Chained Representation Cycling: Learning to Estimate 3D Human Pose and Shape by Cycling Between Representations

Nadine Rüegg,^{1,2} Christoph Lassner,³ Michael Black,² Konrad Schindler¹

¹ETH Zürich, ²Max Planck Institute for Intelligent Systems, Tübingen, ³Amazon
 {nruegg, classner, black}@tue.mpg.de, schindler@ethz.ch

Abstract

The goal of many computer vision systems is to transform image pixels into 3D representations. Recent popular models use neural networks to regress directly from pixels to 3D object parameters. Such an approach works well when supervision is available, but in problems like human pose and shape estimation, it is difficult to obtain natural images with 3D ground truth. To go one step further, we propose a new architecture that facilitates unsupervised, or lightly supervised, learning. The idea is to break the problem into a series of transformations between increasingly abstract representations. Each step involves a cycle designed to be learnable without annotated training data, and the chain of cycles delivers the final solution. Specifically, we use 2D body part segments as an intermediate representation that contains enough information to be lifted to 3D, and at the same time is simple enough to be learned in an unsupervised way. We demonstrate the method by learning 3D human pose and shape from un-paired and un-annotated images. We also explore varying amounts of paired data and show that cycling greatly alleviates the need for paired data. While we present results for modeling humans, our formulation is general and can be applied to other vision problems.

Introduction

A fundamental task of any vision system, whether natural or artificial, is to learn the relationship between the observed “image” and the 3D scene structure. That is, the system must transform pixels into 3D information and, conversely, take 3D hypotheses about the world and confirm them in the image. This can be thought of as a “cycle” from image pixels to 3D models. Learning such a cyclic mapping is challenging and current methods assume paired training data of image pixels and ground truth 3D. We present an approach that requires little, or no, paired data.

To make the problem concrete, we consider the problem of 3D human pose and shape estimation. This is a challenging test domain since the human body is complex, deformable, and articulated. Additionally, it varies dramatically in appearance due to clothing, hairstyles, etc. If we can find a solution to learn 3D body shape and pose from pixels,

then we posit that this approach should be powerful enough to generalize to many other scenarios.

Consider the example in Fig. 1. Here, we have three representations of the human body: pixels, segment-based, semantic 2D maps, and 3D models. The goal is to learn the mapping from image pixels to 3D models, but we take a path through an intermediate representation and argue that this sequence is easier to learn. The key question is: what properties should intermediate representations in the chain have to facilitate learning? We suggest that a good intermediate representation should be pixel-based, so that the system can learn a per-pixel mapping from the image pixels to the representation. That is, our part segmentations are not so abstract that the mapping becomes too hard to learn in an unsupervised way. Second, the intermediate representation should support 3D inference. That is, it should carry enough information that it is easy to learn the mapping to 3D pose and shape. Previous work has shown that it is possible to learn the mapping that “lifts” human part segmentations to 3D human pose (Lassner et al. 2017; Omran et al. 2018). We go further to show that we can learn a mapping from image pixels via the intermediate representation to 3D, and vice versa, with little or no training data.

Specifically, imagine you are given a set of images, $\mathcal{I}$, showing people with arbitrary pose and appearance. Moreover, you have access to a “high-level” 3D human body model. Rendering photo-realistic training images from a 3D body remains a hard problem, while it is easy to render the body parts as a segmentation mask (see Figure 1). But we do not know the relationship between the set of rendered segmentations, $\mathcal{R}$, and the natural image set $\mathcal{I}$. Can we nevertheless use the two unrelated sets of pictures to learn the mapping from raw pixels in $\mathcal{I}$ to 3D body shape and pose? To do this we develop a novel, unsupervised multi-step cycle model—a representation chain—between three different representations: pixels, segments, and 3D parameters.

The presented approach provides a step towards unsupervised learning of 3D object pose from images. A key advantage of such a direction is that it theoretically allows us to train from unlimited amounts of data, which should improve robustness to real imaging conditions and natural variability.

To test the effectiveness of the approach, we experiment

Figure 1: We present a method to learn a mapping between 2D pixels and deformable 3D models in an unsupervised way. We achieve this by using a chained cycle architecture between 2D images, a semantic encoding as latent appearance vector, background and segments, and the parameters of a 3D body model. Our model can be used in both directions: to automatically recover 3D parameters from 2D data and to generate renderings with varying appearance.

on human pose estimation using both synthetic and real datasets. While not yet as accurate as fully supervised methods, we show that it is feasible to learn with no paired image and 3D data. To our knowledge, this is the first approach to tackle this problem for highly deformable 3D models with such large variability in appearance. We also evaluate the impact of small (*i.e.*, practical) amounts of supervision, which improves performance. Additionally, we show how the model learns a representation that can be used for other tasks like sampling synthetic bodies, transferring clothing, and reposing. Our hope is that the framework of *chained representation cycling* inspires other applications and research into new intermediate representations.

Related Work

Domain adaptation with CycleGAN. A number of recent works employ cyclic adversarial learning to translate between different image domains (Hoffman et al. 2017; Zhu et al. 2017a; Lee et al. 2018; Hoshen and Wolf 2018; Mueller et al. 2018). *CyCADA* (Hoffman et al. 2017), the most relevant representative of that line of work for our task, is trained on pairs of synthetic street scenes and rendered images as well as unpaired real images. (Zhu et al. 2017b) introduce *BicycleGAN*, a method for cross-domain image-to-image translation that includes a latent vector to encode appearance information, as in our work. In contrast to ours, *BicycleGAN* needs supervision by paired data from the two domains. *Augmented CycleGAN* (Almahairi et al. 2018) and *MUNIT* (Huang et al. 2018) also advocate the use of latent vectors to separate appearance from content. We have tested the *MUNIT* architecture for our problem, but found that it did not perform better than a vanilla *CycleGAN*.

3D Human Pose and Shape Estimation. A full review of 3D pose and shape estimation models is beyond the scope of this paper. Recent, representative examples include (Sun et al. 2017; Tome, Russell, and Agapito 2017; Pavlakos et al. 2017; Popa, Zanfir, and Sminchisescu 2017; Pavlakos, Zhou, and Daniilidis 2018; Kanazawa et al. 2018; Omran et al. 2018). Typically, the output is a set of key joints

linked corresponding to a stick figure. These reconstructions often result in unrealistic 3D bone lengths. One solution is to fit a complete 3D mesh instead of just a skeleton. This makes it possible to exploit silhouettes to better constrain the part localization in image space (Lassner et al. 2017), and can also include a prior that favors plausible shapes and poses (Bogo et al. 2016), but requires a supervised method to estimate key locations.

Even when using an elaborate body model, the intermediate 2D representation of the body is predominantly a set of keypoints (Tome, Russell, and Agapito 2017; Martinez et al. 2017; Kanazawa et al. 2018). We deviate from that tradition and instead model the body as a collection of 2D segments corresponding to body parts, like in (Omran et al. 2018).

Another method in the context of unsupervised learning is (Balakrishnan et al. 2018), which performs direct synthesis without explicit 3D reconstruction. A different, more expressive approach is to model directly the volumetric body shape, but also deviate from true human shapes (Varol et al. 2018).

Generative Person Models. One of the first end-to-end trainable generative models of people was presented in (Lassner, Pons-Moll, and Gehler 2017). (Ma et al. 2017) and (Esser, Sutter, and Ommer 2018) focus on transferring person appearance to a new pose, which we also showcase with our model. (Ma et al. 2018) aims for a factorization of foreground, background and pose, similar to our approach, but only tackles the generative part. In contrast to the aforementioned models, (Zanfir et al. 2018) takes a high-fidelity approach to human appearance transfer, combining deep learning with computer graphics models. To our best knowledge, all existing methods for appearance transfer rely on a keypoint detector or part segments, and therefore require either labeled data or multiple images of the same person for training.

Semi-Supervised Approaches. A major difficulty for 3D body fitting is to obtain enough ground truth 3D poses. Several recent works therefore aim to do the 2D-to-3D lifting with little or no 3D supervision. (Kanazawa et al. 2018) cast

the lifting problem as 3D pose regression and train it by minimizing a reprojection loss in 2D and an adversarial loss on the predicted 3D pose. Related is (Kudo et al. 2018), which follows the same general idea, but reprojects the predicted 3D joints into a number of 2D views and places the adversarial loss on the resulting 2D joint configurations. (Rhodin, Salzmann, and Fua 2018) train an encoder-decoder architecture to synthesize novel viewpoints of a person. Their pretraining only needs multi-view data, but no annotations. (Tulsiani, Efros, and Malik 2018) follows the same idea and uses solely multiview consistency as a supervisory signal, for rigid objects.

Recently, (Lorenz et al. 2019) have presented a method for unsupervised, part-based disentangling of object shape and appearance, given video data or images with a plain white background. (Jakab et al. 2018; 2019; Minderer et al. 2019) also propose strategies for unsupervised detection of 2D human keypoints in videos. None of those approaches treats the background explicitly, so training is only possible on videos with almost no camera motion.

Foundations

In this section we briefly introduce two prior methods used as components of our models.

Skinned Multi-Person Linear Model (SMPL). Our method learns a mapping between the parameterization space of a 3D model and 2D images of object instances of the model. As 3D model in our work, we choose SMPL, the “Skinned Multi-Person Linear Model” as proposed in (Loper et al. 2015). This model establishes a pose and shape space for human bodies in 3D space. Concretely, it is parameterized by the pose parameters $\theta \in \mathbb{R}^{72}$ (one 3D axis angle rotation vector for 23 body parts plus a root rotation vector) and the shape parameters $\beta \in \mathbb{R}^{10}$ (ten PCA components to account for variation in body shape). The result of posing and shaping SMPL is a 6890 vertex mesh with 10 faces.

Neural Body Fitting (NBF). NBF describes an approach for supervised 3D human pose estimation (Omran et al. 2018). It uses (different) body part segments as an intermediate representation and maps these segments to a parameterization of SMPL. This mapping is performed using a neural network that takes a $224 \times 224 \times 3$ color coded segmentation map as input and predicts the full SMPL parameters. This module can be treated as a black box for our approach, and we use it to implement the translation between intermediate representation and 3D space, with changes only to input and output encoding and losses.

Approach

We establish a connection between the 2D and 3D representations of objects, and present a novel, unsupervised approach for mapping 3D objects to images and back. We do note that for many object types 3D models exist, or can be created from existing 3D data in an automated manner. This includes objects with limited shape variability and articulation (e.g., cars, furniture), but more recently also common animals (Zuffi et al. 2017). As we aim to handle articulated

pose, including 3D orientation and translation, the projected images can vary strongly in appearance. To still be able to use unsupervised training, we need to reduce the complexity of the task as much as possible. An ideal intermediate representation would describe the image and, at the same time, reduce the complexity of the following, chained step. We introduce part segments as this main, intermediate representation to connect the two domains. With them, we can “factorize” an image into background, a segment-based description of object pose and shape, and residual appearance. With this representation, (1) it is trivial to extract the background; (2) object pose and shape are encoded more richly than by keypoints (Omran et al. 2018) for the next step; and (3) residual appearance reduces to precisely the information needed to “paint” a person with known segmentation. Furthermore, the part segments reside in the same domain as images: the 2D plane. They are spatially related to the image pixels and are therefore well-predictable by a CNN.

Domains

For clarity we refer to the 2D image space as domain $\mathcal{I}$ (image), to the intermediate representation as domain $\mathcal{P}$ (proxy), and to the observed 3D world as domain $\mathcal{T}$ (target). Domain $\mathcal{I}$ consists of images of the object class of interest, here people. Domain $\mathcal{T}$ is represented with a fully parameterized, deformable 3D model of that class. For our scenario we use the SMPL body model (Loper et al. 2015) with pose and body shape parameters. Domain $\mathcal{P}$ factors the image space into a *background*, explicit *part segments* and a *latent appearance code*. Technically, we store this information as (1) a 4-channel image that contains in the first three channels the natural image of the background region overlaid by a colour-coded segment map of the foreground (with arbitrary, well-separated colours) as well as in channel four a binary foreground-background mask and (2) a vector z that encodes appearance.

Cycle $\mathcal{I}$ - $\mathcal{P}$

The cycle $\mathcal{I}$ - $\mathcal{P}$ consists of a *parsing GAN* model, translating images to part segments, background and appearance, and a *generative GAN* model to translate this intermediate representation to images. We base the design of this loop on the CycleGAN architecture (Zhu et al. 2017a).

CycleGANs. are a model family designed such that a mapping from one image domain $\mathcal{A}$ to another image domain $\mathcal{B}$ can be found in an unsupervised way. The mappings $G : \mathcal{A} \rightarrow \mathcal{B}$ and $F : \mathcal{B} \rightarrow \mathcal{A}$ are learned using two types of loss functions: (i) Adversarial losses encourage images from $G(\mathcal{A})$ to be indistinguishable from images sampled from domain $\mathcal{B}$. Similarly, an adversarial loss enforces matching distributions for $F(\mathcal{B})$ and domain $\mathcal{A}$. (ii) Cycle consistency losses enforce $F(G(\mathcal{A})) \approx \mathcal{A}$ and $G(F(\mathcal{B})) \approx \mathcal{B}$. Through the combination of the two losses, it is possible to sidestep annotations altogether: the adversarial losses ensure that “translated” representations match the target domain, and the consistency losses enforce a consistent mapping between domains. In the following, we extend this basic Cy-

Figure 2: *The mapping function from $\mathcal{IM}$ to $\mathcal{PR}$. We use a classical encoding structure, followed by (1) a fully connected layer predicting mean and variance of the residual appearance as well as (2) a set of deconvolutions to produce a part segmentation map in the original resolution.*

Figure 3: *The mapping function from domain $\mathcal{PR}$ to domain $\mathcal{IM}$. The part segment image and the background are split and the part segments are used together with the residual appearance encoding to produce the object image before fusing it with the background.*

cleGAN model for our purpose, to enforce the intended factorization of images from $\mathcal{IM}$ in $\mathcal{PR}$.

A Complete Segmentation Cycle. We develop a cycle between real-world images and the combination of background, part segments and residual appearance. In contrast to traditional image translation, we need multiple components in $\mathcal{PR}$ to capture all relevant content of natural images in $\mathcal{IM}$. Only then we can establish a full translation cycle without loss of information.

Generator from $\mathcal{IM}$ to $\mathcal{PR}$. The structure of the generator that maps an image from domain $\mathcal{IM}$ to domain $\mathcal{PR}$ is shown in Fig. 2. The image is encoded by convolutional layers and processed by a ResNet (He et al. 2016). The output is processed by several fractionally strided convolutions to produce a segmentation and background map at input resolution. In the first three channels, this map contains the color coded semantic segments of the person where applicable and the background colors where not. In the fourth channel, it contains a semantic segmentation with solely foreground and background labels. In a second branch, we use a fully connected layer to predict the mean and variance of the a latent person appearance encoding vector. This parameterization is inspired by variational autoencoders (Kingma and Welling 2013) and enables us to sample person appearance vectors for the cycle part $\mathcal{PR}\text{-}\mathcal{IM}$ at test time. This allows us to fully sample representation for $\mathcal{PR}$ randomly.

Generator from $\mathcal{PR}$ to $\mathcal{IM}$. The cycle part to generate an image from its semantic description receives as input a background image, the part segment image and a vector z . The

segments together with the z vector contain all information needed to generate an image of a human in the desired pose and appearance. As illustrated in Fig. 3, z is tiled and concatenated with a resized version of the part segmentation map. We use a ResNet (He et al. 2016) followed by fractionally strided convolutions to produce an image in the original size. The fusion with the background is guided by the mask from domain $\mathcal{PR}$.

All design choices for building the generators are driven by two needs: to be able to randomly sample instances in both domains and to create a representation in domain $\mathcal{PR}$ that is helpful for the following chain link to domain $\mathcal{TR}$, but still ‘complete’ in the sense that it contains enough information to generate an example for domain $\mathcal{IM}$.

Mitigating Information Hiding. We find that, when translating across domains with strongly differing appearance, models tend to “hide” texture information by reducing its contrast. This happens in subtle ways even when translating between natural image domains.

Whereas information hiding in regular CycleGANs is mainly a cosmetic problem, it is fundamentally problematic in our setting: the network can hide appearance in the part segments instead of using the vector z to encode it. This (1) leads to polluted silhouettes with noisy segment information and (2) the network becomes agnostic to changes in z . The ability to sample different appearance cues is important, because it enables us to generate different images for the cycle $\mathcal{PR}\text{-}\mathcal{IM}\text{-}\mathcal{PR}\text{-}(\mathcal{TR})$. Hence, we take several steps to prevent the network from hiding appearance in the segmentations: first, we suppress residual color variations within each predicted part by flooding it with the assigned part color. Second, in the cycle $\mathcal{IM}\text{-}\mathcal{PR}\text{-}\mathcal{IM}$, both, the originally predicted image and the thresholded and inpainted image are passed to the transformation $\mathcal{PR}\text{-}\mathcal{IM}$.

The Leap to 3D

The second part of the chained model is trained to bridge the domains $\mathcal{PR}$ and $\mathcal{TR}$, *i.e.*, to match 3D object parameters and part segment maps. Whereas the link $\mathcal{TR}\text{-}\mathcal{PR}$ is implemented as rendering of a segment color-coded SMPL mesh, we implement the link $\mathcal{PR}\text{-}\mathcal{TR}$ as an adaptation of neural body fitting (NBF (Omran et al. 2018)), with the part segments in front of a neutral background as input. Global translation, rotation and shape parameters are estimated in addition to one rotation matrix per joint. In order to handle the ambiguity between model size and camera viewpoint, we height-normalize our 3D model templates and use depth as the only way to influence the size of the depicted model.

The $\mathcal{PR}\text{-}\mathcal{TR}$ part of our model is first trained on sampled pairs of 3D parameters and rendered part segment maps. However, the sampled part segmentation maps in $\mathcal{PR}$ are noise-free, as opposed to the ones produced by our encoder $\mathcal{IM}\text{-}\mathcal{PR}$. Still, we can train our full network by using the chain $(\mathcal{TR}\text{-})\mathcal{PR}\text{-}\mathcal{IM}\text{-}\mathcal{PR}\text{-}\mathcal{TR}$ in an unsupervised way and enforce a “reconstruction loss” for the 3D model parameters.

For this extended cycle $\mathcal{PR}\text{-}\mathcal{IM}\text{-}\mathcal{PR}\text{-}\mathcal{TR}$, we use all loss terms for the $\mathcal{IM}\text{-}\mathcal{PR}$ cycle and an additional term for the reconstruction of 3D parameters. The initial z for the transition $\mathcal{PR}\text{-}\mathcal{IM}$ is *not* sampled randomly, but copied from a

different input image picked at random, to ensure as realistic person appearance as possible. Furthermore, we stop the 3D reconstruction loss gradients at domain $\mathcal{I}\mathcal{M}$, such that the network is not encouraged to paint silhouette information into images from domain $\mathcal{I}\mathcal{M}$.

Experiments

In the following, we first present a detailed evaluation on the synthetic *SURREAL* (Varol et al. 2017) dataset and discuss in a second part results on images of real people. Those images come from the *Human3.6M* (Ionescu et al. 2014) and *Penn Action* (Zhang, Zhu, and Derpanis 2013) datasets.

Results on the SURREAL Dataset

SURREAL contains images of people rendered with 772 clothing textures and 193 CAESAR dataset textures (Robinet, Daanen, and Paquet 1999). Corresponding annotations, including pose, depth maps and segmentation masks are provided. To create a realistic scenario, we process the training dataset such that we have the two unpaired sets of images, strictly separated by motion sequence+rotation. From one set we draw images for domain $\mathcal{I}\mathcal{M}$, from the other poses for domain $\mathcal{P}\mathcal{R}$. First, our method is evaluated in terms of semantic body part segmentation and 3D pose error. Second, we provide qualitative fitting results and show examples for appearance transfer and body shape estimation.

Quantitative Results. To our knowledge there are few different papers that present pose estimation results on *SURREAL*, all using full supervision (Tung et al. 2017; Varol et al. 2017; 2018; Madadi, Bertiche, and Escalera 2018). As a connecting baseline we run a fully supervised training of our model. Note, though, as our version of the dataset is split to be unpaired, we use only half as much data for training. Moreover, the chained architecture is not ideal for the fully supervised setting, due to less direct gradient propagation.

We first evaluate the semantic segmentation results of our model, which correspond to the results from the cycle part $\mathcal{I}\mathcal{M}\text{-}\mathcal{P}\mathcal{R}$. Tab. 1 shows results for segmentation into 14 parts as well as 3D pose estimation. (Madadi, Bertiche, and Escalera 2018) and (Varol et al. 2018) list the mean per joint position error (MPJPE) in millimeter (mm) after the positions are subtracted from the root joint. We additionally provide the translation normalized root mean squared error (t-RMSE). Even though we use a smaller training set and the chained model is not designed for supervised learning, its segmentation results are on par with (Varol et al. 2017) and (Varol et al. 2018), pose and shape reconstructions are slightly worse.

1=1.2

Unsupervised and Semi-Supervised Training. We now turn to the unsupervised setting. As an ablation study, we include a semi-supervised variant with small amounts of supervision added. The results can be found in Tab. 2. We used a fully unsupervised trained model for initialization and then ran training using an L2 loss for the part segment maps of every k -th example ($k \in \{500, 200, 100\}$). The 1% data split contains the 0.5% as well as 0.2% data splits completely. For evaluation, we used the metrics:

Segmentation. We provide intersection over union (IoU) scores for the predicted segmentation masks, ignoring the dominant background category. We additionally provide segmentation results after combining segments, resulting in a total of 14, 4 or 1 segment. *4seg* corresponds to the segments $d = \{\text{arms}, \text{legs}, \text{upper_body}, \text{head}\}$, while *1seg* corresponds to foreground-background segmentation.

3D Error (normal). We report the root mean square error (RMSE) between predicted and ground truth 3D SMPL joint locations. Furthermore, we calculate the RMSE with corrected translation (t-RMSE) and corrected translation as well as rotation (tr-RMSE). tr-RMSE is equal to the pose error calculated after a Procrustes matching.

3D Error (best of 4 swaps). To analyze the effect of the well-known left-right swapping ambiguity, we provide this additional score. To calculate it, we swap the left and right arms, respectively legs, in the segmentation; fit the 3D model for each of the four possible combinations; and use the best fit.

Discussion. We find that the fully unsupervised learning of segments as well as 3D fits works surprisingly well, but is often subject to left-right ambiguities. A main source of information for the model in the unsupervised case is the outer silhouette of the person. Since left-right swaps preserve the outer silhouette, it is hard for the model to identify these swapping errors. This becomes obvious when comparing, e.g., t-RMSE (best of 4) unsupervised with the t-RMSE (normal) with 100% supervision: the unsupervised model learns to estimate body pose and shape nearly as well as in the supervised case, but the left-right swaps remain a hurdle. Importantly, we find that adding very little supervision can give just the right hints to overcome this hurdle. When adding even just 1% of supervision, the t-RMSE (normal) is halved. Also, with 1% supervision the segmentation (14 segments) reaches 90% of the fully supervised performance.

Qualitative Results. Our representation bridges all three domains $\mathcal{I}\mathcal{M}$, $\mathcal{P}\mathcal{R}$ and $\mathcal{T}\mathcal{R}$. We show examples for different modes of operation, (1) 3D fits from images (the direction $\mathcal{I}\mathcal{M}\text{-}\mathcal{T}\mathcal{R}$); and (2) appearance transfer (the direction $\mathcal{T}\mathcal{R}\text{-}\mathcal{I}\mathcal{M}$). All results provided in this section are obtained with unsupervised training on the *SURREAL* dataset.

From $\mathcal{I}\mathcal{M}$ to $\mathcal{T}\mathcal{R}$ – 3D Fits and Segments. Randomly selected results of the unsupervised, chained cycle model are shown in Fig. 4. The explanation of the human silhouette is usually successful, however the model is susceptible to left-right swaps that have a strong influence on the 3D fit, as discussed above. Fig. 5 shows 3D predictions for people with varying body shape. The silhouette provides a strong signal to estimate body shape from all directions. This enables the model to predict shape parameters faithfully, even without any supervision.

From $\mathcal{T}\mathcal{R}$ to $\mathcal{I}\mathcal{M}$ – Appearance Transfer. We showcase a different mode of operation for our model and transfer the appearance between people in Fig. 6. Appearance z is predicted for each of the input images on the left. It is transferred to another image as well as multiple random 3D configurations. The results are stable for a fixed vector z , which indicates that the model learns to encode appearance and does not simply hide cues in the part segmentation map. We

Figure 4: *Qualitative results of our unsupervised model on the SURREAL dataset (Varol et al. 2017). Per example, per column: input image, predicted segments, fitted 3D body.* Even though many body fits are correct w.r.t. the segmentation, the model is susceptible to left-right swaps. The overall body shape is typically well-explained.

	Seg. (IoU)	t-RMSE	MPJPE
SURREAL (Varol et al. 2017)	0.69	-	-
BodyNet (Varol et al. 2018)	0.69	-	40.8
SMPLR (Madadi, Bertiche, and Escalera 2018)	-	-	55.8
Ours (fully supervised)	0.69	85.0	69.5
Ours (unsupervised)	0.36	204.0	167.3

Table 1: Quantitative comparison with supervised methods. We only use half of the dataset for supervised training and still manage to match the performance of SURREAL and BodyNet models in segmentation.

		no sup.	supervision			
		-	0.2%	0.5%	1%	100%
segmen-tation	14seg	0.364	0.488	0.563	0.622	0.688
	4seg	0.713	0.791	0.755	0.777	0.779
	1seg	0.887	0.924	0.923	0.932	0.943
3D error (normal)	RMSE	252.90	182.95	161.74	149.72	129.96
	t-RMSE	204.29	141.90	113.79	100.71	84.97
	tr-RMSE	135.59	104.18	93.80	84.31	71.90
3D error (best of 4 swaps)	RMSE	175.81	145.88	142.45	135.47	121.75
	t-RMSE	125.88	104.28	99.69	90.84	81.51
	tr-RMSE	103.18	86.00	83.89	77.03	69.29

Table 2: *Unsupervised and semi-supervised evaluation results.* Ablation study for varying amounts of supervision (0%, 0.2%, 0.5%, 1% and 100%).

Figure 5: *3D pose and shape predictions of our unsupervised model for people with strongly varying body shape.*

are able to transfer appearance to any 3D configuration and obtain good results, for varying poses and body shapes. The ability to generate people in previously unseen poses is important for the cycle $\mathcal{IM}\text{-}\mathcal{PR}$ to work.

Figure 6: *Example results for appearance transfer.* **1st column:** input images. **2nd column:** appearance transfer to a new image. The new image is visible in the top row, the altered appearance below. **Remaining columns:** the top row shows plain SMPL instance, below are results of mapping $\mathcal{TR}\text{-}\mathcal{IM}$ with the extracted appearance.

Results on Real Images

Results on the H36M Dataset. H36M is a dataset recorded in a lab setting with 3D ground truth available through motion capture data. We initialize with our model

	IoU, 14 Seg.	IoU, 4 Seg.	IoU, 1 Seg.	tr-MPJPE
(Ramakrishna, Kanade, and Sheikh 2012)*	-	-	-	157.3
SMPLify (Bogo et al. 2016)*	-	-	-	82.3
NBF (Omran et al. 2018)*	-	-	-	59.9
HMR (Kanazawa et al. 2018)*	-	-	-	56.8
Ours (27 labeled images)	0.292	0.567	0.797	142.2
Ours (66 labeled images)	0.326	0.568	0.800	140.1

Table 3: Results on the H36M dataset (Ionescu et al. 2014). IoU for predicted body part segments (14 parts, 4 parts and fg/bg) and 3D error (mean per joint reprojection error after translation and rotation alignment) for varying amounts of supervision. *These results are computed on the full H36M test set and trained with supervision. We provide them to give an impression of the current supervised state of the art.

Figure 7: Random examples of estimated 3D pose and shape on the H36M dataset (Ionescu et al. 2014). These results are generated by a model that uses partial supervision with only 66 labeled images.

from the unsupervised training on the SURREAL dataset and adapt it to H36M with minimal supervision (only 27 and 66 labeled images) to limit the impact of the static background on our discriminator. During training, we sample poses for domain $\mathcal{PR}$ from the same Gaussian mixture prior as in (Lassner et al. 2017). Background images are taken from H36M directly. The input image crops are generated following the procedure in (Omran et al. 2018). We perform data augmentation by rotating, mirroring and rescaling the images. For the few labeled images, we extend the augmentation by pasting the segmented persons in front of varying backgrounds. We use subject 1 as a held-out test subject and subjects 5, 6, 7 and 8 as training subjects.

Tab. 3 shows quantitative results. We list results from supervised methods for comparison. Notice that they are not strictly comparable, as the mean per joint reprojection error after translation and rotation alignment (tr-MPJPE) is calculated on H36M joints instead of SMPL joints and a different training / test split is used. For our method, left-right swaps are a major obstacle. However, the tr-MPJPE shows that the rough pose of the person is extracted faithfully. Under-represented poses like sitting strongly impact the average error. In Fig. 7, we provide randomly picked qualitative results to give a direct impression of the predictions.

Qualitative Results on In-The-Wild Images. H3.6M is a studio dataset, thus the variety of person appearances and backgrounds is very limited. To show that our model can be used on in-the-wild images, we provide additional qualitative results for less constrained images. We fine-tune the

unsupervised SURREAL model on the Penn Action training set (Zhang, Zhu, and Derpanis 2013). The dataset provides bounding boxes; we use the regions around the people to create background samples for domain $\mathcal{PR}$. 3D pose and shape are sampled from the same distribution as for SURREAL. Even though Penn Action is a video dataset, we only train on individual, unpaired images. Fig. 8 shows example results. This highlights that our model does not over-fit to characteristics of the synthetic SURREAL images, but does learn to represent humans and recognize them in images.

Conclusion

The 2D-3D generation-parsing cycle is ubiquitous in vision systems. In this paper we have investigated a novel model that aims to capture the entire cycle and use its structure to enable unsupervised learning of the underlying relationships. For this purpose, we have developed a chain of cycles that link representations at different levels of abstraction. We demonstrate the feasibility of *representation cycling*, which we think will be applicable to a variety of problems. In this work we have focused on learning the extraction and generation of human bodies and, to our knowledge, are the first to show entirely unsupervised results on this problem. The complexity of the human body, self-occlusion, and varying appearance all make this problem challenging, even with paired training data. So why attack such a hard problem? While getting supervised data for 3D human pose and shape is still possible, it is much harder for other object classes like animals. Our model pushes how far we can go without

Figure 8: *Qualitative results of our unsupervised model on in-the-wild images (Zhang, Zhu, and Derpanis 2013). Per example, per column: input image, predicted segments, fitted 3D body.*

paired data. It does not yet match state-of-the-art results on human pose benchmarks—getting there will likely require an engineering effort that supervised models already have behind them. Nonetheless, we see considerable potential in the longer term. Approaches that learn on unlabeled data enable us to exploit almost unlimited amounts of data and thus have the potential to yield more robust models. Furthermore, an unsupervised approach can directly be applied to new object categories, avoiding repeating labeling efforts.

Acknowledgements. *This research was supported by the Max Planck ETH Center for Learning Systems. MJB has received research gift funds from Intel, NVIDIA, Adobe, Facebook, and Amazon. While MJB is a part-time employee of Amazon, this research was performed solely at MPI.*

References

- Almahairi, A.; Rajeswar, S.; Sordoni, A.; Bachman, P.; and Courville, A. 2018. Augmented CycleGAN: Learning many-to-many mappings from unpaired data. *arXiv preprint arXiv:1802.10151*.
- Balakrishnan, G.; Zhao, A.; Dalca, A. V.; Durand, F.; and Guttag, J. 2018. Synthesizing images of humans in unseen poses. *arXiv preprint arXiv:1804.07739*.
- Bogo, F.; Kanazawa, A.; Lassner, C.; Gehler, P.; Romero, J.; and Black, M. J. 2016. Keep it SMPL: Automatic estimation of 3D human pose and shape from a single image. In *Proc. ECCV*.
- Esser, P.; Sutter, E.; and Ommer, B. 2018. A variational u-net for conditional appearance and shape generation. In *Proc. CVPR*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proc. CVPR*.
- Hoffman, J.; Tzeng, E.; Park, T.; Zhu, J.-Y.; Isola, P.; Saenko, K.; Efros, A. A.; and Darrell, T. 2017. Cycada: Cycle-consistent adversarial domain adaptation. *arXiv preprint arXiv:1711.03213*.
- Hoshen, Y., and Wolf, L. 2018. Nam - unsupervised cross-domain image mapping without cycles or GANs.
- Huang, X.; Liu, M.-Y.; Belongie, S.; and Kautz, J. 2018. Multimodal unsupervised image-to-image translation. *arXiv preprint arXiv:1804.04732*.
- Ionescu, C.; Papava, D.; Olaru, V.; and Sminchisescu, C. 2014. Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *PAMI*.
- Jakab, T.; Gupta, A.; Bilen, H.; and Vedaldi, A. 2018. Unsupervised learning of object landmarks through conditional image generation. In *Adv. NeurIPS*.
- Jakab, T.; Gupta, A.; Bilen, H.; and Vedaldi, A. 2019. Learning human pose from unaligned data through image translation. In *Proc. CVPR*.
- Kanazawa, A.; Black, M. J.; Jacobs, D. W.; and Malik, J. 2018. End-to-end recovery of human shape and pose. In *Proc. CVPR*.
- Kingma, D. P., and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kudo, Y.; Ogaki, K.; Matsui, Y.; and Odagiri, Y. 2018. Unsupervised adversarial learning of 3d human pose from 2d joint locations. *arXiv preprint arXiv:1803.08244*.
- Lassner, C.; Romero, J.; Kiefel, M.; Bogo, F.; Black, M. J.; and Gehler, P. V. 2017. Unite the people: Closing the loop between 3d and 2d human representations. In *Proc. CVPR*.
- Lassner, C.; Pons-Moll, G.; and Gehler, P. V. 2017. A generative model of people in clothing. In *Proc. ICCV*.
- Lee, H.-Y.; Tseng, H.-Y.; Huang, J.-B.; Singh, M. K.; and Yang, M.-H. 2018. Diverse image-to-image translation via disentangled representations. *arXiv preprint arXiv:1808.00948*.
- Loper, M.; Mahmood, N.; Romero, J.; Pons-Moll, G.; and Black, M. J. 2015. Smpl: A skinned multi-person linear model. *ACM TOG*.
- Lorenz, D.; Bereska, L.; Milbich, T.; and Ommer, B. 2019. Unsupervised part-based disentangling of object shape and appearance. *arXiv preprint arXiv:1903.06946*.
- Ma, L.; Jia, X.; Sun, Q.; Schiele, B.; Tuytelaars, T.; and Van Gool, L. 2017. Pose guided person image generation. In *Adv. NeurIPS*.
- Ma, L.; Sun, Q.; Georgoulis, S.; Van Gool, L.; Schiele, B.; and Fritz, M. 2018. Disentangled person image generation. In *Proc. CVPR*.
- Madadi, M.; Bertiche, H.; and Escalera, S. 2018. Smplr: Deep smpl reverse for 3d human pose and shape recovery. *arXiv preprint arXiv:1812.10766*.
- Martinez, J.; Hossain, R.; Romero, J.; and Little, J. J. 2017. A simple yet effective baseline for 3d human pose estimation. In *Proc. ICCV*.
- Minderer, M.; Sun, C.; Villegas, R.; Cole, F.; Murphy, K.; and Lee, H. 2019. Unsupervised learning of object structure and dynamics from videos. *arXiv preprint arXiv:1906.07889*.
- Mueller, F.; Bernard, F.; Sotnychenko, O.; Mehta, D.; Sridhar, S.; Casas, D.; and Theobalt, C. 2018. Generated hands for real-time 3d hand tracking from monocular rgb. In *Proc. CVPR*.
- Omran, M.; Lassner, C.; Pons-Moll, G.; Gehler, P. V.; and Schiele, B. 2018. Neural body fitting: Unifying deep learning and model-based human pose and shape estimation. In *Proc. 3DV*.

Pavlakos, G.; Zhou, X.; Derpanis, K. G.; and Daniilidis, K. 2017. Coarse-to-fine volumetric prediction for single-image 3D human pose. In *Proc. CVPR*.

Pavlakos, G.; Zhou, X.; and Daniilidis, K. 2018. Ordinal depth supervision for 3D human pose estimation. In *Proc. CVPR*.

Popa, A.; Zanfir, M.; and Sminchisescu, C. 2017. Deep Multitask Architecture for Integrated 2D and 3D Human Sensing. In *Proc. CVPR*.

Ramakrishna, V.; Kanade, T.; and Sheikh, Y. 2012. Reconstructing 3d Human Pose from 2d Image Landmarks. *Proc. ECCV*.

Rhodin, H.; Salzmann, M.; and Fua, P. 2018. Unsupervised geometry-aware representation for 3d human pose estimation. *arXiv preprint arXiv:1804.01110*.

Robinette, K. M.; Daanen, H.; and Paquet, E. 1999. The caesar project: a 3-d surface anthropometry survey. In *2nd 3DIM*.

Sun, X.; Shang, J.; Liang, S.; and Wei, Y. 2017. Compositional human pose regression. In *Proc. ICCV*.

Tome, D.; Russell, C.; and Agapito, L. 2017. Lifting from the deep: Convolutional 3d pose estimation from a single image. *Proc. CVPR*.

Tulsiani, S.; Efros, A. A.; and Malik, J. 2018. Multi-view consistency as supervisory signal for learning shape and pose prediction. In *Proc. CVPR*.

Tung, H.-Y.; Tung, H.-W.; Yumer, E.; and Fragkiadaki, K. 2017. Self-supervised learning of motion capture. In *Adv. NeurIPS*.

Varol, G.; Romero, J.; Martin, X.; Mahmood, N.; Black, M. J.; Laptev, I.; and Schmid, C. 2017. Learning from synthetic humans. In *Proc. CVPR*.

Varol, G.; Ceylan, D.; Russell, B.; Yang, J.; Yumer, E.; Laptev, I.; and Schmid, C. 2018. BodyNet: Volumetric inference of 3D human body shapes. In *Proc. ECCV*.

Zanfir, M.; Popa, A.-I.; Zanfir, A.; and Sminchisescu, C. 2018. Human appearance transfer. In *Proc. CVPR*.

Zhang, W.; Zhu, M.; and Derpanis, K. G. 2013. From actemes to action: A strongly-supervised representation for detailed action understanding. In *Proc. CVPR*.

Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017a. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proc. ICCV*.

Zhu, J.-Y.; Zhang, R.; Pathak, D.; Darrell, T.; Efros, A. A.; Wang, O.; and Shechtman, E. 2017b. Toward multimodal image-to-image translation. In *Adv. NeurIPS*.

Zuffi, S.; Kanazawa, A.; Jacobs, D. W.; and Black, M. J. 2017. 3d menagerie: Modeling the 3d shape and pose of animals. In *Proc. CVPR*.