

Synthesizing Training Images for Boosting Human 3D Pose Estimation

Wenzheng Chen¹ Huan Wang¹ Yangyan Li² Hao Su² Changhe Tu¹
 Dani Lischinski³ Daniel Cohen-Or⁴ Baoquan Chen¹

¹ Shandong University

² Stanford University ³ Hebrew University ⁴ Tel Aviv University

Abstract

Human 3D pose estimation from a single image is a challenging task with numerous applications. Convolutional Neural Networks (CNNs) have recently achieved superior performance on the task of 2D pose estimation from a single image, by training on images with 2D annotations collected by crowd sourcing. This suggests that similar success could be achieved for direct estimation of 3D poses. However, 3D poses are much harder to annotate, and the lack of suitable annotated training images hinders attempts towards end-to-end solutions. To address this issue, we opt to automatically synthesize training images with ground truth pose annotations. We find that pose space coverage and texture diversity are the key ingredients for the effectiveness of synthetic training data. We present a fully automatic, scalable approach that samples the human pose space for guiding the synthesis procedure and extracts clothing textures from real images. We demonstrate that CNNs trained with our synthetic images out-perform those trained with real photos on 3D pose estimation tasks.

1. Introduction

Recovering the 3D geometry of objects in an image is one of the longstanding and most fundamental tasks in computer vision. In this paper, we address a particularly important and challenging instance of this task: the estimation of human 3D pose from a single (monocular) still RGB image of a human subject, which has a multitude of applications [17, 45, 44, 12].

Most of the existing work in human pose estimation produces a set of 2D locations corresponding to the joints of an articulated human skeleton [26, 36]. Additional processing is then required in order to estimate the 3D pose, which is necessary e.g., for fitting a 3D model to the human subject [39, 4, 13, 33, 9, 2]. However, errors are accumulated in this two-stage 3D pose estimation system. Inspired by the




Figure 1. Our training data generation pipeline. The 3D pose space is sampled and the samples are used for deforming SCAPE models. Meanwhile, various clothes textures are mapped onto the human models. The deformed textured models are rendered using a variety of viewpoints and light sources, and finally composited over real image backgrounds.

recent success of training CNNs in an end-to-end fashion, one might expect that direct estimation of 3D poses should be more effective.

In this paper, we directly estimate 3D poses, with a focus on synthesis of effective training data for boosting the performance of deep CNN networks. This task of direct 3D pose estimation is more challenging than the 2D case due to a larger number of parameters to estimate and due to the ambiguities that arise when a 3D articulated figure is projected onto a 2D image plane. Specifically, for 3D pose estimation from monocular images, one needs a large number of human bodies with different genders and fitness levels, seen from a wide variety of viewing angles, featuring a diversity of poses, clothing, and backgrounds.

Therefore, an effective CNN has to be trained by a large number of training examples, which cover well the huge space of appearance variations. Obtaining a sufficiently diverse training set is a major bottleneck. Crowd-sourcing is not a practical option here, because manually annotating a multitude of images with 3D skeletons by human workers is not a feasible task: the annotations must be marked in 3D,

and, furthermore, it is inherently hard for humans to estimate the depth of each joint given only a single 2D image. Massive amounts of 3D poses may be captured by Motion Capture (MoCap) systems, however these systems are not designed to capture the accompanying appearance, and it is difficult to achieve the necessary diversity of the training data.

Recently, synthesized images have been shown to be effective for training CNNs [38, 25]. These works mostly target man-made objects, where rigid transformations are applied to generate pose variation, and very limited work has been done to address texture variation. In this work, we also use synthetic images to address the bottleneck of training data. However, the challenges we face are rather different and much more difficult. Unlike static man-made objects, human bodies are non-rigid, richly articulated, and wear varied clothing. To make CNNs trained on synthesized images effective when applied to real ones, the space of human body types, poses and the diversity of clothing textures must be well represented in the synthetic training set.

To address these issues, we propose to drive the synthesis procedure with real data. We build a statistical model from a large number of 3D poses from a MoCap system, or inferred from human annotated 2D poses, from which we can sample as many body types and poses as needed for training. We further present an automatic method for transferring clothing textures from real product images onto a human body. Without complicated physical simulation as in traditional cloth modeling [5], this data-driven texture synthesis approach is highly scalable, but still retains visual details such as wrinkles and micro-structures. Effectively, we generate 10,556 human models with unique and quality textured clothing. Given a sampled (articulated) pose, we render the textured human body, overlaid on a randomly chosen background image to generate a synthetic image. In our experiments we generated 5,099,405 training images.

We argue that our synthetic training set is richer and more diverse than any of the existing datasets with ground truth 3D annotations, which results in better performance in 3D pose estimation. To demonstrate this, we train several state-of-the-art CNNs with our synthetic images, and evaluate their performance on the human 3D pose estimation task using different datasets, observing consistent and significant improvements over the published state-of-the-art results. Our synthetic training data and the code to generate it will be made publicly available.

2. Related Work

Analyzing human bodies in images and videos has been a research topic for many decades, with particular attention paid to estimation of human body poses [14, 21]. While some earlier works were based on local descriptors [27, 10],

the recent emergence of CNNs has led to significant improvements in body pose estimation from a single image.

Human Pose Datasets FLIC [34], MPI [3], and LSP [19, 20] are the largest available datasets. FLIC has 5003 annotated human bodies, MPI has 2179 fully-marked annotated human bodies, and LSP has 2000 annotated human bodies. Generally speaking, CNNs, even those trained for object classification tasks, can extract high quality image descriptors, which can then be adapted for various other recognition tasks. However, fine-tuning a CNN for a specific task, such as pose estimation, still requires a large number of annotated images. Existing datasets listed above are still too limited in scale and diversity.

CNN-based 2D Pose Estimation Toshev and Szegedy [41] proposed a cascade of CNN-based regressors for predicting 2D joints in a coarse to fine manner. Both Fan et al. [8] and Li et al. [24] proposed to combine body-part detection and 2D joints localization tasks. Gkioxari et al. [11] also explored a multi-task CNN, where an action classifier is combined with a 2D joints detector. Both Tompson et al. [40] and Chen and Yuille [6] proposed to represent the spatial relationships between joints by graphical models, while training a CNN for predicting 2D joint positions. Jain et al. [18] extended FLIC to FLIC-motion by adding optical flow between FLIC frames, and proposed a CNN that takes pairs of images together with the motion features between them as input for predicting human pose from videos.

3D Pose Estimation Since all of the above methods estimate 2D poses, several methods have been proposed for recovering 3D joints from their 2D locations [39, 33, 9, 2]. However, such methods operate only the 2D joint locations, ignoring all other information in the input image, which might contain important cues for inferring the 3D pose. CNN solutions are likely to work better, since they are capable of taking advantage of additional information in the image. We found that CNNs trained with our synthetic images outperform 2D-pose-to-3D-pose methods, even when the latter are provided with the ground truth 2D joint locations, not to mention when they start from automatically estimated 2D poses.

Li and Chan [23] proposed a multi-task CNN for jointly detecting the body parts and regressing the poses, trained using the Human3.6M dataset [15, 16], where the ground truth 3D poses were captured by a MoCap system. Their method achieves high performance on subjects from the same dataset that were put aside as test data. However, we found that the performance of their CNN drops significantly when tested on other datasets, which indicates a strong overfit on the Human3.6M dataset. The reason for this may be

that while there are millions of frames and 3D poses in this dataset, their variety is rather limited.

Human Pose Data Synthesis Several recent works synthesize human body images from 3D models for training algorithms. However, these works are limited in scalability, pose variation or viewpoint variation. Both Pishchulin et al. [31] and Zhou et al. [45] fit 3D models to images, and deform the model to synthesize new images. However, they either request user to supply a good 3D skeleton and segmentation, or need considerable user interaction. Vazquez et al. [42] collect synthesized images with annotations from game engines, thus it is restricted to certain scenes and people. Park et al. [30] use layering to reconstruct images with different pose, leading to imprecise and poor resolution synthetic images. Since none of the above methods can generate large-scale training images, they can hardly satisfy the demand of CNNs.

Other methods exist that recover 3D pose by adding extra information beside 2D annotations. Agarwal et al. [1] predict 3D pose from silhouette. Radwan et al. [32] use both kinematic and orientation constraints to estimate self-occlusion 3D pose. However, the additional constraints may also introduce error and decrease the reliability of the full system.

3. Synthesis of Training Data

Training effective CNNs for 3D pose estimation requires a large number of 3D-annotated images. There are an infinite number of possible combinations of viewing angles, human poses, clothing articles, and backgrounds; thus, a brute-force synthesis approach can generate literally an infinite number of unique training images with human 3D pose annotations by blindly combining the above properties. Clearly, such a brute-force approach will not work, as only a tiny fraction of the synthesized images will resemble real images of humans, which the trained CNN is supposed to work on.

Thus, the generated combinations should be chosen such that (i) the result resembles a real image of a human (we refer to this as the **alignment principle**; and (ii) the synthesized images should be diverse enough to sample well the space of real images of humans (the **variation principle**). Below we describe our synthesis process, which attempts to comply with both of these principles.

Our training data generation approach is illustrated in Figure 1, which consists of sampling the pose space, as described in Section 3.1, and using the results to generate a large collection of articulated 3D human models of different body types with SCAPE [4]. The models are textured with realistic clothing textures extracted from real images, as described in Section 3.2. A sample of synthesized 3D human models in various body types, clothes and poses is




Figure 2. A sample of poses drawn from the learned non-parametric Bayesian network and t-SNE 2D visualization of the high dimensional pose space. Note that the 3D poses inferred from human annotated 2D poses (red) are complementary to MoCap 3D poses (green). New poses (blue) can be sampled from the prior learned from both the MoCap and inferred 3D poses, and have better coverage of the pose space.

shown in Figure 5. Finally, the textured models are rendered and composited over real image backgrounds (Section 3.3).

3.1. Body Pose Space Modelling

Faithful modeling of pose space is essential for our task. The pose distribution of the synthetic images should agree with that of real-world images. It is not practical to design a parametric model for generating poses that conform to the distribution of valid poses. In case there are enough poses available that cover the entire pose space, we can keep sampling poses from this pool, and select a large number of poses whose distribution approximates that of real images well. Unfortunately, we found the poses from the existing datasets only sparsely cover a small portion of the pose space. The existing datasets, e.g. CMU MoCap dataset [7] and Human3.6M [16], are classified by different actions. Even though they cover many common human actions, they can hardly represent the entire pose space. If we simply sample poses from existing MoCap datasets, a large portion of the pose space will not be covered, leaving “holes” in the pose space. CNNs trained with such training data would not perform as well for real images containing poses that happen to fall into these “holes”.

To better cover the pose space, unseen poses should also be generated, rather than only repeating poses from a pose pool. The key challenge is to make sure the generated unseen poses are valid. The idea is to learn the variations of parts that frequently occur together and produce new poses by combining these parts. We learn a sparse and non-parametric Bayesian network from a set of input poses, to factorize the pose representation, and then composite substructures for generating new poses, as proposed in [22].



Figure 3. A sample of clothing images used for transferring texture onto 3D human models. They are from Google Image Search.

Pose samples drawn from the learned Bayesian network exhibit richer variations due to the substructure composition; meanwhile, the poses stay valid as substructures are composited only when appropriate. Note that the degree to which the substructures may be composited is captured by the network and learned from input poses as well. Please refer to [22] and the supplementary material for more detail.

We learn the Bayesian network from both MoCap 3D poses and 3D poses inferred from human annotated 2D poses. There are two large MoCap datasets which contain a large variety of 3D poses: CMU MoCap dataset [7] and Human3.6M [16]. We use the CMU MoCap dataset because it contains more types of actions. The MoCap 3D poses are captured in highly controlled settings with limited number of performers and organized by different actions, which limits the variability of the poses. In contrast, images are taken in much less controlled settings with ubiquitous devices, depicting more people performing a much wider range of activities. Moreover, human 2D pose annotations for images can be obtained from affordable crowd sourcing approaches. Though the 3D pose inferred from these 2D poses might not be accurate, they are more scalable than the MoCap approach, and thus an important source for 3D poses. We use LSP [19, 20] as a 2D pose source because it contains different sports actions, which are more diverse in pose. We use Akhter et al. [2] to recover 3D poses from 2D annotations. The complementary nature of MoCap 3D poses and inferred 3D poses is demonstrated in Figure 2.

Note that the poses sampled from the learned Bayesian network cover the input MoCap and inferred 3D poses well. Moreover, since the prior is learned from both the MoCap and inferred 3D poses, the “interpolation” between the MoCap and inferred poses can be sampled from the learned Bayesian network as well, due to the compositionality.

Each sample of the pose space yields a set of 3D joints. The 3D joints, together with other parameters, such as gender and fitness level, are provided as input to SCAPE [4]



Figure 4. Contour matching for clothing texture transfer. We render 3D human models in a few candidate poses (a), and try to match their contour to those in real clothing images (b). Next, textures from real images are warped according to the best contour matching (c,d,e) and projected onto the corresponding 3D model. Finally, the textures are mirrored (left-right as well as front-back) to cover the entire 3D model (f). Note that the wrinkles on clothes of the 3D body are transferred from product images and are still relatively natural.

for yielding richly varied articulated human models. The fitness levels are supplied based on an empirical distribution, though a distribution learnt from real data might be even better.

3.2. Clothing Texture Transfer

Humans wear a wide variety of clothing. In the real world, clothes are designed for a variety of purposes, may be made of various materials, and exhibit many different colors and textures, resulting in a wide diversity of appearances. Our goal is to synthesize clothed humans whose appearance mimics that seen in real images. However, it is hard to design or learn a parametric model for generating suitable textures. Instead, we propose a fully automatic approach that transfers large amount of clothes textures from images onto human 3D models.

Realistic textures can be easily transferred from an image onto a 3D model, if the model is properly aligned with the corresponding object in the image. However, in general human images, it is extremely hard to find the matching between the clothes and the 3D human model, because clothes are deformed according to the pose of the person who wears it. Even worse, there might be significant foreground occlusion and background clutter. Fortunately, there are many product images of clothes available online, in which clothes are often imaged in canonical poses with little or no foreground occlusion or background clutter (see Figure 3). Our approach is to collect and analyze such images, and then use them to transfer realistic clothing textures onto our 3D human models. The transfer is done by establishing a matching between the contour of the clothing article in the image and the corresponding part of a rendered human model (see Figure 4).

Firstly, we collect a large set of images of sportswear of

various styles using Google and Bing image search. Next, we apply the method of Wang et al. [43] for extracting the foreground clothing from these images, resulting in 2,000 segmented images (1,000 for upper body and 1,000 for lower body). Correspondingly, we split a 3D human model into overlapping upper and lower parts for matching. These two parts are projected onto the clothing images, where they are matched to the clothes. We use continuous dynamic time warping [28] for computing the dense correspondences $M(p)$ between the contours of the two human body parts $\mathcal{P} = \{p_i\}$ and those of the imaged clothing articles $\mathcal{Q} = \{q_j\}$:

$$M = \operatorname{argmin}_{M: \mathcal{P} \rightarrow \mathcal{Q}} \left\{ \sum_p \operatorname{dist}(p, M(p)) \right\}. \quad (1)$$

Once the dense correspondences between the contours are available, we warp the image of the article to fit the projected contour, and the warped image is then used to define the texture for the corresponding portion of the 3D human model. We want the warping to be smooth for minimizing artifacts, and we achieve this by applying MLS image deformation [35]. The resulting textures on the 3D model are mirrored both left-right and front-back for better model coverage.

Since we now have textures for overlapping upper and lower body parts, we can increase the diversity of the clothing textures by randomly adjusting the seam between the textures assigned to the upper and the lower parts. We place 6–12 control points of the seam in the overlapping region, whose vertical positions are randomly adjusted, such that variation is added to the seam between upper and lower body wear.

We found that the clothing images can be matched better, and result in less deformation when the 3D human model is provided in multiple candidate poses, and we pick the pose that results in minimal deformation:

$$\underset{\substack{A \in \mathcal{A} \\ M: \mathcal{C}(I) \rightarrow \mathcal{C}(R_A)}}}{\operatorname{minimize}} \sum_{p \in \mathcal{C}(I)} \operatorname{dist}(p, M(p)), \quad (2)$$

where $\mathcal{C}(I)$ is the contour of clothing image I , A is the pose of 3D human body and $\mathcal{C}(R_A)$ is the contour of the rendering R_A of the 3D human model at pose A . In practice, we picked three upper body candidate poses. Note that rather than pursuing realistic clothing effect, we try to generate human 3D models of high diversity to prevent CNNs from picking up unreliable patterns.

Note that manual texture mapping is an extremely tedious task. We conducted an informal user study among 20 professional CG artists, each of whom was asked to texture map 20 clothing articles onto a human 3D model. It took about 1.5 hour for an artist to texture map one clothing article, on average. Although we found the results of manual



Figure 5. A sample of synthesized 3D human models. Our pose variation is large. Also note that wrinkles and micro-structures are present in these renderings. Besides clothes texture and pose, 3D human model exhibits rich variations in shape due to gender and fitness levels. Such factors are taken into consideration in the 3D human model generation process for enriching the variations.

texture mapping by the artists to be of higher quality, this approach does not scale up due to its labor intensiveness. Our automatic clothing texture transfer method, produces results of somewhat lower quality, but it scales up well.

The head, feet, and hands, which may not be covered by clothes, are texture mapped with a small set of head, shoes and skin textures. Their colors are further perturbed by blending to generate more variations before clothing textures are transferred onto the models. Since the area of these regions is relatively small, their appearance is less important than the clothes, thus we opt for this simple strategy that is also scalable.

3.3. Rendering and Composition

Finally, textured human models in various poses are ready to be rendered and composited into synthetic images for CNN training. Three factors are important in the rendering process: camera viewpoint, lighting, and materials. The camera viewpoint is specified with three parameters: elevation, azimuth, and in-plane rotation. Typically, perturbations are added to the in-plane rotation parameter, by rotating the training images, to augment and generate more training data. Perturbations can also be added to the elevation and azimuth parameters. Starting from the camera viewpoint associated with each 3D pose, we add Gaussian perturbations with standard deviations of 15, 45, and 15 degrees to the elevation, azimuth and in-plane rotation parameters, respectively. Various lighting models, number and energy of light sources are used during the rendering. The color tone of the body skin is also perturbed to represent different types of skin colors. Each rendered image is composited over a randomly chosen sports background image. We collected 796 natural images to serve as background. They were collected manually from image repositories and search engines. As shown in Figure 6, the synthetic images exhibit a wide variety of clothing textures, as well as poses, and present comparable complexity to real images, even though a human observer can easily identify them as synthetic.



Figure 6. A sample of synthetic training images (3 top rows) and real testing images (bottom row). The synthetic images may look fake to a human, but exhibit a rich diversity of poses and appearance for pushing CNNs to learn better.

4. Results and Discussion

We first introduce the datasets for evaluating 3D pose estimation in Section 4.1. Then we demonstrate the effectiveness of our synthetic training data in 3D pose estimation task by feeding it into a number of different CNNs in Section 4.2. We study the performance of our synthetic datasets with some additional experiments in Section 4.3.

4.1. Evaluation Datasets

The lack of images with 3D pose annotations is not only posing a problem for training, but also for evaluating 3D pose estimation methods. Existing datasets with 3D pose annotations, such as Human3.6M [16] and HumanEva [37], have been captured in controlled indoor scenes, and are not as rich in their variability (clothing, lighting, background) as real-world images of humans.

Thus, we have created Human3D+, a new richer dataset of images with 3D annotation, captured in both indoor and outdoor scenes such as room, playground, and park, and containing general actions such as walking, running, playing football, and so on. The dataset consists of 1,574 images, captured with **Perception Neuron** MoCap system by Noitom Ltd. [29]. These images are richer in appearance and background, better representing human images in real-world scenarios, and thus are better suited for evaluating 3D pose estimation methods¹. See our supplementary material for a sample of the images from our evaluation dataset.

4.2. Evaluations on 3D Pose Estimation Task

The method proposed by Li and Chan [23] is the state-of-the-art in human 3D pose estimation from 2D images. They design a network which combines detection and pose estimation in one model, which is less deep than AlexNet and VGG. Their CNN model is directly trained, rather than fine-tuned, on the Human3.6M dataset, and evaluated on

¹The sensors mounting strips are artificial, but necessary for accurate capturing. However, since such strips do not appear in Human3.6M or in our synthetic images, it is not harmful for the comparison fairness.




Figure 7. 3D pose estimation evaluated on Human3.6M (left) and our Human3D+ (right). Various deep learning models (Li and Chan [23], AlexNet and VGG) trained on our data, Human3.6M, or a mixture of them are evaluated. We also compare against the 3D pose estimation methods of Ramakrishna et al. [33] and Akhter et al. [2] which reconstruct a 3D pose from 2D joint locations (left). To compare the generalizability of models trained on Human3.6M and our synthetic data, we evaluate these networks on a new dataset — Human3D+ (right). We observe that models trained by our synthetic data perform significantly better, i.e., our synthetic data facilitates the training of networks with better generalizability, compared against Human3.6M real data.

the same dataset. They train six models corresponding to six actions. The authors kindly provided us all their models trained on Human3.6M.

Since the focus of this paper is on the generation of the training data, and it is not our intention to advocate a new network architecture for pose estimation, we test the effectiveness of our data using “off-the-shelf” image classification CNNs. Specifically, we adapt both AlexNet and VGG for the task of human 3D pose estimation by modifying the last fully connected layer to output the 3D coordinates, appended with an Euclidean loss, and fine-tuning all the fully connected layers to adapt them to the new task.

More specifically, our synthesis process outputs a large set of images $\{\mathcal{I}_i\}$, each associated with a vector $\mathcal{P}_i \in \mathbb{R}^{45}$: the 15 ground truth 3D joint positions (in camera coordinates). The vector \mathcal{P}_i defines the relative spatial relationships between the 3D joints (the human 3D pose), and also the camera viewpoint direction relative to a canonical human coordinate system (e.g, from which side of the human the camera is looking at it). We normalize the joint coordinates such that the sum of skeleton lengths is equal to a constant. We train the CNNs to estimate the 3D pose from a single input image. That is, given an image with a full human subject visible in it, the CNN yields 3D joint positions in camera coordinates. We denote the joint predictions as $\mathcal{P}'_i \in \mathbb{R}^{45}$. We measure the prediction error with an Euclidean loss: $E = \sum_i \|\mathcal{P}_i - \mathcal{P}'_i\|_2$.

We compare the performance of our simple adaptations of AlexNet and VGG with Li and Chan [23]. As aforementioned, Li and Chan train six models on different actions. We test all the testing images by their six models, and select the best one. Both our adaptations and Li and Chan out-




Figure 8. 3D pose estimation performance increases with the size of the synthetic training set (left), and the number of different clothes textures used (right).

put poses that are given in camera view. We first normalize and align the estimated 3D poses towards the ground truth, and then compare the results by plotting the percentage of detected joint points when different error thresholds are used against the ground truth annotations. Note that the normalization and alignment is fair as it does not change the relative spatial relationships between the 3D joints — it only means to bring the output from different methods into a comparable format.

We train these three networks (VGG, AlexNet, and Li and Chan) on three training image datasets (our synthetic images, Human3.6M, and their mixture), and evaluate their performance on two evaluation datasets (Human3.6M and Human3D+). The performance of these variants is plotted in Figure 7.

Several interesting observations can be made from the comparisons in Figure 7. First, training on Human3.6M leads to over-fitting. While the models trained on Human3.6M apparently perform comparably or better than those trained on our synthetic images, when tested on Human3.6M (Figure 7 left), they perform less well when tested on Human3D+ (Figure 7 right), which is more varied than Human3.6M. Another evidence of the over-fitting is that VGG, which is generally considered to have larger learning capability than AlexNet, performs worse than AlexNet when trained on Human3.6M and tested on it (Figure 7 left), since it suffers from stronger over-fitting due to its larger learning capability. Second, it is clear that training with our synthetic data, rather than Human3.6M, leads to better performance on Human3D+ images, which exhibit richer variations (Figure 7 right). This shows a clear advantage of our synthetic images. Third, our synthetic images, when combined together with Human3.6M in the training, consistently improve the performance on both Human3.6M. This is an indication that our synthetic images and real images have complementary characteristics for the training of CNN models. We suspect our synthetic images cover larger pose space and texture variations, while Human3.6M images still have some characteristics, e.g. the realism, that are closer to real images than our synthetic images.

To get a better reference of the performance, we also




Figure 9. Performance of various models evaluated on our synthetic data.

compare against the methods of Ramakrishna et al. [33] and Akhter et al. [2] which reconstruct a 3D pose from 2D joint locations. We found these methods to perform significantly worse, even when provided with the ground truth 2D poses (Figure 7 left)². This is not surprising, as these methods take only the 2D joint positions as input, while ignoring the appearance. In contrast, CNN models effectively consume all the information in the input images.

4.3. Parameter Analysis

Importance of scalability. To investigate how important the number of synthetic images is for the 3D pose estimation performance, we train the same models using different synthetic training set sizes and report their performance in Figure 8 (left). It is clear that increasing the number of synthetic images improves the performance of CNN models.

Importance of texture variability. Similarly, we also study the impact of the number of different clothes textures used for “dressing” the human 3D models in Figure 8 (right). Note that the richness of the clothes textures also plays important role in the overall performance, thus it is critical for the texturing steps to be as automatic as possible. In our case, only a modest amount of user input is required in the clothes images collection step, which actually can be further automated by collecting images from online clothing shops, or by a classifier trained for this task.

²Due to the technical limitation of [29], the ground truth 2D poses are not available in Human3D+, thus this experiment could not be done on Human3D+.

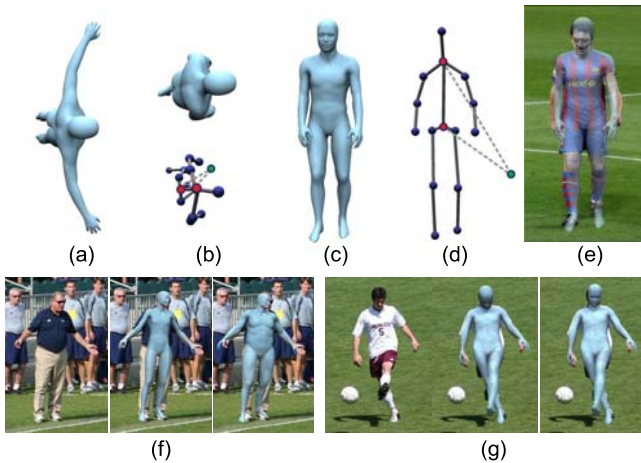



Figure 10. 3D human reconstruction from single image. The SCAPE model in rest pose (a), can be articulated to (b) according to the pose estimated from the image (d). The rigid transformation between (b) and (d) can be computed from corresponding joints to align (b) to the human in the image (c). The reconstruction is visualized in (e).

Evaluation on synthetic images. To better understand the influence of our data for networks’ generalizability, we test the various deep learning models on our synthetic data. The results are summarized in Figure 9. We see that the performance gap between models trained on real images (cyan) and on our synthetic images (green) is much more notable than in Figure 7. It implies that, when the test data and the training data are from different sources, models trained on Human3.6M perform worse than those trained on our synthetic data. This asymmetry in the gaps is also another indication that our synthetic images have more variations than that in Human3.6M — data with less variation is more likely to result in a model that performs well on itself but bad on new data.

4.4. 3D Reconstruction

Human 3D pose estimation from a single image is an important step towards human 3D reconstruction from a single image. The estimated 3D pose can be used to articulate a SCAPE model, as well as align it to the human in the image. The articulated aligned model can already serve as a feasible 3D reconstruction, as shown in Figure 10, and more in Figure 11.

However, more faithful 3D reconstruction requires recovering additional 3D properties from the input image. As shown in Figure 10 (f) and (g), body shape and gaze also play important roles. Similarly to pose, such 3D properties can be hard to annotate, but come free from the synthesis pipeline. We believe our work will encourage more research along these directions.



S

Figure 11. A sample of 3D human reconstruction from single image results, based on a 3D pose estimation model train on our synthetic images.

5. Future Work and Conclusions

Training data for inferring 3D human pose is costly to collect. In our system to synthesize training images from 3D models, the association between the images and the 3D ground truth data is available for free. We found the richness of the clothing textures and the distribution of the poses to be of particular importance. However, constructing a model for realistic clothing synthesis from scratch is a difficult challenge in itself, so we propose instead to sidestep this challenge by transferring clothing textures from real images. We show that the CNNs trained on our synthetic data advance the state-of-the-art performance in the 3D human pose estimation task. We plan to make all of our data and software publicly available to encourage and stimulate further research.

References

- [1] A. Agarwal and B. Triggs. Recovering 3d human pose from monocular images. *TPAMI*, 28(1):44–58, Jan 2006.
- [2] I. Akhter and M. J. Black. Pose-conditioned joint angle limits for 3d human pose reconstruction. In *CVPR*, pages 1446–1455, 2015.
- [3] M. Andriluka, L. Pishchulin, P. Gehler, and B. Schiele. 2d human pose estimation: New benchmark and state of the art analysis. In *CVPR*, June 2014.

- [4] D. Anguelov, P. Srinivasan, D. Koller, S. Thrun, J. Rodgers, and J. Davis. SCAPE: Shape completion and animation of people. *ACM ToG*, 24(3):408–416, July 2005.
- [5] D. Baraff and A. Witkin. Large steps in cloth simulation. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '98, pages 43–54, New York, NY, USA, 1998. ACM.
- [6] X. Chen and A. Yuille. Articulated pose estimation by a graphical model with image dependent pairwise relations. In *NIPS*, pages 1736–1744, 2014.
- [7] CMU. Cmu graphics lab motion capture database. <http://mocap.cs.cmu.edu/>.
- [8] X. Fan, K. Zheng, Y. Lin, and S. Wang. Combining local appearance and holistic view: Dual-source deep neural networks for human pose estimation. In *CVPR15*.
- [9] X. Fan, K. Zheng, Y. Zhou, and S. Wang. Pose locality constrained representation for 3d human pose reconstruction. In *ECCV*, pages 174–188. Springer, 2014.
- [10] V. Ferrari, M. Marin-Jimenez, and A. Zisserman. Progressive search space reduction for human pose estimation. In *CVPR*, pages 1–8, 2008.
- [11] G. Gkioxari, B. Hariharan, R. Girshick, and J. Malik. R-cnn for pose estimation and action detection. *arXiv preprint arXiv:1406.5212*, 2014.
- [12] P. Guan, L. Reiss, D. A. Hirshberg, A. Weiss, and M. J. Black. DRAPE: Dressing any person. *ACM ToG*, 31(4):35:1–35:10, 2012.
- [13] P. Guan, A. Weiss, A. O. Bălan, and M. J. Black. Estimating human shape and pose from a single image. In *ICCV*, pages 1381–1388. IEEE, 2009.
- [14] D. Hogg. Model-based vision: a program to see a walking person. *Image and Vision computing*, 1(1):5–20, 1983.
- [15] C. Ionescu, F. Li, and C. Sminchisescu. Latent structured models for human pose estimation. In *ICCV*, pages 2220–2227, 2011.
- [16] C. Ionescu, D. Papava, V. Olaru, and C. Sminchisescu. Human3.6M: Large scale datasets and predictive methods for 3d human sensing in natural environments. *TPAMI*, 36(7):1325–1339, July 2014.
- [17] A. Jain, T. Thormählen, H.-P. Seidel, and C. Theobalt. MovieReshape: Tracking and reshaping of humans in videos. *ACM ToG*, 29(6):148:1–148:10, 2010.
- [18] A. Jain, J. Tompson, Y. LeCun, and C. Bregler. Modeep: A deep learning framework using motion features for human pose estimation. In *ACCV*. 2014.
- [19] S. Johnson and M. Everingham. Clustered pose and nonlinear appearance models for human pose estimation. In *BMVC*, pages 12.1–12.11, 2010.
- [20] S. Johnson and M. Everingham. Learning effective human pose estimation from inaccurate annotation. In *CVPR*, pages 1465–1472, June 2011.
- [21] H.-J. Lee and Z. Chen. Determination of 3d human body postures from a single view. *Computer Vision, Graphics, and Image Processing*, 30(2):148–168, 1985.
- [22] A. M. Lehrmann, P. V. Gehler, and S. Nowozin. A non-parametric bayesian network prior of human pose. In *ICCV, ICCV '13*, pages 1281–1288, Washington, DC, USA, 2013. IEEE Computer Society.
- [23] S. Li and A. B. Chan. 3d human pose estimation from monocular images with deep convolutional neural network. In *ACCV*, pages 332–347. Springer, 2014.
- [24] S. Li, Z.-Q. Liu, and A. Chan. Heterogeneous multi-task learning for human pose estimation with deep convolutional neural network. In *CVPR14 Workshops*.
- [25] F. Massa, B. Russell, and M. Aubry. Deep exemplar 2d-3d detection by adapting from real to rendered views. *arXiv preprint arXiv:1512.02497*, 2015.
- [26] G. Mori and J. Malik. Estimating human body configurations using shape context matching. In *ECCV, ECCV '02*, pages 666–680, 2002.
- [27] G. Mori and J. Malik. Recovering 3d human body configurations using shape contexts. *TPAMI*, 28(7):1052–1062, 2006.
- [28] M. E. Munich and P. Perona. Continuous dynamic time warping for translation-invariant curve alignment with applications to signature verification. In *ICCV*, volume 1, pages 108–115. IEEE, 1999.
- [29] Noitom. Noitom mocap system. <http://www.noitom.com/>.
- [30] D. Park and D. Ramanan. Articulated pose estimation with tiny synthetic videos. In *CVPR Workshops*, pages 58–66, June 2015.
- [31] L. Pishchulin, A. Jain, M. Andriluka, T. Thormählen, and B. Schiele. Articulated people detection and pose estimation: Reshaping the future. In *CVPR12*.
- [32] I. Radwan, A. Dhall, and R. Goecke. Monocular image 3d human pose estimation under self-occlusion. In *ICCV*, pages 1888–1895, Dec 2013.
- [33] V. Ramakrishna, T. Kanade, and Y. Sheikh. Reconstructing 3D human pose from 2D image landmarks. In *ECCV, ECCV '12*, pages 573–586, 2012.
- [34] B. Sapp and B. Taskar. MODEC: Multimodal decomposable models for human pose estimation. In *CVPR*, pages 3674–3681, June 2013.
- [35] S. Schaefer, T. McPhail, and J. Warren. Image deformation using moving least squares. *ACM Trans. Graph.*, 25(3):533–540, July 2006.
- [36] G. Shakhnarovich, P. Viola, and T. Darrell. Fast pose estimation with parameter-sensitive hashing. In *ICCV, ICCV '03*, pages 750–757, 2003.
- [37] L. Sigal, A. O. Balan, and M. J. Black. Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated human motion. *IJCV*, 87(1):4–27, 2009.
- [38] H. Su, C. R. Qi, Y. Li, and L. J. Guibas. Render for CNN: viewpoint estimation in images using CNNs trained with rendered 3d model views. December 2015.
- [39] C. J. Taylor. Reconstruction of articulated objects from point correspondences in a single uncalibrated image. *Comput. Vis. Image Underst.*, 2000.
- [40] J. J. Tompson, A. Jain, Y. LeCun, and C. Bregler. Joint training of a convolutional network and a graphical model for human pose estimation. In *NIPS*. 2014.
- [41] A. Toshev and C. Szegedy. DeepPose: Human pose estimation via deep neural networks. In *CVPR*, pages 1653–1660, 2014.

- [42] D. Vazquez, A. M. Lopez, J. Marin, D. Ponsa, and D. Geronimo. Virtual and real world adaptation for pedestrian detection. *TPAMI*, 36(4):797–809, April 2014.
- [43] M. Wang, L. Shen, and Y. Yuan. Automatic foreground extraction of clothing images based on grabcut in massive images. In *ICIST*, pages 238–242. IEEE, 2012.
- [44] A. Weiss, D. Hirshberg, and M. J. Black. Home 3D body scans from noisy image and range data. In *ICCV*, pages 1951–1958, 2011.
- [45] S. Zhou, H. Fu, L. Liu, D. Cohen-Or, and X. Han. Parametric reshaping of human bodies in images. *ACM ToG*, 29(4):126:1–126:10, July 2010.