

Unsupervised Part-Based Disentangling of Object Shape and Appearance

Dominik Lorenz Leonard Bereska Timo Milbich Björn Ommer
 Heidelberg Collaboratory for Image Processing / IWR, Heidelberg University

Abstract

Large intra-class variation is the result of changes in multiple object characteristics. Images, however, only show the superposition of different variable factors such as appearance or shape. Therefore, learning to disentangle and represent these different characteristics poses a great challenge, especially in the unsupervised case. Moreover, large object articulation calls for a flexible part-based model. We present an unsupervised approach for disentangling appearance and shape by learning parts consistently over all instances of a category. Our model for learning an object representation is trained by simultaneously exploiting invariance and equivariance constraints between synthetically transformed images. Since no part annotation or prior information on an object class is required, the approach is applicable to arbitrary classes. We evaluate our approach on a wide range of object categories and diverse tasks including pose prediction, disentangled image synthesis, and video-to-video translation. The approach outperforms the state-of-the-art on unsupervised keypoint prediction and compares favorably even against supervised approaches on the task of shape and appearance transfer.

1. Introduction

A grand goal of computer vision is to automatically, without supervision information, learn about the characteristics of an object in the world. Typically, images show the interplay of multiple such factors of variation. We want to disentangle [9, 2, 5, 17, 10] the effects of these different characteristics and imagine, i.e., synthesize, new images where they are altered individually. For instance, after observing a number of different unlabeled instances of an object category, we want to learn their variations in shape (such as pose relative to the viewer and body articulation) and appearance, e.g., texture and color differences in fur/clothing or skin color. Disentangling shape and appearance is particularly challenging because object deformation typically leads to complicated “recoloring” of image pixels

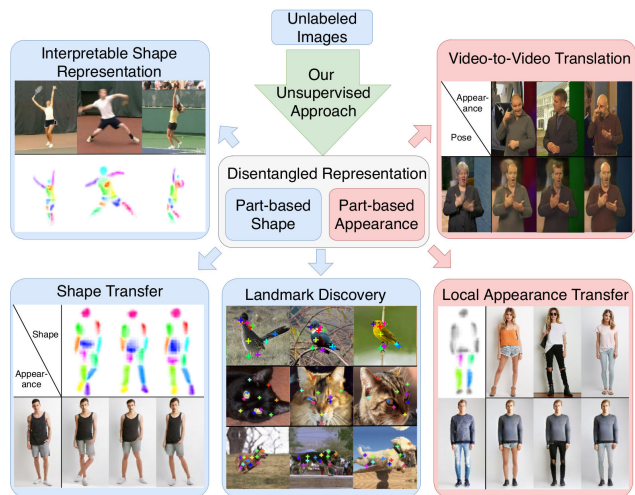


Figure 1: Our unsupervised learning of a disentangled part-based shape and appearance enables numerous tasks ranging from unsupervised pose estimation to image synthesis and retargeting. For more results visit the project page ¹

[40, 12]: moving a limb may change the color of former background pixels into foreground and vice versa.

To address the disentangling problem for shape and appearance, several supervised methods have been proposed lately [29, 28, 7, 12, 41, 1]. By conditioning generative models on a pre-specified shape representation, they are able to successfully explain away appearance. However, they are limited to object categories, for which pose labels are readily available such as human bodies and faces, but they cannot be applied to the vast amounts of unlabelled data of arbitrary object classes.

For unsupervised learning, instead of taking a known shape to capture all non-shape factors, both shape and appearance need to be learned simultaneously. Recently some unsupervised approaches have been proposed to disentangle these factors [40, 52]. However, these works have only shown results for rather rigid objects, like human faces or require several instances of the same person [8].

Object variation can be global, such as difference in

¹<https://compvis.github.io/unsupervised-disentangling/>

viewpoint, but it is oftentimes local (animal tilting its head, person with/without jacket), thus calling for a local, disentangled object representation. The traditional answer are compositions of rigid parts [15, 14, 13]. In the context of recent unsupervised shape learning an instantiation of this are landmarks [45, 58, 21]. In this paper, we propose the first approach to learn a part-based disentangled representation of shape *and* appearance for articulated object classes without supervision and from scratch. In the spirit of analysis-by-synthesis [54], we learn the factors by a generative process. We formulate explicit equivariance and invariance constraints an object representation should fulfill and incorporate them in a fully differentiable autoencoding framework.

Our approach yields significant improvements upon the state-of-the-art in unsupervised object shape learning, evaluated on the task of landmark regression. We compare to competitors on a wide range of diverse datasets both for rigid and articulated objects, with particularly large gains for strong articulations. Furthermore, our disentangled representation of shape and appearance competes favorably even against state-of-the-art supervised results. We also show disentangling results on the task of video-to-video translation, where fine-grained articulation is smoothly and consistently translated on a frame-to-frame level. Lastly, since our representation captures appearance locally, it is also possible to transfer appearance on the level of individual object parts. An overview of the scope of possible applications is given in Fig. 1.

2. Related Work

Disentangling shape and appearance. Factorizing an object representation into shape and appearance is a popular ansatz for representation learning. Recently, a lot of progress has been made in this direction by conditioning generative models on shape information [12, 29, 7, 28, 41, 1]. While most of them explain the object holistically, only few also introduce a factorization into parts [41, 1]. In contrast to these shape-supervised approaches, we learn both shape and appearance without any supervision.

For unsupervised disentangling, several generative frameworks have been proposed [17, 5, 24, 8, 40, 52]. However, these works use holistic models and show results on rather rigid objects and simple datasets, while we explicitly tackle strong articulation with a part-based formulation.

Part-based representation learning. Describing an object as an assembly of parts is a classical paradigm for learning an object representation in computer vision [38, 32, 6, 11]. What constitutes a part, is the defining question in this scheme. Defining parts by visual and semantic features or by geometric shape and its behavior under viewpoint changes and object articulation in general leads to a different partition of the object. Recently, part learning has

been mostly employed for discriminative tasks, such as in [13, 33, 42, 30, 53, 22]. To solve a discriminative task, parts will encode their semantic connection to the object and can ignore the spatial arrangement and articulation. In contrast, our method is driven by an image modelling task. Hence, parts have to encode both spatial structure *and* visual appearance accurately.

Landmark learning. There is an extensive literature on landmarks as compact representations of object structure. Most approaches, however, make use of manual landmark annotations as supervision signal [50, 36, 55, 60, 61, 59, 47, 34, 18, 46, 35, 49, 31, 25, 3].

To tackle the problem without supervision, Thewlis *et al.* [45] proposed enforcing equivariance of landmark locations under artificial transformations of images. The equivariance idea had been formulated in earlier work [23] and has since been extended to learn a dense object-centric coordinate frame [44]. However, enforcing only equivariance encourages consistent landmarks at discriminable object locations, but disregards an explanatory coverage of the object.

Zhang *et al.* [58] addresses this issue: the equivariance task is supplemented by a reconstruction task in an autoencoder framework, which gives visual meaning to the landmarks. However, in contrast to our work, he does not disentangle shape and appearance of the object. Furthermore, his approach relies on a separation constraint in order to avoid the collapse of landmarks. This constraint results in an artificial, rather grid-like layout of landmarks, that does not scale to complex articulations.

Jakab *et al.* [21] proposes conditioning the generation on a landmark representation from another image. A global feature representation of one image is combined with the landmark positions of another image to reconstruct the latter. Instead of considering landmarks which only form a representation for spatial object structure, we factorize an object into local parts, each with its own shape *and* appearance description. Thus, parts are learned which meaningfully capture the variance of an object class in shape as well as in appearance.

Additionally, and in contrast to all these works ([45, 58, 21]) we take the extend of parts into account, when formulating our equivariance constraint. Furthermore, we explicitly address the goal of disentangling shape and appearance on a part-based level by introducing invariance constraints.

3. Approach

Let $x : \Lambda \rightarrow \mathbb{R}$ be an image portraying an object and background clutter. $\Lambda \subset \mathbb{N}^2$ is the space of image coordinates. Now consider an image $x' : \Lambda \rightarrow \mathbb{R}$ showing another instance of the same object category. Despite drastic differences of their image pixels, you can recognize both to be

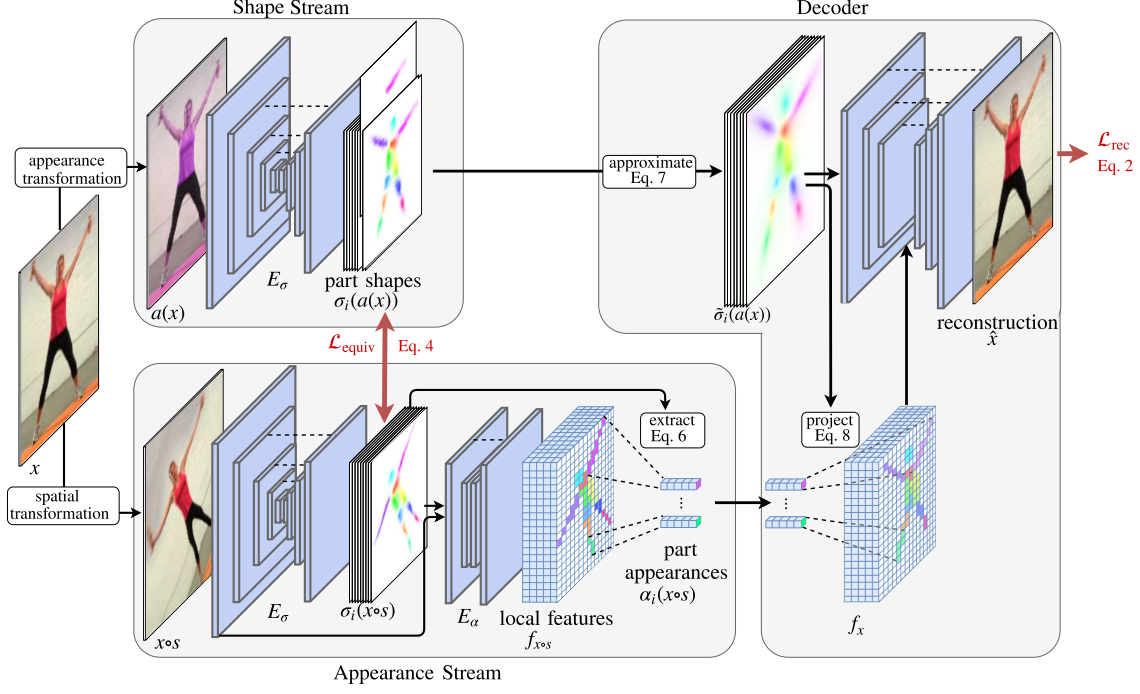


Figure 2: Two-stream autoencoding architecture for unsupervised learning of object shape and appearance.

related. What renders both images similar although no two pixels are identical? What are the characteristic, salient differences? And how can we obtain a representation ϕ that maps images to vectors $\phi(x)$ which retain both, these similarities and also the characteristic differences?

3.1. Part-based Representation

Numerous causes may have led x to be changed into x' (change in articulation, viewpoint, object color or clothing, lighting conditions, etc.). But we can approximate and summarize their effects as a combination of a change in appearance and a change in shape. The effect of a change in object shape on an image x can be expressed in terms of a spatial image transformation $s : \Lambda \rightarrow \Lambda$ acting on the underlying image coordinates, such that the image $x \circ s$ depicts the object with altered shape. Similarly, we denote the effect of a change in object appearance on an image x as an image transformation a such that the image $a(x)$ depicts the object with altered appearance.

Note that many of the image changes are local in nature, affecting only part of the image. For instance, animals may only move an individual body part. Similarly, only part of their appearance may vary, *e.g.*, by switching a shirt but not the pants. This motivates a part-based factorization of the representation, $\phi(x) := (\phi_1(x), \phi_2(x), \dots)^\top$, so that local changes in appearance and shape stay local and do not alter the overall representation. Nevertheless, global changes can also be accounted for by representing them as a compo-

sition of changes in the individual part representations ϕ_i .

3.2. Invariance and Equivariance

Let us now carefully observe differences between images x and x' to derive constraints for the representation ϕ that is to be learned. *i)* Changes in the appearance of an object (*e.g.* in its color or texture), should not impact its shape. *ii)* Similarly, changes in shape (*e.g.* through articulation), should not alter the appearance. Therefore, the representation needs to separate appearance and shape of the object, so that both can vary individually, *i.e.* the representation of a part is disentangled into two components $\phi_i(x) = (\alpha_i(x), \sigma_i(x))$. Part appearance is modeled as an n -dimensional feature vector $\alpha_i(x) \in \mathbb{R}^n$. Whereas part shape is modeled as a part activation map $\sigma_i(x) : \Lambda \rightarrow \mathbb{R}^+$. We visualize these maps as colored images (cf. Fig. 2, Fig. 3), where each color denotes a single part activation map.

The invariance of our representation under changes in object appearance and shape can be summarized by the invariance constraints *i)* $\alpha_i(x \circ s) = \alpha_i(x)$ and *ii)* $\sigma_i(a(x)) = \sigma_i(x)$. In addition, changes in shape should obviously be captured by the shape representation. Thus, for spatial transformations s we obtain the equivariance constraint *iii)* $\sigma_i(x \circ s) = \sigma_i(x) \circ s$. The equivariance constraint simply states that the part activation maps have to consistently track the object part they represent (cf. $\sigma_i(a(x))$ and $\sigma_i(x \circ s)$ in Fig. 2).

3.3. Objective Function for Learning

Learning of the representation ϕ is driven by integrating invariance and equivariance constraints from the previous section into a reconstruction task. The invariance constraints *i*) and *ii*) imply

$$\phi_i(x) = [\alpha_i(x), \sigma_i(x)] \stackrel{!}{=} [\alpha_i(x \circ s), \sigma_i(a(x))]. \quad (1)$$

Let $D([\phi_i(x)]_{i=1,\dots})$ be a reconstruction of the original image x from the encoded part representations $\phi_1(x), \phi_2(x), \dots$ using a decoder D . We seek to reconstruct x and, simultaneously, demand the representation to obey the invariance constraints summarized in (1),

$$\mathcal{L}_{\text{rec}} = \left\| x - D \left([\alpha_i(x \circ s), \sigma_i(a(x))]_{i=1,\dots} \right) \right\|_1. \quad (2)$$

Moreover, the representation of part shape $\sigma_i(x)$ should be equivariant under deformations. However, simply minimizing equivariance on the scale of pixels, i.e.

$$\sum_i \sum_{u \in \Lambda} \left\| \sigma_i(x \circ s)[u] - \sigma_i(x)[s(u)] \right\|, \quad (3)$$

is unstable in practice and favors to the trivial solution of uniform part activations. Therefore, we establish an equivariance loss

$$\begin{aligned} \mathcal{L}_{\text{equiv}} = & \sum_i \lambda_\mu \left\| \mu[\sigma_i(x \circ s)] - \mu[\sigma_i(a(x)) \circ s] \right\|_2 \\ & + \lambda_\Sigma \left\| \Sigma[\sigma_i(x \circ s)] - \Sigma[\sigma_i(a(x)) \circ s] \right\|_1, \end{aligned} \quad (4)$$

where $\mu[\sigma_i(x)]$ and $\Sigma[\sigma_i(x)]$ denote the mean and covariance over coordinates of $\sigma_i(x)/\sum_{u \in \Lambda} \sigma_i(x)[u]$. Note that we have employed invariance *ii*) so that we can use the same shape encoding $\sigma_i(a(x))$ as in (2). The overall training objective of our model is to minimize the reconstruction and equivariance loss,

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \mathcal{L}_{\text{equiv}}. \quad (5)$$

Note that object parts a priori unknown, but in order to reconstruct the object, the representations ϕ_i automatically learn to structure it into meaningful parts which capture the variance in shape and appearance. In particular, we do not need to introduce artificial prior assumptions about the relations between parts, such as the separation constraints employed in [58, 45]. Instead, the local modelling of the part representation (cf. sec. 3.4) as disentangled components of shape and appearance drives our representation to meaningfully structure the object and learn natural relations between parts.

3.4. Unsupervised Learning of Part-based Shape and Appearance

Subsequently, we discuss the network architecture in Fig. 2 for unsupervised learning of an appearance and shape

representation using the reconstruction (2) and equivariance loss (4). Learning considers image pairs $x \circ s$ and $a(x)$. The leading design principle of our architecture is to model the local interplay between part shape and part appearance. In a fully differentiable procedure equivariance of part activation maps is used to extract part appearances from $x \circ s$ and assign them to corresponding image regions in x .

Part Shape. In a shape stream (cf. Fig. 2), an hourglass network [31] E_σ learns to *localize* parts i by means of part activation maps $\sigma_i(a(x)) \in \mathbb{R}^{h \times w}$. The hourglass model nicely suits this task, since it preserves pixel-wise locality and integrates information from multiple scales [31]. Multi-scale context is essential to learn the relations between parts and consistently assign them to an object.

Part Appearance. Let us now localize the parts by detecting $\sigma_i(x \circ s)$ in a spatially transformed image $x \circ s$ using the same network E_σ (cf. Fig. 2 Appearance Stream). To learn representations of part appearance $\alpha_i(x \circ s)$, we first stack all normalized part activations $\sigma_i(x \circ s)/\sum_{u \in \Lambda} \sigma_i(x \circ s)[u]$ and an image encoding, i.e., the output of the first convolution filters of the network E_σ applied to $x \circ s$. A second hourglass network E_α takes this stack as input and maps it onto a localized image appearance encoding $f_{x \circ s} \in \mathbb{R}^{h \times w \times n}$. To obtain local part appearances, we average pool these features at all locations where part i has positive activation

$$\alpha_i(x \circ s) = \frac{\sum_{u \in \Lambda} f_{x \circ s}[u] \sigma_i(x \circ s)[u]}{\sum_{u \in \Lambda} \sigma_i(x \circ s)[u]}. \quad (6)$$

Reconstructing the Original Image. Next we reconstruct x from part appearances $\alpha_i(x \circ s)$ and part activations $\sigma_i(a(x))$ using a U-Net [37] (cf. Fig. 2). The encoder of the U-Net is simply a set of fixed downsampling layers. Only its decoder is learned. We approximate part activations $\sigma_i(a(x))$ by their first two moments,

$$\tilde{\sigma}_i(a(x))[u] = \frac{1}{1 + (u - \mu_i)^T \Sigma_i^{-1} (u - \mu_i)}, \quad (7)$$

where μ_i and Σ_i denote the mean and covariance of the normalized part activation maps $\sigma_i(a(x))/\sum_{u \in \Lambda} \sigma_i(a(x))[u]$. Thus, extra information present in part activations is neglected, forcing the shape encoder E_σ to concentrate on an unambiguous part localization (or else reconstruction loss would increase). The second input to the decoder D in Eq. 2 are part appearances $\alpha_i(x \circ s)$. Note that $\alpha_i(x \circ s)$ are feature vectors devoid of localization. We exploit the fact that the corresponding part activations $\tilde{\sigma}_i(a(x))$ designate the regions of parts i in image x (cf. Fig 2) to project the part appearances onto a localized appearance encoding f_x :

$$f_x[u] = \sum_i \frac{\alpha_i(x \circ s) \cdot \tilde{\sigma}_i(a(x))[u]}{1 + \sum_j \tilde{\sigma}_j(a(x))[u]}. \quad (8)$$

Table 1: Difficulties of datasets: articulation, intra-class variance, background clutter and viewpoint variation.

Dataset	Articul.	Var.	Backgr.	Viewp.
CelebA				
Cat Head		✓		
CUB-200-2011		✓	✓	
Human3.6M	✓			✓
BBC Pose	✓		✓	
Dogs Run	✓	✓	✓	
Penn Action	✓	✓	✓	✓

To reconstruct x , the U-Net can then exploit the local correspondence between f_x , $\tilde{\sigma}_i(a(x))$ and x .

3.5. Implementation Details

For appearance transformation a we apply changes in brightness, contrast, and hue. For image datasets, s are thin plate spline (TPS) transformations. On video datasets, in addition to applying synthetic TPS transformations we randomly sample another frame from the same video sequence which acts as $x \circ s$. Selecting the number of parts is uncritical, since our model is robust for different numbers of parts, Tab. 2. For image synthesis in Sect. 4.3 we train the decoder D with an adversarial loss [20]. We refer to the supplementary for further details on the architecture and the experimental setup.

4. Experiments

In this section we evaluate our unsupervised approach for learning disentangled representation of appearance and shape. Sect. 4.2 evaluates and visualizes the shape representation on the task of unsupervised landmark discovery. Sect. 4.3 investigates the disentangling of our representation. On the task of conditional image generation, we compare our unsupervised shape/appearance disentanglement performance against a state-of-the-art disentanglement method that utilizes groundtruth shape annotations. Moreover, on the task of frame-to-frame video translation we show the robustness of our representation across multiple frames. Additionally, we evaluate the ability of our method to disentangle parts and their local appearance and shape using a part-wise appearance transfer.

4.1. Datasets

CelebA [27] contains ca. 200k celebrity faces of 10k identities. We resize all images to 128×128 and exclude the training and test set of the MAFL subset, following [45]. As [45, 58], we train the regression (to 5 ground truth landmarks) on the MAFL training set (19k images) and test on the MAFL test set (1k images).

Cat Head [56] has nearly 9k images of cat heads. We




Figure 3: Learned shape representation on Penn Action. For visualization, 13 of 16 part activation maps are plotted in one image. (a) Different instances, showing intra-class consistency and (b) video sequence, showing consistency and smoothness under motion, although each frame is processed individually.

use the train-test split of [58] for training (7,747 images) and testing (1,257 images). We regress 5 of the 7 (same as [58]) annotated landmarks. The images are cropped by bounding boxes constructed around the mean of the ground truth landmark coordinates and resized to 128×128 .

CUB-200-2011 [48] comprises ca. 12k images of birds in the wild from 200 bird species. We excluded bird species of seabirds, roughly cropped using the provided landmarks as bounding box information and resized to 128×128 . We aligned the parity with the information about the visibility of the eye landmark. For comparing with [58] we used their published code.

BBC Pose [4] contains videos of sign-language signers with varied appearance in front of a changing background. Like [21] we loosely crop around the signers. The test set includes 1000 frames and the test set signers did not appear in the train set. For evaluation, as [21], we utilized the provided evaluation script, which measures the PCK around $d = 6$ pixels in the original image resolution.

Human3.6M [19] features human activity videos. We adopt the training and evaluation procedure of [58]. For proper comparison to [58] we also removed the background using the off-the-shelf unsupervised background subtraction method provided in the dataset.

Penn Action [57] contains 2326 video sequences of 15 different sports categories. For this experiment we use 6 categories (tennis serve, tennis forehand, baseball pitch, baseball swing, jumping jacks, golf swing). We roughly cropped

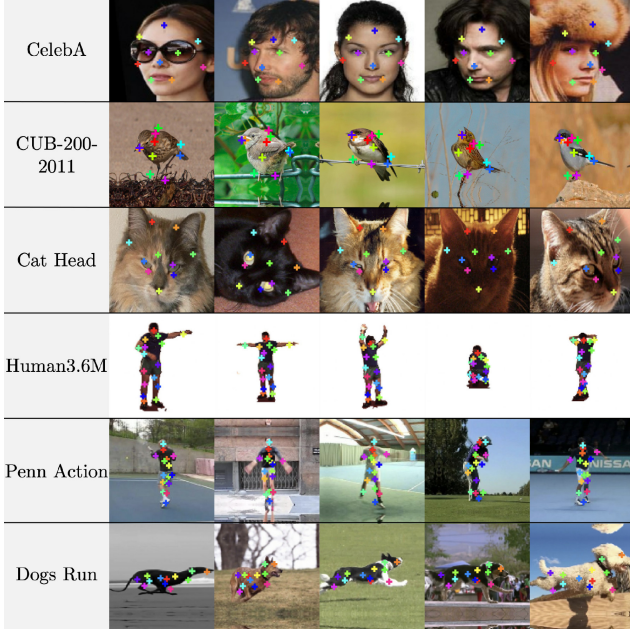


Figure 4: Unsupervised discovery of landmarks on diverse object classes such as human or cat faces and birds and for highly articulated human bodies and running dogs.

Table 2: Error of unsupervised methods for landmark prediction on the Cat Head, MAFL (subset of CelebA), and CUB-200-2011 testing sets. The error is in % of inter-ocular distance for Cat Head and MAFL and in % of edge length of the image for CUB-200-2011.

Dataset	Cat Head		MAFL	CUB
# Landmarks	10	20	10	10
Thewlis [45]	26.76	26.94	6.32	-
Jakab [21]	-	-	4.69	-
Zhang [58]	15.35	14.84	3.46	5.36
Ours	9.88	9.30	3.24	3.91

the images around the person, using the provided bounding boxes, then resized to 128×128 .

Dogs Run is made from dog videos from YouTube totaling in 1250 images under similar conditions as in Penn Action. The dogs are running in one direction in front of varying backgrounds. The 17 different dog breeds exhibit widely varying appearances.

Deep Fashion [26] consists of ca. 53k in-shop clothes images in high-resolution of 256×256 . We selected the images which are showing a full body (all keypoints visible, measured by [3]) and used the provided train-test split. For comparison with Esser *et al.* [12] we used their published code.




Figure 5: Comparing discovered keypoints against [58] on CUB-200-2011. We improve on object coverage and landmark consistency. Note our flexible part placement compared to a rather rigid placement of [58] due to their part separation bias.

4.2. Evaluating Unsupervised Learning of Shape

Fig. 3 visualizes the learned shape representation. To quantitatively evaluate the shape estimation, we measure how well groundtruth landmarks (only during testing) are predicted from it. The part means $\mu[\sigma_i(x)]$ (cf. (4)) serve as our landmark estimates and we measure the error when linearly regressing the human-annotated groundtruth landmarks from our estimates. For this, we follow the protocol of Thewlis *et al.* [45], fixing the network weights after training the model, extracting unsupervised landmarks and training a single linear layer without bias. The performance is quantified on a test set by the mean error and the percentage of correct landmarks (PCK). We extensively evaluate our model on a diverse set of datasets, each with specific challenges. An overview over the challenges implied by each dataset is given in Tab. 1. On all datasets we outperform the state-of-the-art by a significant margin.

Diverse Object Classes. On the object classes of human faces, cat faces, and birds (datasets CelebA, Cat Head, and CUB-200-2011) our model predicts landmarks consistently across different instances, cf. Fig. 4. Tab. 2 compares against the state-of-the-art. Due to different breeds and species the Cat Head, CUB-200-2011 exhibit large variations between instances. Especially on these challenging datasets we outperform competing methods by a large margin. Fig. 5 also provides a direct visual comparison to [58] on CUB-200-2011. It becomes evident that our predicted landmarks track the object much more closely. In contrast, [58] have learned a slightly deformable, but still rather rigid grid. This is due to their separation constraint, which forces landmarks to be mutually distant. We do not need such a problematic bias in our approach, since the localized, part-based representation and reconstruction guides the shape learning and captures the object and its articulations more closely.

Articulated Object Pose. Object articulation makes consistent landmark discovery challenging. Fig. 4 shows that our model exhibits strong landmark consistency under articulation and covers the full human body meaning-

Table 3: Performance of landmark prediction on BBC Pose test set. As upper bound, we also report the performance of supervised methods. The metric is % of points within 6 pixels of groundtruth location.

BBC Pose		Accuracy
supervised	Charles [4]	79.9%
	Pfister [35]	88.0%
unsupervised	Jakab [21]	68.4%
	Ours	74.5%

Table 4: Comparing against supervised, semi-supervised and unsupervised methods for landmark prediction on the Human3.6M test set. The error is in % of the edge length of the image. All methods predict 16 landmarks.

Human3.6M		Error w.r.t. image size
supervised	Newell [31]	2.16
semi-supervised	Zhang [58]	4.14
unsupervised	Thewlis [45]	7.51
	Zhang [58]	4.91
	Ours	2.79

fully. Even fine-grained parts such as the arms are tracked across heavy body articulations, which are frequent in the Human3.6M and Penn Action datasets. Despite further complications such as viewpoint variations or blurred limbs our model can detect landmarks on Penn Action of similar quality as in the more constrained Human3.6M dataset. Additionally, complex background clutter as in BBC Pose and Penn Action, does not hinder finding the object. Experiments on the Dogs Run dataset underlines that even completely dissimilar dog breeds can be related via semantic parts. Tab. 3 and Tab. 4 summarize the quantitative evaluations: we outperform other unsupervised and semi-supervised methods by a large margin on both datasets. On Human3.6M, our approach achieves a large performance gain even compared to methods that utilize optical flow supervision. On BBC Pose, we outperform [21] by 6.1%, reducing the performance gap to supervised methods significantly.

4.3. Disentangling Shape and Appearance

Disentangled representations of object shape and appearance allow to alter both properties individually to synthesize new images. The ability to flexibly control the generator allows, for instance, to change the pose of a person or their clothing. In contrast to previous work [12, 8, 28, 29, 7, 21], we achieve this ability without requiring supervision *and* using a flexible part-based model instead of a holistic representation. This allows to explicitly control the parts of an object that are to be altered. We quantitatively compare



Figure 6: Transferring shape and appearance on Deep Fashion. Without annotation the model estimates shape, 2nd column. Target appearance is extracted from images in top row to synthesize images. Note that we trained without image pairs only using synthetic transformations. All images are from the test set.

against *supervised* state-of-the-art disentangled synthesis of human figures. Also we qualitatively evaluate our model on unsupervised synthesis of still images, video-to-video translation, and local editing for appearance transfer.

Conditional Image Generation. On Deep Fashion [27, 26], a benchmark dataset for supervised disentangling methods, the task is to separate person ID (appearance) from body pose (shape) and then synthesize new images for previously unseen persons from the test set in eight different poses. We randomly sample the target pose and appearance conditioning from the test set. Fig. 6 shows qualitative results. We quantitatively compare against supervised state-of-the-art disentangling [12] by evaluating *i*) invariance of appearance against variation in shape by the re-identification error and *ii*) invariance of shape against variation in appearance by the distance in pose between generated and pose target image.

i) To evaluate appearance we fine-tune an ImageNet-pretrained [39] Inception-Net [43] with a re-identification (ReID) algorithm [51] via a triplet loss [16] to the Deep Fashion training set. On the generated images we evaluate the standard metrics for ReID, mean average precision (mAP) and rank-1, -5, and -10 accuracy in Tab. 5. Although our approach is unsupervised it is competitive compared to the supervised VU-Net [12].

Table 5: Mean average precision (mAP) and rank-n accuracy for person re-identification on synthesized images after performing shape/appearance swap. Input images from Deep Fashion test set. Note [12] is supervised w.r.t. shape.

	mAP	rank-1	rank-5	rank-10
VU-Net [12]	88.7%	87.5%	98.7%	99.5%
Ours	90.3%	89.4%	98.2%	99.2%

Table 6: Percentage of Correct Keypoints (PCK) for pose estimation on shape/appearance swapped generations. α is pixel distance divided by image diagonal. Note that [12] serves as upper bound, as it uses the groundtruth shape estimates.

α	2.5%	5%	7.5%	10%
VU-Net [12]	95.2%	98.4%	98.9%	99.1%
Ours	85.6%	94.2%	96.5%	97.4%



Figure 7: Video-to-video translation on BBC Pose. Top-row: target appearances, left: target pose. Note that even fine details in shape are accurately captured. Visit the project page for the video.

ii) To evaluate shape, we extract keypoints using the pose estimator [3]. Tab. 6 reports the difference between generated and pose target in percentage of correct keypoints (PCK). As would be expected, VU-Net performs better, since it is trained with exactly the keypoints of [3]. Still our approach achieves an impressive PCK without supervision underlining the disentanglement of appearance and shape.

Video-to-Video Translation. To evaluate the robustness of our disentangled representation, we synthesize a video sequence frame-by-frame without temporal consistency constraints. On BBC Pose [4], one video provides a sequence of target poses, another video a sequence of source appearances to then perform retargeting, Fig. 7. Although there is no temporal coupling, the generated sequences are smooth and pose estimation is robust. Secondly, the training on the natural spatial deformations in video data enables the model to encapsulate realistic transitions such as out-of-plane rotation and complex 3D artic-

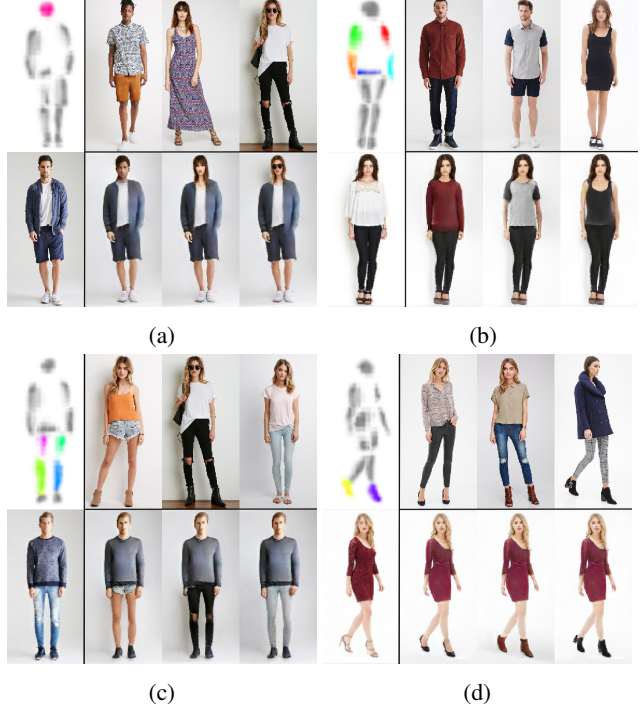


Figure 8: Swapping part appearance on Deep Fashion. Appearances can be exchanged for parts individually and without altering shape. We show part-wise swaps for (a) head (b) torso (c) legs, (d) shoes. All images are from the test set.

ulation of hands and even fingers. Due to the local nature of the part based representation, the model is robust to variations in the background and focuses on the object whilst the background is only roughly reconstructed.

Part Appearance Transfer. The flexible part-based representation allows to explicitly control local appearance. Fig. 8 shows swaps of appearance for shirt, pants, etc. In contrast to holistic representations [12, 21, 28, 29, 7], we can guarantee the transfer to be focused on selected object parts.

5. Conclusion

We have presented an unsupervised approach to learning the compositional part structure of objects by disentangling shape from appearance. We incorporate invariance and equivariance constraints in a generative framework. The model discovers consistent parts without requiring prior assumptions. Experiments show our approach significantly improves upon previous unsupervised methods.

This work has been supported in part by DFG grant OM81/1-1 and a hardware donation from NVIDIA Corporation.

References

- [1] G. Balakrishnan, A. Zhao, A. V. Dalca, F. Durand, and J. V. Guttag. Synthesizing images of humans in unseen poses. *arXiv preprint arXiv:1804.07739*, 2018. 1, 2
- [2] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *TPAMI*, 2013. 1
- [3] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017. 2, 6, 8
- [4] J. Charles, T. Pfister, D. R. Magee, D. C. Hogg, and A. Zisserman. Domain adaptation for upper body pose tracking in signed tv broadcasts. In *BMVC*, 2013. 5, 7, 8
- [5] X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *NIPS*, 2016. 1, 2
- [6] T. F. Cootes, G. J. Edwards, and C. J. Taylor. Active appearance models. In *ECCV*, 1998. 2
- [7] R. de Bem, A. Ghosh, T. Ajanthan, O. Miksik, N. Siddharth, and P. H. S. Torr. Dgpose: Disentangled semi-supervised deep generative models for human body analysis. *arXiv preprint arXiv:1804.06364*, 2018. 1, 2, 7, 8
- [8] E. L. Denton and V. Birodkar. Unsupervised learning of disentangled representations from video. In *NIPS*, 2017. 1, 2, 7
- [9] G. Desjardins, A. Courville, and Y. Bengio. Disentangling factors of variation via generative entangling. *arXiv preprint arXiv:1210.5474*, 2012. 1
- [10] C. Eastwood and C. K. Williams. A framework for the quantitative evaluation of disentangled representations. *ICLR*, 2018. 1
- [11] A. Eigenstetter, M. Takami, and B. Ommer. Randomized max-margin compositions for visual recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3590–3597. IEEE, IEEE, 2014. 2
- [12] P. Esser, E. Sutter, and B. Ommer. A variational u-net for conditional appearance and shape generation. *CVPR*, 2018. 1, 2, 6, 7, 8
- [13] P. F. Felzenszwalb, R. B. Girshick, D. A. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *TPAMI*, 2010. 2
- [14] R. Fergus, P. Perona, and A. Zisserman. Object class recognition by unsupervised scale-invariant learning. In *CVPR*, 2003. 2
- [15] M. A. Fischler and R. A. Elschlager. The representation and matching of pictorial structures. *IEEE Transactions on computers*, 1973. 2
- [16] A. Hermans, L. Beyer, and B. Leibe. In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737*, 2017. 7
- [17] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner. Beta-vae: Learning basic visual concepts with a constrained variational framework. *ICLR*, 2017. 1, 2
- [18] C. Ionescu, F. Li, and C. Sminchisescu. Latent structured models for human pose estimation. In *ICCV*, 2011. 2
- [19] 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*, 2014. 5
- [20] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint*, 2017. 5
- [21] T. Jakab, A. Gupta, H. Bilen, and A. Vedaldi. Conditional image generation for learning the structure of visual objects. *NIPS*, 2018. 2, 5, 6, 7, 8
- [22] M. Lam, B. Mahasseni, and S. Todorovic. Fine-grained recognition as hsnet search for informative image parts. In *CVPR*, 2017. 2
- [23] K. Lenc and A. Vedaldi. Learning covariant feature detectors. In *ECCV Workshops*, 2016. 2
- [24] Z. Li, Y. Tang, and Y. He. Unsupervised disentangled representation learning with analogical relations. In *IJCAI*, 2018. 2
- [25] J. Lim, Y. Yoo, B. Heo, and J. Y. Choi. Pose transforming network: Learning to disentangle human posture in variational auto-encoded latent space. *Pattern Recognit. Lett.*, 2018. 2
- [26] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *CVPR*, 2016. 6, 7
- [27] Z. Liu, P. Luo, X. Wang, and X. Tang. Deep learning face attributes in the wild. In *ICCV*, 2015. 5, 7
- [28] L. Ma, X. Jia, Q. Sun, B. Schiele, T. Tuytelaars, and L. V. Gool. Pose guided person image generation. In *NIPS*, 2017. 1, 2, 7, 8
- [29] L. Ma, Q. Sun, S. Georgoulis, L. V. Gool, B. Schiele, and M. Fritz. Disentangled person image generation. *CVPR*, 2017. 1, 2, 7, 8
- [30] G. Mesnil, A. Bordes, J. Weston, G. Chechik, and Y. Bengio. Learning semantic representations of objects and their parts. *Mach Learn*, 2013. 2
- [31] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. *ECCV*, 2016. 2, 4, 7
- [32] T. D. Nguyen, T. Tran, D. Q. Phung, and S. Venkatesh. Learning parts-based representations with nonnegative restricted boltzmann machine. In *ACML*, 2013. 2
- [33] D. Novotny, D. Larlus, and A. Vedaldi. AnchorNet: A weakly supervised network to learn geometry-sensitive features for semantic matching. In *CVPR*, 2017. 2
- [34] M. Pedersoli, R. Timofte, T. Tuytelaars, and L. J. V. Gool. Using a deformation field model for localizing faces and facial points under weak supervision. In *CVPR*, 2014. 2
- [35] T. Pfister, J. Charles, and A. Zisserman. Flowing convnets for human pose estimation in videos. In *ICCV*, 2015. 2, 7
- [36] R. Ranjan, V. M. Patel, and R. Chellappa. Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *TPAMI*, 2017. 2
- [37] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015. 4
- [38] D. A. Ross and R. S. Zemel. Learning parts-based representations of data. *JMLR*, 2006. 2

- [39] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *ICCV*, 2015. 7
- [40] Z. Shu, M. Sahasrabudhe, R. A. Güler, D. Samaras, N. Paragios, and I. Kokkinos. Deforming autoencoders: Unsupervised disentangling of shape and appearance. In *ECCV*, 2018. 1, 2
- [41] A. Siarohin, E. Sangineto, S. Lathuilière, and N. Sebe. Deformable gans for pose-based human image generation. *CVPR*, 2018. 1, 2
- [42] S. Singh, A. Gupta, and A. A. Efros. Unsupervised discovery of mid-level discriminative patches. In *ECCV*, 2012. 2
- [43] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 7
- [44] J. Thewlis, H. Bilen, and A. Vedaldi. Unsupervised learning of object frames by dense equivariant image labelling. In *NIPS*, 2017. 2
- [45] J. Thewlis, H. Bilen, and A. Vedaldi. Unsupervised learning of object landmarks by factorized spatial embeddings. In *ICCV*, 2017. 2, 4, 5, 6, 7
- [46] A. Toshev and C. Szegedy. Deeppose: Human pose estimation via deep neural networks. In *CVPR*, 2014. 2
- [47] N. Ufer and B. Ommer. Deep semantic feature matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6914–6923, 2017. 2
- [48] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. Technical report, California Institute of Technology, 2011. 5
- [49] S.-E. Wei, V. Ramakrishna, T. Kanade, and Y. Sheikh. Convolutional pose machines. In *CVPR*, 2016. 2
- [50] Y. Wu and Q. Ji. Robust facial landmark detection under significant head poses and occlusion. *CVPR*, 2015. 2
- [51] T. Xiao, S. Li, B. Wang, L. Lin, and X. Wang. Joint detection and identification feature learning for person search. In *CVPR. IEEE*, 2017. 7
- [52] X. Xing, R. Gao, T. Han, S.-C. Zhu, and Y. N. Wu. Deformable generator network: Unsupervised disentanglement of appearance and geometry. *arXiv preprint arXiv:1806.06298*, 2018. 1, 2
- [53] W. Yang, W. Ouyang, H. Li, and X. Wang. End-to-end learning of deformable mixture of parts and deep convolutional neural networks for human pose estimation. In *CVPR*, 2016. 2
- [54] I. Yildirim, T. D. Kulkarni, W. Freiwald, and J. B. Tenenbaum. Efficient analysis-by-synthesis in vision: A computational framework, behavioral tests, and modeling neuronal representations. In *CogSci*, 2015. 2
- [55] X. Yu, F. Zhou, and M. Chandraker. Deep deformation network for object landmark localization. In *ECCV*, 2016. 2
- [56] W. Zhang, J. Sun, and X. Tang. Cat head detection - how to effectively exploit shape and texture features. In *ECCV*, 2008. 5
- [57] W. Zhang, M. Zhu, and K. G. Derpanis. From actemes to action: A strongly-supervised representation for detailed action understanding. In *ICCV*, 2013. 5
- [58] Y. Zhang, Y. Guo, Y. Jin, Y. Luo, Z. He, and H. Lee. Unsupervised discovery of object landmarks as structural representations. In *CVPR*, 2018. 2, 4, 5, 6, 7
- [59] Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Facial landmark detection by deep multi-task learning. In *ECCV*, 2014. 2
- [60] Z. Zhang, P. Luo, C. C. Loy, and X. Tang. Learning deep representation for face alignment with auxiliary attributes. *TPAMI*, 2016. 2
- [61] S. Zhu, C. Li, C. C. Loy, and X. Tang. Face alignment by coarse-to-fine shape searching. In *CVPR*, 2015. 2