Int. Conf. on 3D Vision (3DV). pp. 643–653 (2019)
20. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) Computer Vision – ECCV 2014. pp. 740–755. Springer (2014)
21. Liu, W., Anguelov, D., Erhan, D., Szegedy, C., Reed, S., Fu, C.Y., Berg, A.C.: Ssd: Single shot multibox detector. In: Proc. European Conf. on Computer Vision. pp. 21–37. Springer (2016)
22. Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: SMPL: a skinned multi-person linear model. ACM Transactions on Graphics **34**(6), 248:1–248:16 (2015)

23. Masi, I., Hassner, T., Tran, A.T., Medioni, G.: Rapid synthesis of massive face sets for improved face recognition. In: 2017 12th IEEE Int. Conf. on Automatic Face & Gesture Recognition (FG 2017). pp. 604–611 (2017)
24. Michel, O.: Cyberbotics ltd. webots™: professional mobile robot simulation. *Int. Journal of Advanced Robotic Systems* **1**(1), 5 (2004)
25. Movshovitz-Attias, Y., Kanade, T., Sheikh, Y.: How useful is photo-realistic rendering for visual learning? In: *Computer Vision – ECCV 2016 Workshops*. pp. 202–217. Springer Int. Publishing (2016)
26. Osokin, D.: Real-time 2d multi-person pose estimation on cpu: Lightweight open-pose. In: *arXiv preprint arXiv:1811.12004* (2018)
27. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: *Proc. Conf. on Computer Vision and Pattern Recognition*. pp. 10975–10985 (2019)
28. Pishchulin, L., Jain, A., Wojek, C., Andriluka, M., Thormählen, T., Schiele, B.: Learning people detection models from few training samples. In: *Proc. Conf. on Computer Vision and Pattern Recognition*. pp. 1473–1480 (2011)
29. Richter, S., Vineet, V., Roth, S., Koltun, V.: Playing for data: Ground truth from computer games. In: *Computer Vision – ECCV 2016*. pp. 102–118. Springer (2016)
30. Rodriguez, M., Laptev, I., Sivic, J., Audibert, J.Y.: Density-aware person detection and tracking in crowds. In: *Proc. Int. Conf. on Computer Vision*. pp. 2423–2430 (2011)
31. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: *Proc. Conf. on Computer Vision and Pattern Recognition* (2016)
32. Saito, S., Huang, Z., Natsume, R., Morishima, S., Kanazawa, A., Li, H.: Pifu: Pixel-aligned implicit function for high-resolution clothed human digitization. In: *Proc. Conf. on Computer Vision and Pattern Recognition* (2019)
33. Sengupta, S., Cheng, J., Castillo, C., Patel, V., Chellappa, R., Jacobs, D.: Frontal to profile face verification in the wild. In: *IEEE Conf. on Applications of Computer Vision* (2016)
34. Roberto de Souza, C., Gaidon, A., Cabon, Y., Manuel Lopez, A.: Procedural generation of videos to train deep action recognition networks. In: *Proc. Conf. on Computer Vision and Pattern Recognition* (2017)
35. Tremblay, J., To, T., Birchfield, S.: Falling things: A synthetic dataset for 3d object detection and pose estimation. In: *Proc. Conf. on Computer Vision and Pattern Recognition Workshops (June 2018)*
36. Wang, H., Wang, Y., Zhou, Z., Ji, X., Gong, D., Zhou, J., Li, Z., Liu, W.: Cosface: Large margin cosine loss for deep face recognition (2018)
37. Wang, Q., Gao, J., Lin, W., Yuan, Y.: Learning from synthetic data for crowd counting in the wild. In: *Proc. Conf. on Computer Vision and Pattern Recognition* (2019)
38. Yang, W., Luo, P., Lin, L.: Clothing co-parsing by joint image segmentation and labeling. In: *Proc. Conf. on Computer Vision and Pattern Recognition* (2014)
39. Zhang, F., Zhu, X., Dai, H., Ye, M., Zhu, C.: Distribution-aware coordinate representation for human pose estimation. *arXiv preprint arXiv:1910.06278* (2019)
40. Zimmermann, C., Welschehold, T., Dornhege, C., Burgard, W., Brox, T.: 3d human pose estimation in rgbd images for robotic task learning. In: *Proc. IEEE Int. Conf. on Robotics and Automation*. pp. 1986–1992 (2018)