

A Global-Local Embedding Module for Fashion Landmark Detection

Sumin Lee¹Sungchan Oh²Chanho Jung³Changick Kim¹¹Korea Advanced Institute of Science and Technology(KAIST)²Electronics and Telecommunications Research Institute³Hanbat National University

{shum_ming, changick}@kaist.ac.kr, sungchan.oh@etri.re.kr, peterjung@hanbat.ac.kr

Abstract

Detecting fashion landmarks is a fundamental technique for visual clothing analysis. Due to the large variation and non-rigid deformation of clothes, localizing fashion landmarks suffers from large spatial variances across poses, scales, and styles. Therefore, understanding contextual knowledge of clothes is required for accurate landmark detection. To that end, in this paper, we propose a fashion landmark detection network with a global-local embedding module. The global-local embedding module is based on a non-local operation for capturing long-range dependencies and a subsequent convolution operation for adopting local neighborhood relations. With this processing, the network can consider both global and local contextual knowledge for a clothing image. We demonstrate that our proposed method has an excellent ability to learn advanced deep feature representations for fashion landmark detection. Experimental results on two benchmark datasets show that the proposed network outperforms the state-of-the-art methods.

1. Introduction

Visual fashion analysis has attracted research attention in recent years because of its huge potential usefulness in industry. With a development of large fashion datasets [7, 8], deep learning-based methods have achieved significant progresses in many tasks, such as clothes classification [7, 10, 6, 1], clothing retrieval [2, 7, 1, 4], and clothes generation [13, 3]. The fashion analysis is a challenging task due to the large variation and non-rigid deformation of clothes in images. To deal with this issue, recent methods [7, 10, 6, 1] utilized fashion landmarks. As illustrated in Fig 1, the fashion landmarks are key-points describing clothing structures such as collars, sleeves, waistlines, and hemlines. Since these methods improved their performance considerably, fashion landmark detection is one of the key issues. For localizing fashion landmarks, Liu et al. [8, 7]



Figure 1. Examples of fashion landmarks. The landmarks represent collars (C), sleeves (S), waistlines (W), and hemlines (H). The red and yellow circles are used to indicate right and left sides in a image. Note that different types of clothes have different number of landmarks.

utilized a regression method. Wang et al. [10] designed a model to capture kinematic and symmetry grammar of clothing landmarks. However, the fashion landmarks exhibit large spatial variances across poses, scales, and styles of clothing items. Due to this property, the models require to understand comprehensive semantic information of clothes for accurate landmark detection.

To address this issue, we propose a fashion landmark detection network with a global-local embedding module to exploit rich contextual knowledge for clothes. The global-local embedding module is specifically designed for embedding fashion landmark information by employing a non-local block [11] followed by convolutions. Note that, in our proposed module, the non-local operation captures long-range global dependencies within a clothing image, and the following convolution operation enhances the local representation power of output features. With this processing, the output features are well facilitated to have global as well as local contextual information for input clothing image. Lastly, by upsampling the feature map to the same size of the input fashion image, our network predicts high-

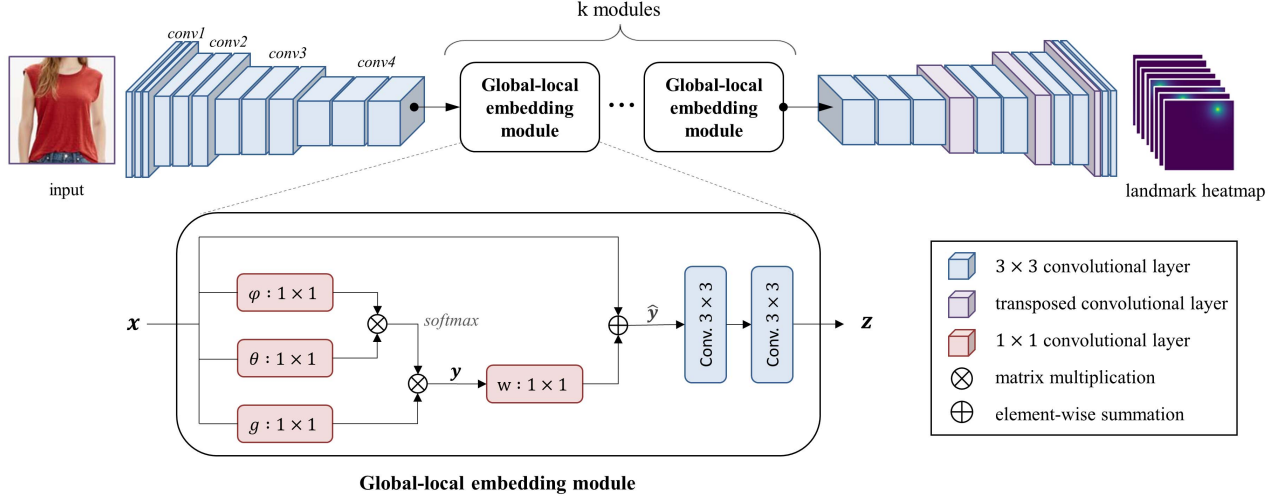


Figure 2. Illustration of our proposed network. Our network consists of basic convolutional network VGG-16 for feature extraction, the global-local embedding module for extracting rich landmark embedding features, and the upsampling network for predicting the landmark heatmaps.

resolution heatmaps for more accurate landmark localization. To demonstrate the strength of our proposed network, we conduct experiments by using two datasets, Deepfashion [7] and FLD [8]. Experimental results show that our network outperforms the state-of-the-art methods.

2. Methods

2.1. Network architecture

Our network consists of three parts: a feature extractor, a global-local embedding module, and an upsampling network. The overall architecture of the proposed method is illustrated in Fig. 2. First, an input fashion image is resized to 224×224 and fed into the feature extractor. For the feature extractor, we use VGG-16 network [9] except the last convolutional layer and initialize weights with parameters pretrained on ImageNet [5]. After the conv4_3 layer of VGG-16, the global-local embedding module is employed to generate rich landmark embedding features. We describe details of the global-local embedding module in 2.2.

From generated rich features, the landmark localization network predicts landmark heatmaps. For more accurate landmark estimation, we produce high-resolution landmark heatmaps which have the same size to the input image. Each value of the heatmaps represents the probability that there is a landmark. To increase the spatial size of the feature map, we use several transposed convolutions with kernel size 4, stride 2, and padding 1. At the end of the upsampling network, we utilize a 1×1 convolution to produce a heatmap for each landmark.

2.2. Global-local embedding module

We introduce a global-local embedding module to exploit rich contextual knowledge of a clothing item. As shown in the bottom of Fig. 2, the global-local embedding module consists of a non-local block [11] and two convolutional layers.

Let \mathbf{x} denote input features of the global-local embedding module, which is the conv4_3 feature map in our network. The output feature is first represented by the non-local block. Compared to a conventional convolution operation, the non-local operation calculates long-range dependencies between any two different points, and sums up the weighted input features. This operation is formulated as:

$$\mathbf{y} = \frac{1}{C(\mathbf{x})} f(\theta(\mathbf{x}), \phi(\mathbf{x})) g(\mathbf{x}), \quad (1)$$

where $C(\mathbf{x})$ is a normalization factor of \mathbf{x} , and $\theta(\cdot)$, $\phi(\cdot)$, and $g(\cdot)$ are 1×1 convolutions. We define the function f as an embedded Gaussian function as follows:

$$f(\theta(\mathbf{x}), \phi(\mathbf{x})) = \exp(\theta(\mathbf{x})^T \phi(\mathbf{x})). \quad (2)$$

By utilizing the non-local operation with a residual connection, we obtain the $\hat{\mathbf{y}}$:

$$\hat{\mathbf{y}} = w(\mathbf{y}) + \mathbf{x}. \quad (3)$$

where $w(\cdot)$ is a 1×1 convolution. Finally, $\hat{\mathbf{y}}$ has rich global knowledge of the clothing item by weighted with long-range relationship.

For locating landmark positions, it is necessary to assimilate the global-informative features in a local manner. To

FLD									
	L.Collar	R.Collar	L.Sleeve	R.Sleeve	L.Waistline	R.Waistline	L.Hem	R.Hem	Avg.
FashionNet [7]	0.0781	0.0803	0.0975	0.0923	0.0874	0.0821	0.0802	0.0893	0.0859
DFA [8]	0.0480	0.0480	0.0910	0.0890	-	-	0.0710	0.0720	0.0680
DLAN [12]	0.0531	0.0547	0.0705	0.0735	0.0752	0.0748	0.0693	0.0675	0.0672
BCRNNs [10]	0.0463	0.0471	0.0627	0.0614	0.0635	0.0692	0.0635	0.0527	0.0583
Ours	0.0386	0.0391	0.0675	0.0672	0.0576	0.0605	0.0615	0.0621	0.0568
DeepFashion									
	L.Collar	R.Collar	L.Sleeve	R.Sleeve	L.Waistline	R.Waistline	L.Hem	R.Hem	Avg.
FashionNet [7]	0.0854	0.0902	0.0973	0.0935	0.0854	0.0845	0.0812	0.0823	0.0872
DFA [8]	0.0628	0.0637	0.0658	0.0621	0.0726	0.0702	0.0658	0.0663	0.0660
DLAN [12]	0.0570	0.0611	0.0672	0.0647	0.0703	0.0694	0.0624	0.0627	0.0643
BCRNNs [10]	0.0415	0.0404	0.0496	0.0449	0.0502	0.0523	0.0537	0.0551	0.0484
Liu et al. [6]	0.0332	0.0346	0.0487	0.0519	0.0422	0.0429	0.0620	0.0639	0.0474
Ours	0.0312	0.0324	0.0427	0.0434	0.0361	0.0373	0.0442	0.0475	0.0393

Table 1. Quantitative results for clothing landmark detection on Deepfashion and FLD with respect to normalized error (NE). The best results are marked in **bold**.

improve the local representation power of the output feature \mathbf{z} , two convolutional layers perform a weighted sum in local neighborhoods:

$$\mathbf{z} = F_2(F_1(\hat{\mathbf{y}})), \quad (4)$$

where F_i denotes the i -th convolutional layer including a 3×3 convolution, a batch normalization, and a nonlinear function (i.e. ReLU) sequentially. By this process, the output feature \mathbf{z} is manipulated to have not only global but also local contextual information of the input clothing image.

We use sequentially k global-local embedding modules to gradually produce more advanced deep feature representation. We empirically set k as 2 for Deepfashion and 3 for FLD.

3. Experiments

3.1. The benchmark datasets and evaluation

We evaluate the proposed network on two datasets, DeepFashion [7] and FLD [8]. DeepFashion is a large fashion dataset with 289,222 images. Those images are composed of 209,222 images for training, 40,000 images for validation, and 40,000 images for testing. FLD [8] is a fashion landmark dataset with 123,016 images with more diverse variations in poses, scales, and background. FLD images are divided into 83,033 images for training, 19,991 images for validation, and remaining 19,991 images for testing. Each image of both two datasets is annotated with bounding boxes and landmarks. There are 8 landmarks for full-body clothes, 6 landmarks for upper clothes, and 4 landmarks for lower clothes.

We adopt normalized error (NE) metric [8] for evaluation of our proposed method and comparisons to the other state-of-the-art methods. NE is the L_2 distance between predicted

and ground-truth landmarks in the normalized coordinate space, formulated as:

$$NE = \frac{1}{N} \sum_{i=1}^N \frac{\|p_i - \hat{p}_i\|}{h_i \times w_i}, \quad (5)$$

where p_i and \hat{p}_i are ground-truth and predicted landmarks of i -th sample respectively, N is the number of samples, and h_i and w_i are the height and width of i -th sample.

3.2. Results

We conduct experiments on two large datasets and compare results of the proposed network with the state-of-the-art methods [7, 8, 12, 10, 6]. Table 1 summaries comparison results. From the table, our proposed network outperforms all the competitors at 0.0568 on FLD and 0.0393 on DeepFashion. Among the landmarks, it is hard to discriminate between waistlines and hems landmarks which have the largest error rate in other methods. Our proposed method achieves performance improvements in detecting waistlines and hems landmarks. From these results, it is proved that our network learns more advanced feature representations for fashion landmark detection with the aid of the proposed global-local embedding module.

We also visualize landmark detection results in Fig. 3. We can see that the proposed model discriminates right- and left-side landmarks even in the back-view and side-view images. According to these results, it is demonstrated that our network is robust to view-point variations and deformations.

4. Conclusion

In this paper, we have proposed a fashion landmark detection network with a global-local embedding module.



Figure 3. Qualitative results on DeepFashion (first row) and FLD (second row). For each example, ground-truth landmarks are shown in right image and predicted landmarks are shown in left image. Red and yellow circles indicate right-side and left-side landmarks, respectively.

Our proposed global-local embedding module is based on a non-local operation and a convolution operation. Utilizing global-local embedding module facilitated to exploit not only global but also local contextual knowledge of a clothing item. We evaluated our method on two benchmark datasets, and achieved the state-of-the-art performance over recent methods. Experimental results demonstrated that our proposed method improves the feature representation of a clothing item for fashion landmark detection.

Acknowledgement

This work was supported by Electronics and Telecommunications Research Institute(ETRI) grant funded by the Korean government. [19ZS1100, Core Technology Research for Self-Improving Artificial Intelligence System]

References

- [1] Y. Ge, R. Zhang, X. Wang, X. Tang, and P. Luo. Deep-fashion2: A versatile benchmark for detection, pose estimation, segmentation and re-identification of clothing images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2019.
- [2] K. M. Hadi, X. Han, S. Lazebnik, A. C. Berg, and T. L. Berg. Where to buy it: Matching street clothing photos in online shops. In *IEEE International Conference on Computer Vision (ICCV)*. 2015.
- [3] X. Han, Z. Wu, Z. Wu, R. Yu, and L. S. Davis. Viton: An image-based virtual try-on network. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.
- [4] J. Huang, R. S. Feris, Q. Chen, and S. Yan. Cross-domain image retrieval with a dual attribute-aware ranking network. In *IEEE International Conference on Computer Vision (ICCV)*. 2015.
- [5] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems (NIPS)*. 2012.
- [6] J. Liu and H. Lu. Deep fashion analysis with feature map up-sampling and landmark-driven attention. In *European Conference on Computer Vision (ECCV) Workshop*. Springer, 2018.
- [7] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2016.
- [8] Z. Liu, S. Yan, P. Luo, X. Wang, and X. Tang. Fashion landmark detection in the wild. In *European Conference on Computer Vision (ECCV)*. Springer, 2016.
- [9] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *International Conference on Learning Representations (ICLR)*. 2015.
- [10] W. Wang, Y. Xu, J. Shen, and S. Zhu. Attentive fashion grammar network for fashion landmark detection and clothing category classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.
- [11] X. Wang, R. Girshick, A. Gupta, and K. He. Non-local neural networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2018.
- [12] S. Yan, Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang. Unconstrained fashion landmark detection via hierarchical recurrent transformer networks. In *ACM International Conference on Multimedia*. 2017.
- [13] S. Zhu, R. Urtasun, S. Fidler, D. Lin, and C. Change Loy. Be your own prada: Fashion synthesis with structural coherence. In *IEEE International Conference on Computer Vision (ICCV)*. 2017.