

Fashion Landmark Detection in the Wild

Ziwei Liu¹, Sijie Yan¹, Ping Luo^{1,2}(✉), Xiaogang Wang^{1,2}, and Xiaoou Tang^{1,2}

¹ Department of Information Engineering, The Chinese University of Hong Kong, Hong Kong, China

{lz013,siyan,pluo,xtang}@ie.cuhk.edu.hk, xgwang@ee.cuhk.edu.hk

² Shenzhen Key Laboratory for Computer Vision and Pattern Recognition, Shenzhen Institutes of Advanced Technology, CAS, Shenzhen, China

Abstract. Visual fashion analysis has attracted many attentions in the recent years. Previous work represented clothing regions by either bounding boxes or human joints. This work presents fashion landmark detection or fashion alignment, which is to predict the positions of functional key points defined on the fashion items, such as the corners of neckline, hemline, and cuff. To encourage future studies, we introduce a fashion landmark dataset (The dataset is available at <http://mmlab.ie.cuhk.edu.hk/projects/DeepFashion/LandmarkDetection.html>.) with over 120K images, where each image is labeled with eight landmarks. With this dataset, we study fashion alignment by cascading multiple convolutional neural networks in three stages. These stages gradually improve the accuracies of landmark predictions. Extensive experiments demonstrate the effectiveness of the proposed method, as well as its generalization ability to pose estimation. Fashion landmark is also compared to clothing bounding boxes and human joints in two applications, fashion attribute prediction and clothes retrieval, showing that fashion landmark is a more discriminative representation to understand fashion images.

Keywords: Clothes landmark detection · Cascaded deep convolutional neural networks · Attribute prediction · Clothes retrieval

1 Introduction

Visual fashion analysis has drawn lots of attentions recently, due to its wide spectrum of applications such as clothes recognition [1–3], retrieval [3–5], and recommendation [6, 7]. It is a challenging task because of the large variations presented in the clothing items, such as the changes of poses, scales, and appearances. To reduce these variations, existing works tackled the problem by looking for informative regions, *i.e.* detecting the clothes bounding boxes [1, 2] or the human joints [8, 9]. We go beyond the above by studying a more discriminative representation, fashion landmark, which is the key-point located at the functional region of clothes, for example the neckline and the cuff.

The first two authors contribute equally and share first authorship.

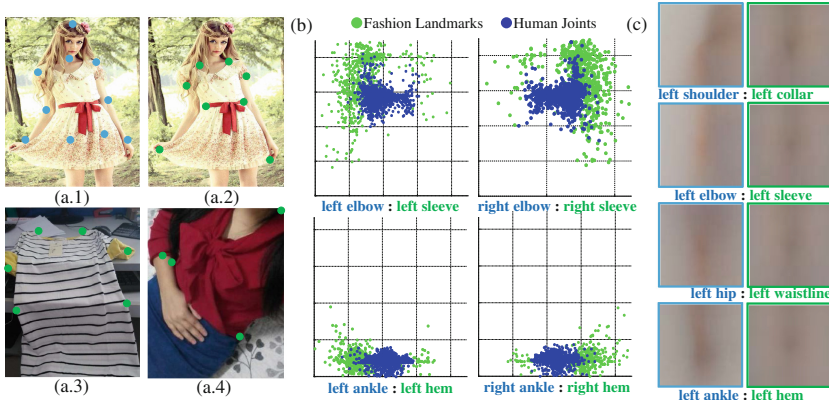


Fig. 1. Comparisons of fashion landmarks and human joints: (a.1) sample annotations for human joints, (a.2) sample annotations for fashion landmarks (a.3–4) typical deformation and scale variations present in clothing items, (b) spatial variances of human joints (in blue) and fashion landmarks (in green) and (c) appearance variances of human joints (in blue) and fashion landmarks (in green). (Color figure online)

This work addresses fashion landmark detection or fashion alignment in the wild. Different from human pose estimation, which detects human joints such as neck and elbows as shown in Fig. 1(a.1), fashion alignment localizes fashion landmarks as shown in (a.2). These landmarks facilitate fashion analysis in the sense that they not only implicitly capture bounding boxes of clothes, but also indicate their functional regions, which can better distinguish design/pattern/category of the clothes. Therefore, features extracted from these landmarks are more discriminative than those extracted from human joints. For example, when search for a dress with ‘V-neck and fringed-hem’, it is more desirable to extract features from collar and hemline.

To fully benchmark the task of fashion landmark detection, we select a large subset of images from the DeepFashion database [3] to constitute a fashion landmark dataset (FLD). These images have large pose and scale variations. With FLD, we show that fashion landmark detection in clothes images is a more challenging task than human joint detection in three aspects. First, clothes undergo *non-rigid deformations* or *scale variations* as shown in Fig. 1(a.3–4), while rigid deformations are usually presented in human joints. Second, fashion landmarks exhibit much larger *spatial variances* than human joints, as illustrated in Fig. 1(b), where we plot the positions of the landmarks and the relative human joints in the test set of the FLD dataset. For instance, the positions of ‘left sleeve’ are more diverse than those of the ‘left elbow’ in both the vertical and horizontal directions. Third, the local regions of fashion landmarks also have larger *appearance variances* than those of human joints. As shown in Fig. 1(c), we average the patches centered at the fashion landmarks and human joints respectively, resulting in several visual comparisons. The patterns of the mean patches of human joints are still recognizable, but those of the mean patches of fashion landmarks are not.

To tackle the above challenges, we propose a deep fashion alignment (DFA) framework, which cascades three deep convolutional networks (CNNs) for landmark estimation. It has three appealing properties. First, to ensure the CNNs have high discriminative powers, unlike existing works [10–14] that only estimated the landmarks’ positions, we train the cascaded CNNs to predict both the landmarks’ positions and the pseudo-labels, which encode the similarities between training samples to boost the estimation accuracy. In each stage of the network cascade, the scheme of pseudo-label is carefully designed to reduce different variations presented in the fashion images. Second, instead of training multiple networks for each body part as previous work did [10, 15], the DFA framework trains CNNs using the full image as input to significantly reduce computations. Third, in the DFA framework, an auto-routing strategy is introduced to partition the challenging and easy samples, such that different samples can be handled by different branches of CNNs.

Extensive experiments demonstrate the effectiveness of the proposed method, as well as its generalization ability to pose estimation. Fashion landmark is also compared to clothing bounding boxes and human joints in two applications, fashion attribute prediction and clothes retrieval, showing that fashion landmark is a more discriminative representation to understand fashion images.

1.1 Related Work

Visual Fashion Understanding. Visual fashion understanding has been a long-pursuing topic due to the many human-centric applications it enables. Recent advances include predicting semantic attributes [3, 8, 9, 16, 17], clothes recognition and retrieval [2–4, 18–20] and fashion trends discovery [6, 7, 21]. To better capture discriminative information in fashion items, previous works have explored the usage of full image [2], general object proposals [2], bounding boxes [9, 17] and even masks [18, 22–24]. However, these representations either lack sufficient discriminative ability or are too expensive to obtain. To overcome these drawbacks, we introduce the problem of clothes alignment in this work, which is a necessary step toward robust fashion recognition.

Human Pose Estimation. We further convert the problem of clothes alignment into fashion landmark estimation. Though there are no prior work in fashion landmarks, approaches from similar fields (*e.g.* human pose estimation [10, 11, 25–28]) serve as good candidates to explore. Recently, deep learning [10–12, 29] has shown great advantages in locating human joints and there are generally two directions here. The first direction [10, 13, 30] utilizes the power of cascading for iterative error correction. DeepPose [10] employs divide-and-conquer strategy and designs a cascaded deep regression framework on part level while Iterative Error Feedback [13] emphasizes more on the target scheduling in each stage. The second direction [11, 12, 31, 32], on the other hand, focuses on the explicit modeling of landmark relationships using graphical models. Chen *et al.* [11] proposed the combination of CNN and structural SVM [33]

to model the tree-like relationships among landmarks while Thompson *et al.* [12] plugged Markov Random Field (MRF) [34] into CNN for joint training. Here, our framework attempts to absorb the advantages of both directions. The cascading and auto-routing mechanisms enable both stage-wise and branch-wise variation reductions while pseudo-labels encode multi-level sample relationships which depict typical global and local landmark configurations.

2 Fashion Landmark Dataset (FLD)

To benchmark fashion landmark detection, we select a subset of images with large pose and scale variations from DeepFashion database [3], to constitute FLD. We label and refine landmark annotations in FLD to make sure each image is correctly labeled with 8 fashion landmarks along with their visibility¹. Overall, FLD contains more than 120K images. Sample images and annotations are shown in Fig. 2(a). To characterize the properties of FLD, we divide the images into five subsets according to the positions and visibility of their ground truth landmarks, including the subsets of normal/medium/large poses and medium/large zoom-ins. The ‘normal’ subset represents images with frontal

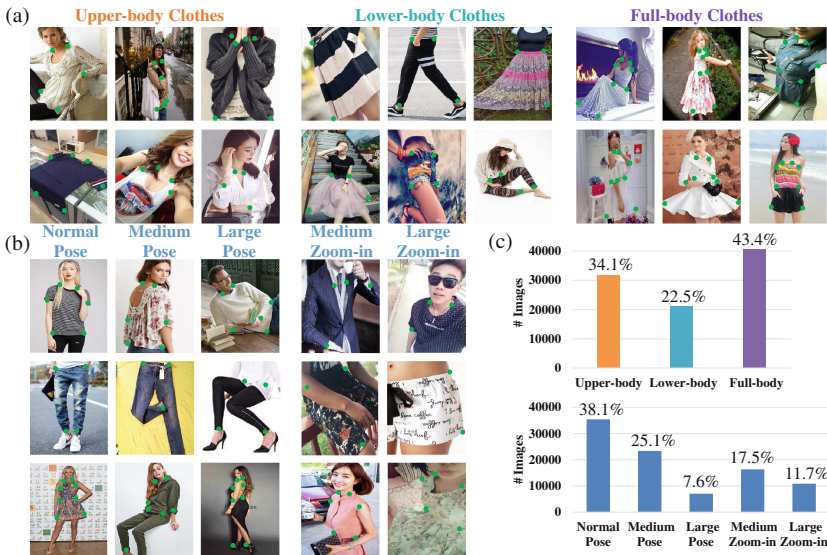


Fig. 2. Illustration of the Fashion Landmark Dataset (FLD): (a) sample images and annotations for different types of clothing items, including upper/lower/full-body clothes, (b) sample images and annotations for different subsets, including normal/medium/large poses and medium/large scales, (c) quantitative data distributions.

¹ Three states of visibility are defined for each landmark, including visible (located inside of the image and visible), invisible (inside of the image but occluded), and truncated/cut-off (outside of the image).

pose and no cut-off landmarks. The subsets of ‘medium’ and ‘large’ poses contain images with side or back views, while the subsets of ‘medium’ and ‘large’ zoom-ins contain clothing items with more than one or three cut-off landmarks, respectively. Sample images of the five subsets are illustrated in Fig. 2(b) and their statistics are demonstrated in Fig. 2(c), which shows that FLD contains substantial percentages of images with large poses/scales.

3 Our Approach

Fashion landmark exhibits large variations in both spatial and appearance domain (see Fig. 1(b, c)). Figure 2(c) further shows that more than 30 % images have large pose and zoom-in variations. To account for these challenges, we propose a deep fashion alignment (DFA) framework as shown in Fig. 3(a), which consists of three stages, where each stage subsequently refines previous predictions. Unlike the existing representative model for human joints prediction, such as DeepPose [10] as shown in (b), which trained multiple networks for all localized parts in each stage, the proposed DFA framework that functions on full image is able to achieve superior performance with much lower computations.

Framework Overview. As shown in Fig. 3(a), DFA has three stages. In each stage, VGG-16 [35] is employed as network architecture. In the first stage, DFA takes the raw image I as input and predicts rough landmark positions, denoted as \hat{l}^1 , as well as pseudo-labels, denoted as \hat{f}^1 , which represent landmark configurations such as clothing categories and poses. In the second stage, both the input image I and the predictions of stage-1, \hat{l}^1 , are fed in. The whole network is

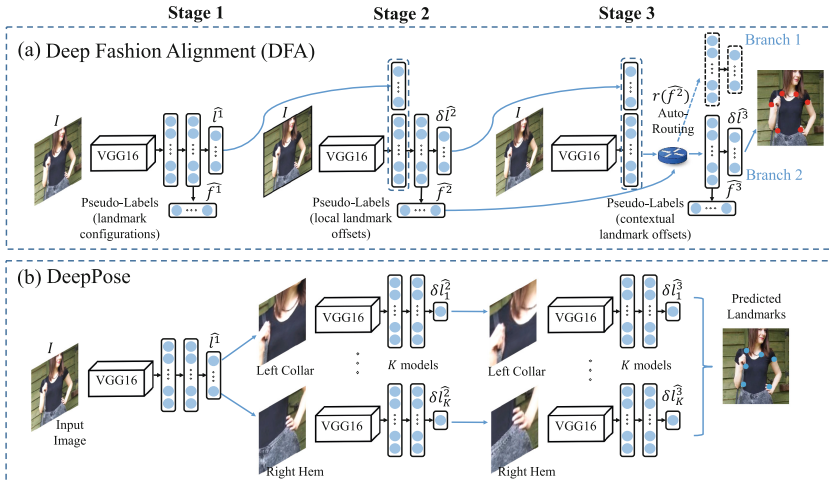


Fig. 3. Pipeline of (a) Deep Fashion Alignment (DFA) and (b) DeepPose [10]. DFA leverages pseudo-labels and auto-routing mechanism to reduce the large variations presented in fashion images with much less computational costs.

required to predict landmark offsets, signified as $\delta\hat{l}^2$, and pseudo-labels f^2 that represent local landmark offsets. The landmark prediction of stage-2 is computed as $l^2 = l^1 + \delta\hat{l}^2$. The third stage has two CNNs as two branches, which have identical input and output. Similar to the second stage, each CNN employs image I as input and learns to estimate landmark offsets $\delta\hat{l}^3$ and pseudo-labels f^3 , which contains information about contextual landmark offsets. In stage-3, each image is passed through one of these two branches. The selection of branch is determined by the predicted pseudo-labels f^2 in stage-2. The final prediction is computed as $l^3 = l^2 + \delta\hat{l}^3$.

Network Cascade. Cascade [36] has been proven an effective technique for reducing variations sequentially in pose estimation [10]. Here, we build the DFA system by cascading three CNNs. The first CNN directly regresses landmark positions and visibility from input images, aided by pseudo-labels in the space of landmark configuration. These pseudo-labels are achieved by clustering absolute landmark positions, indicating different clothing categories and poses, as shown in Fig. 4 stage 1. For example, cluster 1 and 4 represents ‘short T-shirt in front view’ and ‘long coat in side view’, respectively. The second CNN takes both the input image and the predictions from stage-1 as input, and estimates the offsets that should be made to correct the predictions of the first stage. In this case, we are learning in the space of landmark offset. Thus, the pseudo-labels generated here represent typical error patterns and their magnitudes, as shown in Fig. 4 stage 2. For instance, cluster 1 represents the corrections should be made along upside and downside, while cluster 2 suggests left/right-direction corrections. In stage-3, we partition data into two subsets, according to the error patterns predicted in the second stage as shown in Fig. 4 stage 3, where branch one deals with ‘easy’ samples such as frontal T-shirt with different sleeve length, while branch two accounts for ‘hard’ samples such as selfie pose (cluster 1), back view (cluster 2) and large zoom-in (cluster 3).

Pseudo-Label. In DFA, each training sample is associated with a pseudo-label that reflects its relationship to the other samples. Let the ground truth positions of fashion landmarks denote as l , where l_i specifies the pixel location of landmark i . The pseudo-label, $f(k) \in \mathbb{R}^{k \times 1}, k = 1 \dots K$, of each sample i is a K -dim vector and can be calculated as




Fig. 4. Visualization of pseudo-labels obtained for each stage. Pseudo-labels in stage-1 indicate clothing categories and poses; pseudo-labels in stage-2 represent typical error patterns; pseudo-labels in stage-3 partition within ‘easy’ and ‘hard’ samples.

$$f(k) = \exp\left(-\frac{\text{dist}(i, C)}{T}\right), \quad (1)$$

where $\text{dist}(\cdot)$ is a distance measure. C_k is a set of k cluster centers, obtaining by k-means algorithm on the spaces of landmark coordinates (stage-1) or offsets (stage-2 and 3). T is the temperature parameter to soften pseudo-labels. We adopt $K = 20$ for all three stages.

Here we explain the pseudo-label in each of the three stages. Cluster centers C_k^1 in stage-1 are obtained in landmark configuration space l_i , where l_i is the ground truth landmark positions for sample i . Then pseudo-label f_i^1 of sample i in stage-1 can be written as $f_i^1(k) = \exp\left(-\frac{\|l_i - C_k^1\|_2}{T}\right)$. We now have a landmark estimation \hat{l}_i^1 from stage-1 for sample i . In stage-2, we first define the landmark offset $\delta l_i^2 = \hat{l}_i^1 - l_i$, which is the correction should be made on stage-1 estimation. Cluster centers C_k^2 in stage-2 are obtained in landmark offset space δl_i^2 . Similarly, pseudo-label f_i^2 of sample i in stage-2 can be written as $f_i^2(k) = \exp\left(-\frac{\|\delta l_i^2 - C_k^2\|_2}{T}\right)$. Since outer product \otimes of two offsets contains the correlations between different fashion landmarks (e.g. ‘left collar’ v.s. ‘left sleeve’), we further include these contextual information into the pseudo-labels of stage-3 f^3 . To make the results of outer product comparable, we convert them into vectors by stacking columns, which is denoted as $\text{lin}(\cdot)$. The landmark offset of stage-3 is defined as $\delta l_i^3 = \hat{l}_i^2 - l_i$, where $\hat{l}_i^2 = \hat{l}_i^1 + \delta \hat{l}_i^2$ is the estimation made by stage-2. Thus, cluster centers C_k^3 in stage-3 are obtained in contextual offset space $\delta_{\text{context}} l_i^3 = \text{lin}(\delta l_i^3 \otimes \delta l_i^3)$. Similarly, pseudo-label f_i^3 of sample i in stage-3 can be written as $f_i^3(k) = \exp\left(-\frac{\|\delta_{\text{context}} l_i^3 - C_k^3\|_2}{T}\right)$. The pseudo-labels used in each stage are summarized in Table 1.

Table 1. Summary of pseudo-labels used in each stage. l_i is the ground truth landmark positions. \hat{l}_i^1 and \hat{l}_i^2 are the landmark estimations from stage-1 and 2, respectively. \otimes is the outer product and $\text{lin}(\cdot)$ is the linearize operation.

Stage	Clustering space	Indication	Dimension
1	l_i	Landmark configuration	20
2	$\delta l_i^2 = \hat{l}_i^1 - l_i$	Landmark offset	20
3	$\delta_{\text{context}} l_i^3 = \text{lin}(\delta l_i^3 \otimes \delta l_i^3)$ $\delta l_i^3 = \hat{l}_i^2 - l_i$	Contextual offset	20

Auto-Routing. Another important building block of DFA is the auto-routing mechanism. It is built upon the fact that the estimated pseudo-labels in stage-2 f^2 reflects the error patterns for each sample. We first associate each cluster center with an average error magnitude $e(C_k^2)$, $k = 1 \dots K$. This can be done by averaging the errors of training samples in each cluster. Then, we define the error function

$G(\cdot)$ within pseudo-label \hat{f}_i^2 for sample i in stage-2: $G\left(\hat{f}_i^2\right) = \sum_{k=1}^K e\left(C_k^2\right) \cdot \hat{f}_i^2(k)$. Therefore, the routing function r_i for sample i is formulated as

$$r_i = \mathbf{1}\left(G\left(\hat{f}_i^2\right) < \epsilon\right), \quad (2)$$

where $\mathbf{1}(\cdot)$ is the indicator function and ϵ is the error threshold for auto-routing. We set $\epsilon = 0.3$ empirically. If $r_i = 1$ indicates sample i will go through branch 1 in stage-3, and $r_i = 0$ indicates otherwise.

Training. Each stage of DFA is trained with multiple loss functions, including landmark estimation $L_{positions}$, visibility prediction $L_{visibility}$, and pseudo-label approximation L_{labels} . The overall loss function $L_{overall}$ is

$$L_{overall} = L_{positions}(l, \hat{l}) + \alpha(t)L_{visibility}(v, \hat{v}) + \beta(t)L_{labels}(f, \hat{f}), \quad (3)$$

where \hat{l} , \hat{v} and \hat{f} are the predicted landmark positions, visibility, and pseudo-labels respectively. We employ the Euclidean loss for $L_{positions}$ and L_{labels} , and the multinomial logistic loss is adopted for $L_{visibility}$. $\alpha(t)$ and $\beta(t)$ are the balancing weights between them. All the VGG-16 networks are pre-trained using ImageNet [37] and the entire DFA cascaded network is fine-tuned by stochastic gradient descent with back-propagation.

The proper scheduling of $\alpha(t)$ and $\beta(t)$ is very important for network performance. If they are too large, it disturbs the training of landmark positions. If they are too small, the training procedure cannot benefit from these auxiliary information. Similar to [38], we design a piecewise adjustment strategy for $\alpha(t)$ and $\beta(t)$ during training process,

$$\alpha(t) = \begin{cases} \alpha & t < t_1, \\ \frac{t-t_1}{t_2-t_1}\alpha & t_1 \leq t < t_2, \\ 0 & t_2 \leq t, \end{cases} \quad (4)$$

where $t_1 = 2000$ iterations and $t_2 = 4000$ iterations in our implementation. The adjustment for $\beta(t)$ takes a similar form.

Computations. For a three stage cascade to predict 8 fashion landmarks, DeepPose is required to train 17 VGG-16 models in total, while only three VGG-16 models need to be trained for DFA. Thus, our proposed approach at least saved 5 times computational costs.

4 Experiments

This section presents evaluation and analytical results of fashion landmark detection, as well as two applications including clothing attribute prediction and clothes retrieval.

Experimental Settings. For each clothing category, we randomly select 5K images for validation and another 5K images for test. The remaining 30K~40K images are used for training. We employ two metrics to evaluate fashion landmark detection, normalized error (NE) and the percentage of detected landmarks (PDL) [10]. NE is defined as the ℓ_2 distance between predicted landmarks and ground truth landmarks in the normalized coordinate space (*i.e.* divided by the width/height of the image), while PDL is calculated as the percentage of detected landmarks under certain overlapping criterion. Typically, smaller values of NE or higher values of PDL indicate better results.

Competing Methods. Since this work is the first study of fashion landmark detection, it is difficult to find direct comparisons. Nevertheless, to fully demonstrate the effectiveness of DFA, we compare it with two deep models, including DeepPose [10] and Image Dependent Pairwise Relations (IDPR) [11], which achieved best-performing results in human pose estimation. They are two representative methods that explored network cascade and graphical model to handle human pose. Specifically, DeepPose designed a cascaded deep regression framework on human body parts, while IDPR combined CNN and structural SVM [33] to model the relations among landmarks. To have a fair comparison, we replace the backbone networks in DeepPose and IDPR with VGG-16 and carefully train them using the same data and protocol as DFA did.

4.1 Ablation Study

We demonstrate the merits of each component in DFA.

Effectiveness of Network Cascade. Table 2 lists the performance of NE among three stages, where we have two observations. First, as shown in stage one, training DFA with both landmark position and visibility, denoted as ‘direct regression’, outperforms training DFA with only landmark position, denoted as ‘visibility’, showing that visibility helps landmark detection because it indicates variations of pose and scale (*e.g.* zoom-in). Second, cascade networks gradually reduce localization errors on all fashion landmarks from stage one to stage three. By predicting the corrections over previous stage, DFA decomposes a complex mapping problem into several subspace regression. For example, Fig. 8(e–g) demonstrates the quantitative stage-wise landmark detection results of DFA on different clothing items. In these cases, stage-1 gives rough predictions with shape constraints, while stage-2 and stage-3 refine results by referring to local and contextual correction patterns.

Effectiveness of Different Pseudo-Labels. Within each stage, choices of different pseudo-labels are explored. We design pseudo-labels representing landmark configurations, local landmark offsets and contextual landmark offsets for three stages respectively. Table 3 shows that pseudo-labels lead to substantial gains beyond direction regression, especially for the first stage. Next, we further

Table 2. Ablation study of DFA on different fashion landmarks. The normalized error (NE) is used here. ‘abs.’ indicates pseudo-labels obtained from absolute landmark positions. ‘offset’ indicates pseudo-labels obtained from local landmark offsets. ‘c. offset’ indicates pseudo-labels obtained from contextual landmark offsets. Intuitively, for the image size of 224×224 , the best prediction is achieved in stage-3 using contextual offset as pseudo-labels, whose errors in pixels are $0.048 \times 224 = 10.752$, 10.752, 20.384, 19.936, 15.904, and 16.128 respectively. ‘ $T = 20$ ’ is found empirically in the experiments.

Stage	Component	L. Collar	R. Collar	L. Sleeve	R. Sleeve	L. Hem	R. Hem	Avg.
1	Direct regression	.071	.071	.134	.130	.103	.102	.102
	– visibility	.104	.102	.213	.212	.141	.143	.153
	+ p.-labels($T = 1$)	.065	.065	.113	.112	.094	.095	.091
	+ p.-labels($T = 20$)	.057	.058	.106	.104	.088	.089	.084
2	Direct regression	.051	.051	.098	.096	.082	.080	.078
	+ p.-labels(abs.)	.050	.051	.097	.096	.080	.079	.076
	+ p.-labels(offset)	.049	.050	.095	.093	.077	.078	.074
3	Direct regression	.049	.049	.094	.093	.074	.076	.073
	+ two-branch	.048	.049	.094	.093	.073	.075	.072
	+ auto-routing	.048	.049	.092	.091	.071	.073	.070
	+ p.-labels(offset)	.048	.048	.091	.090	.071	.073	.069
	+ p.-labels(c. offset)	.048	.048	.091	.089	.071	.072	.068

Table 3. Ablation study of DFA on different evaluation subsets. NE is used here.

Stage	Component	N. Pose	M. Pose	L. Pose	M. Zoom-in	L. Zoom-in
1	Direct regression	.079	.100	.111	.151	.193
	+ p.-label($T = 20$)	.064	.077	.085	.120	.156
2	Direct regression	.055	.072	.081	.104	.151
	+ p.-label(offset)	.053	.069	.078	.102	.148
3	Direct regression	.052	.069	.078	.096	.148
	+ auto-routing	.051	.067	.077	.093	.146
	+ p.-label(c. offset)	.050	.066	.075	.091	.144

justify the forms of different pseudo-labels adopted for each stage. In stage-1, we find that using soft label, denoted as ‘+p. labels ($T = 20$)’, instead of hard label, denoted as ‘+p. labels ($T = 1$)’, results in better performance, because soft label is more informative in identifying landmark configuration of samples. In stage-2, pseudo-label generated from offset landmark positions is superior to that generated from absolute landmark positions since landmark offsets can provide more guidance on the local corrections to be predicted. In stage-3, including contextual landmark offsets help achieve further gains, due to the fact that landmark corrections to be made are generally correlated.

Effectiveness of Auto-Routing. Finally, we show that auto-routing is an effective way to tackle data with different correction difficulties. From Table 2

stage-3, we can see that auto-routing (denoted as ‘+auto-routing’) provides more benefits when compared with averaging the predictions from two branches trained using all data (denoted as ‘+two-branch’). By further inspecting the stage-wise performance on each evaluation set, which is shown in Table 3, we can observe that auto-routing mechanism improves the performance of medium/large zoom-in subsets, showing that the routing function makes one of the branch in stage-3 focus on difficult samples.

4.2 Benchmarking

To illustrate the effectiveness of DFA, we compare it with state-of-the-art human pose estimation methods like DeepPose [10] and IDPR [11]. We also analyze the strengths and weaknesses of each method on fashion landmark detection.

Landmark Types. Figure 5 (the first row) shows the percentage of detection rates on different fashion landmarks, where we have three observations. First, on landmark ‘hem’, DeepPose performs better when distance threshold is small, while IDPR catches up when the threshold is large, because DeepPose is a part-based method which can locally refine the results for easy cases. Second, collars are the easiest landmarks to be detected while sleeves are the hardest. Third, DFA consistently outperforms both DeepPose and IDPR or shows comparable results on all fashion landmarks, showing that the pseudo-labels and auto-routing mechanisms enable robust fashion landmark detection.




Fig. 5. Performance of fashion landmark detection on different fashion landmarks (the first row) and different clothing types (the second row). [px] represents pixels. The percentage of detected landmarks (PDL) is used here.

Clothing Types. Figure 5 (the second row) shows the percentage of detection rates on different clothing types. Again, DFA outperforms all other methods on all distance thresholds. We have two additional observations. First, DFA (stage-1) already achieves comparable results on full-body and lower-body clothes when compared with IDPR and DeepPose (stage-3). Second, upper-body clothes pose most challenges on fashion landmark detection. It is partially due to the various clothing sub-categories contained.

Difficulty Levels. Figure 6 shows the percentage of detection rates on different evaluation subsets, with the distance threshold fixed at 15 pixels. Two observations are made here. First, fashion landmark detection is a challenging task. Even the detection rate for normal pose set is just above 70%. More powerful model needs to be developed. Second, DFA has the most advantages on medium pose/zoom-in subsets. Pseudo-labels provide effective shape constraints for hard cases. Please also note that DFA (stage-3) requires much less computational costs than DeepPose (stage-3).




Fig. 6. Performance of fashion landmark detection on different evaluation subsets. The distance threshold is fixed at 15 pixels. PDL is used here.

For a 300×300 image, DFA takes around $100ms$ to detect full sets of fashion landmarks on a single GTX Titan X GPU. In contrast, DeepPose needs nearly $650ms$ in the same setting. Our framework has large potential in real-life applications. Visual results of fashion landmark detection by different methods are given in Fig. 8.

4.3 Generalization of DFA

To test the generalization ability of the proposed framework, we further apply DFA on a related task, *i.e.* human pose estimation, as reported in Table 4. In the following, DFA is trained and evaluated on LSP dataset [40] as [11] did.

First, we compare DFA system with other state-of-the-art methods on pose estimation task. Without much adaptation, DFA achieves 74.4 mean strict PCP results, with 87, 91, 70, 56, 81, 76 for ‘torso’, ‘head’, ‘u.arms’, ‘l.arms’, ‘u.legs’ and ‘l.legs’ respectively. It shows comparable results to [11] and outperforms several recent works [10, 30, 32, 39], showing that DFA is a general approach for structural prediction problem besides fashion landmark detection.

Table 4. Comparison of strict PCP results on the LSP dataset. DFA shows good generalization ability to human pose estimation.

Method	Torso	Head	U.arms	L.arms	U.legs	L.legs	Mean
Fu <i>et al.</i> [32]	77.7	85.4	75.0	71.9	62.1	48.8	67.7
Ouyang <i>et al.</i> [39]	85.8	83.1	63.3	46.6	76.5	72.2	68.6
Pose Machines [30]	93.1	83.6	76.8	68.1	42.2	85.4	72.0
DeepPose [10]	-	-	56	38	77	71	-
DFA	90.8	87.2	70.4	56.2	80.6	75.8	74.4
IDPR [11]	92.7	87.8	69.2	55.4	82.9	77.0	75.0
+ p.-label	93.5	88.5	72.3	59.0	83.9	78.7	77.0
+ auto-routing	94.1	88.9	74.3	61.5	85.1	80.4	78.6

Then, we show that pseudo-label and auto-routing scheme of DFA can be generalized to improve performance of pose estimation methods, such as IDPR [11]. [11] trained DCNN and achieved 75 mean strict PCP. We add pseudo-labels to this DCNN and include auto-routing in cascading predictions. Training and evaluation of graphical model are kept unchanged. Pseudo-labels leverage the result to 77 mean strict PCP and auto-routing leads to another 1.6 point gain. It demonstrates that pseudo-labels and auto-routing are effective and complementary techniques to current methods.

4.4 Applications

Finally, we show that fashion landmarks can facilitate clothing attribute prediction and clothes retrieval. We employ a subset of DeepFashion dataset [3], which contains 10K images, 50 clothing attributes and corresponding image pairs (*i.e.* the images containing the same clothing item). We compare fashion landmarks with different localization schemes, including the full image, the bounding box (bbox) of clothing item, and the human-body joints, where fashion landmarks are detected by DFA, human joints are obtained by the executable code of [11], and bounding boxes are manually annotated. For both tasks of attribute recognition and clothes retrieval, we use off-the-shelf CNN features as described in [35].

Attribute Prediction. We train a multi-layer perceptron (MLP) using the extracted CNN features as input to predict all 50 attributes. Following [41], we employ the top- k recall rate as measuring criteria, which is obtained by ranking the classification scores and determine how many ground truth attributes have been found in the top- k predicted attributes. Overall, the average top-5 recall rates on 50 attributes of ‘full image’, ‘bbox’, ‘human joints’, and ‘fashion landmarks’ are 27%, 53%, 65% and 73%, respectively, showing that fashion landmarks are the most effective representation for attribute prediction of fashion items. Figure 7(a) shows the top-5 recall rates of ten representative attributes,

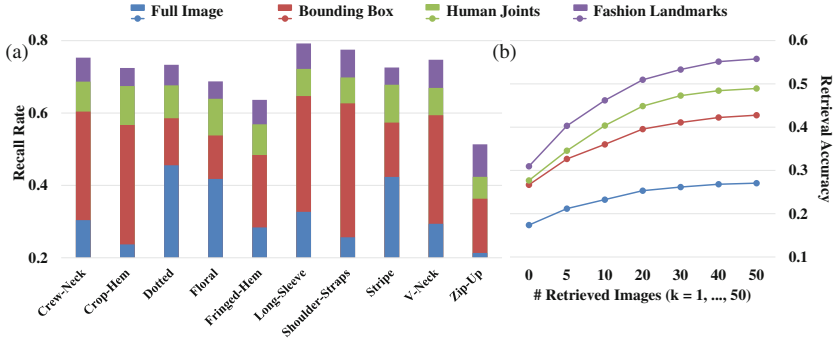


Fig. 7. Experimental results on clothing attribute prediction and clothes retrieval using features extracted from ‘full image’, ‘bbox’, ‘human joints’, and ‘fashion landmarks’: (a) the top-5 recall rates of clothing attributes and (b) the top- k clothes retrieval accuracy.

e.g. ‘stripe’, ‘long-sleeve’, and ‘V-neck’. We observe that fashion landmark outperforms all the other localization schemes in all the attributes, especially for part-based attributes, such as ‘zip-up’ and ‘shoulder-straps’.

Clothes Retrieval. We adopt the ℓ_2 distance between the extracted CNN features for clothes retrieval. The top- k retrieval accuracy is used to measure the performance, such that given a query image, correct retrieval is considered as

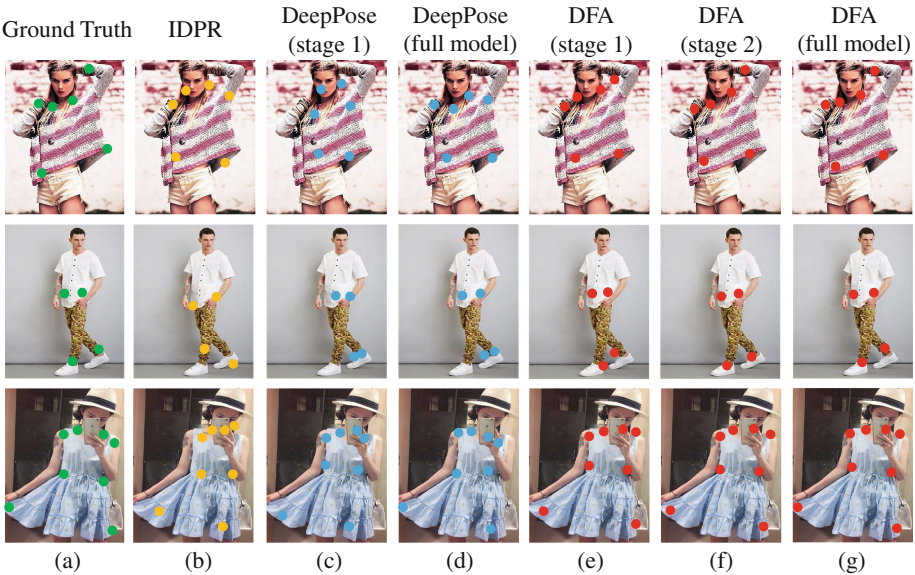


Fig. 8. Visual results of fashion landmark detection by different methods: (a) Ground Truth, (b) IDPR, (c) DeepPose (stage 1), (d) DeepPose (full model), (e) DFA (stage 1), (f) DFA (stage 2) and (g) DFA (full model).

the exact clothing item has been found in the top- k retrieved gallery images. As shown in Fig. 7(b), the top-20 retrieval accuracies for ‘full image’, ‘bbox’, ‘human joints’, and ‘fashion landmarks’ are 25 %, 40 %, 45 %, and 51 % respectively. When $k = 1$ and $k = 5$, features extracted around fashion landmarks still perform better than the other alternatives, demonstrating that fashion landmarks provide more discriminative information beyond traditional paradigms.

5 Conclusions

This paper introduced fashion landmark detection, which is an important step towards robust fashion recognition. To benchmark fashion landmark detection, we introduced a large-scale fashion landmark dataset (FLD). With FLD, we proposed a deep fashion alignment network (DFA) for robust fashion landmark detection, which leverages pseudo-labels and auto-routing mechanism to reduce the large variations presented in fashion images. Extensive experiments showed the effectiveness of different components as well as the generalization ability of DFA. To demonstrate the usefulness of fashion landmark, we evaluated on two fashion applications, clothing attribute prediction and clothes retrieval. Experiments revealed that fashion landmark is a more discriminative representation than clothes bounding boxes and human joints for fashion-related tasks, which we hope could facilitate future research.

Acknowledgements. This work is partially supported by SenseTime Group Limited, the Hong Kong Innovation and Technology Support Programme, the General Research Fund sponsored by the Research Grants Council of the Kong Kong SAR (CUHK 416312), the External Cooperation Program of BIC, Chinese Academy of Sciences (No. 172644KYSB20150019), the Science and Technology Planning Project of Guangdong Province (2015B010129013, 2014B050505017), and the National Natural Science Foundation of China (61503366, 61472410; Corresponding author is Ping Luo).

References

1. Huang, J., Feris, R.S., Chen, Q., Yan, S.: Cross-domain image retrieval with a dual attribute-aware ranking network. In: ICCV (2015)
2. Kiapour, M.H., Han, X., Lazebnik, S., Berg, A.C., Berg, T.L.: Where to buy it: matching street clothing photos in online shops. In: ICCV (2015)
3. Liu, Z., Luo, P., Qiu, S., Wang, X., Tang, X.: Deepfashion: powering robust clothes recognition and retrieval with rich annotations. In: CVPR, pp. 1096–1104 (2016)
4. Liu, S., Song, Z., Liu, G., Xu, C., Lu, H., Yan, S.: Street-to-shop: cross-scenario clothing retrieval via parts alignment and auxiliary set. In: CVPR, pp. 3330–3337 (2012)
5. Di, W., Wah, C., Bhardwaj, A., Piramuthu, R., Sundaresan, N.: Style finder: fine-grained clothing style detection and retrieval. In: CVPR Workshops, pp. 8–13 (2013)
6. Kiapour, M.H., Yamaguchi, K., Berg, A.C., Berg, T.L.: Hipster wars: discovering elements of fashion styles. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8689, pp. 472–488. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10590-1_31](https://doi.org/10.1007/978-3-319-10590-1_31)

7. Simo-Serra, E., Fidler, S., Moreno-Noguer, F., Urtasun, R.: Neuroaesthetics in fashion: modeling the perception of beauty. In: CVPR (2015)
8. Chen, H., Gallagher, A., Girod, B.: Describing clothing by semantic attributes. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012. LNCS, vol. 7574, pp. 609–623. Springer, Heidelberg (2012). doi:[10.1007/978-3-642-33712-3_44](https://doi.org/10.1007/978-3-642-33712-3_44)
9. Bossard, L., Dantone, M., Leistner, C., Wengert, C., Quack, T., Gool, L.: Apparel classification with style. In: Lee, K.M., Matsushita, Y., Rehg, J.M., Hu, Z. (eds.) ACCV 2012. LNCS, vol. 7727, pp. 321–335. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-37447-0_25](https://doi.org/10.1007/978-3-642-37447-0_25)
10. Toshev, A., Szegedy, C.: Deeppose: human pose estimation via deep neural networks. In: CVPR, pp. 1653–1660 (2014)
11. Chen, X., Yuille, A.L.: Articulated pose estimation by a graphical model with image dependent pairwise relations. In: NIPS, pp. 1736–1744 (2014)
12. Tompson, J.J., Jain, A., LeCun, Y., Bregler, C.: Joint training of a convolutional network and a graphical model for human pose estimation. In: NIPS, pp. 1799–1807 (2014)
13. Carreira, J., Agrawal, P., Fragkiadaki, K., Malik, J.: Human pose estimation with iterative error feedback (2015). arXiv preprint [arXiv:1507.06550](https://arxiv.org/abs/1507.06550)
14. Pfister, T., Charles, J., Zisserman, A.: Flowing convnets for human pose estimation in videos. In: ICCV, pp. 1913–1921 (2015)
15. Fan, X., Zheng, K., Lin, Y., Wang, S.: Combining local appearance and holistic view: dual-source deep neural networks for human pose estimation. In: CVPR, pp. 1347–1355 (2015)
16. Wang, X., Zhang, T.: Clothes search in consumer photos via color matching and attribute learning. In: ACM MM, pp. 1353–1356 (2011)
17. Chen, Q., Huang, J., Feris, R., Brown, L.M., Dong, J., Yan, S.: Deep domain adaptation for describing people based on fine-grained clothing attributes. In: CVPR, pp. 5315–5324 (2015)
18. Yamaguchi, K., Kiapour, M.H., Berg, T.: Paper doll parsing: retrieving similar styles to parse clothing items. In: ICCV, pp. 3519–3526 (2013)
19. Kalantidis, Y., Kennedy, L., Li, L.J.: Getting the look: clothing recognition and segmentation for automatic product suggestions in everyday photos. In: ICMR, pp. 105–112 (2013)
20. Fu, J., Wang, J., Li, Z., Xu, M., Lu, H.: Efficient clothing retrieval with semantic-preserving visual phrases. In: Lee, K.M., Matsushita, Y., Rehg, J.M., Hu, Z. (eds.) ACCV 2012. LNCS, vol. 7725, pp. 420–431. Springer, Heidelberg (2013). doi:[10.1007/978-3-642-37444-9_33](https://doi.org/10.1007/978-3-642-37444-9_33)
21. Yamaguchi, K., Berg, T.L., Ortiz, L.E.: Chic or social: visual popularity analysis in online fashion networks. In: ACM MM, pp. 773–776 (2014)
22. Yamaguchi, K., Kiapour, M.H., Ortiz, L.E., Berg, T.L.: Parsing clothing in fashion photographs. In: CVPR, pp. 3570–3577 (2012)
23. Yang, W., Luo, P., Lin, L.: Clothing co-parsing by joint image segmentation and labeling. In: CVPR, pp. 3182–3189 (2014)
24. Liang, X., Xu, C., Shen, X., Yang, J., Liu, S., Tang, J., Lin, L., Yan, S.: Human parsing with contextualized convolutional neural network. In: ICCV, pp. 1386–1394 (2015)
25. Ferrari, V., Marin-Jimenez, M., Zisserman, A.: Progressive search space reduction for human pose estimation. In: CVPR, pp. 1–8 (2008)
26. Yang, Y., Ramanan, D.: Articulated pose estimation with flexible mixtures-of-parts. In: CVPR, pp. 1385–1392 (2011)

27. Dantone, M., Gall, J., Leistner, C., Gool, L.: Human pose estimation using body parts dependent joint regressors. In: CVPR, pp. 3041–3048 (2013)
28. Sapp, B., Taskar, B.: Modec: Multimodal decomposable models for human pose estimation. In: CVPR, pp. 3674–3681 (2013)
29. Belagiannis, V., Ruppel, C., Carneiro, G., Navab, N.: Robust optimization for deep regression. In: ICCV, pp. 2830–2838. IEEE (2015)
30. Ramakrishna, V., Munoz, D., Hebert, M., Andrew Bagnell, J., Sheikh, Y.: Pose machines: articulated pose estimation via inference machines. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8690, pp. 33–47. Springer, Heidelberg (2014). doi:[10.1007/978-3-319-10605-2_3](https://doi.org/10.1007/978-3-319-10605-2_3)
31. Dantone, M., Gall, J., Leistner, C., Van Gool, L.: Body parts dependent joint regressors for human pose estimation in still images. TPAMI **36**(11), 2131–2143 (2014)
32. Fu, L., Zhang, J., Huang, K.: Beyond tree structure models: a new occlusion aware graphical model for human pose estimation. In: ICCV, pp. 1976–1984 (2015)
33. Tsochantaridis, I., Hofmann, T., Joachims, T., Altun, Y.: Support vector machine learning for interdependent and structured output spaces. In: ICML, p. 104 (2004)
34. Jordan, M.I., Ghahramani, Z., Jaakkola, T.S., Saul, L.K.: An introduction to variational methods for graphical models. Mach. Learn. **37**(2), 183–233 (1999)
35. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition (2014). arXiv preprint [arXiv:1409.1556](https://arxiv.org/abs/1409.1556)
36. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: CVPR, pp. 503–511 (2001)
37. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: a large-scale hierarchical image database. In: CVPR, pp. 248–255 (2009)
38. Lee, D.H.: Pseudo-label: the simple and efficient semi-supervised learning method for deep neural networks. In: Workshop on Challenges in Representation Learning, ICML, vol. 3 (2013)
39. Ouyang, W., Chu, X., Wang, X.: Multi-source deep learning for human pose estimation. In: CVPR, pp. 2329–2336 (2014)
40. Johnson, S., Everingham, M.: Clustered pose and nonlinear appearance models for human pose estimation. In: BMVC, vol. 2, p. 5 (2010)
41. Gong, Y., Jia, Y., Leung, T., Toshev, A., Ioffe, S.: Deep convolutional ranking for multilabel image annotation (2013). arXiv preprint [arXiv:1312.4894](https://arxiv.org/abs/1312.4894)