

Pose Guided Attention for Multi-label Fashion Image Classification

Beatriz Quintino Ferreira¹, João P. Costeira¹, Ricardo G. Sousa², Liang-Yan Gui³, João P. Gomes^{1*}

¹ ISR-IST, Universidade Lisboa, ² Farfetch, ³ Carnegie Mellon University

{beatrizquintino, jpc, jpg}@isr.ist.utl.pt, ricardo.sousa@farfetch.com, lgui@andrew.cmu.edu

Abstract

We propose a compact framework with guided attention for multi-label classification in the fashion domain. Our visual semantic attention model (VSAM) is supervised by automatic pose extraction creating a discriminative feature space. VSAM outperforms the state of the art for an in-house dataset and performs on par with previous works on the DeepFashion dataset, even without using any landmark annotations. Additionally, we show that our semantic attention module brings robustness to large quantities of wrong annotations and provides more interpretable results.

1. Introduction

Image classification is a fundamental Computer Vision task, widely applied in the fashion industry to generate rich product descriptions and automate product tagging. This automation is pivotal when dealing with extremely large collections, which naturally arise in most e-commerce platforms. With the advent of Deep Convolutional Neural Networks (CNN's), combined with the availability of massive amounts of data, the ability to extract categories and attributes from visual data gained extreme relevance. However, the performance gains on the computing side usually come with a high cost from human-intensive tasks to generate high quality training data, a key (and expensive) issue in the fashion industry.

Fashion attributes are often associated to specific locations (e.g short sleeve, neckline). This knowledge is either disregarded in purely data-driven approaches [7, 13] or requires the mentioned expensive annotation processes [10, 5]. While the former lacks interpretability and robustness to image artifacts, the latter is impractical at very large scale.

In this work we propose a CNN model that jointly learns to predict fashion categories (multi-class problem) and attributes (multi-label problem) by focusing on the relevant image regions through a guided attention mechanism.

*This work was partially funded by FCT via grant [PD/BD/114430/2016] and project [UID/EEA/50009/2019]. We also thank the guidance and support provided by Professor José F. Moura.



Figure 1. The effect of the pose-guided attention: left - saliency maps highlight neck and arm regions for V-neck and 3/4 sleeves attributes, respectively; right - CAMs highlight the relevant region for Jackets category.

Our approach is premised on the hypothesis that classification tasks benefit if the model identifies salient image regions amplifying their influence, while suppressing irrelevant and potentially confusing information in other regions (see Fig. 1).

We exploit the relation between attribute localization and visual appearance by embedding a semantic attention module guided by body pose. Building on an off-the-shelf architecture (VGG) and adding much less complexity, our model achieves similar performance to state-of-the-art models that use extra (costly) annotations and surpasses by a large margin those who do not. Specifically, our method outperforms previous approaches [7] for their "in-house" dataset and is on par with the state of the art for the DeepFashion dataset [6]. We note that the supervision of the semantic attention comes at a very low cost, as there are readily available pose detectors [1] providing high accuracy and nearly real-time detections.

Finally, we demonstrate that the semantic attention module increases the robustness to large quantities of wrong/missing annotations, that constitutes a prevailing issue in very large datasets.

This supervised attention guides the model to learn a feature space that is more suitable and robust (in terms of inter-class and inter-label confusion) for this specific problem of fashion items classification. In contrast to previous works [13, 10, 5], we learn our attention in a supervised

manner that encodes the underlying context and semantics of the fashion image classification problem.

2. Related work

Currently, CNNs are known to achieve leading performance for the multi-class and multi-label problems, and numerous works have recently addressed these problems in the context of fashion analysis, see for example [6, 10, 7, 2, 5]. The work in [2] proposes binary models, where each model predicts a fashion attribute, and introduces product type classifiers to mitigate outlier impact, as a pre-validation step before deciding on an attribute. However, this pipeline of models is neither efficient nor scalable to large datasets with numerous categories and attributes. On the contrary, in our previous work [7], we contributed with a unified model that jointly outputs category, subcategory and attributes predictions leveraging the hierarchical category tree structure to explore label relations. This model improved performance over a pipeline of state-of-the-art models performing the three classification tasks individually and independently.

Attention mechanisms have become very popular in deep learning to help focusing the learning process on the parts of the inputs that are relevant to the task, contributing not only to significant performance improvements but also to model interpretability. Accordingly, a model with the right attention should concentrate on the image regions that are discriminative to the classification (e.g., to classify a long sleeve, the model is supposed to focus on the regions that contain the sleeve) [13, 10]. Specifically in the fashion domain, [5] proposes an attentive fashion network with up-sampled feature maps for landmark localization, category classification and attribute prediction that is shown to outperform previous models on the DeepFashion dataset [6]. Contrary to [10], the method in [5] generates higher-resolution maps and combines separate attention branches in a unified branch, acting as a soft constraint to the model. However, the inclusion of these branches causes a drastic increase in complexity and parameters.

The previous examples implement self-learned/unsupervised attention modules. Nevertheless, when there is *a priori* knowledge, this knowledge can be used to learn the attention in a supervised manner. In [4] a supervised attention module is also added to off-the-shelf classification CNNs to increase recognition performance. However, as opposed to [4], where the ground-truth attention heatmaps are crowd-sourced (human derived click masks), we learn our attention heatmaps in an automatic manner resorting to the pose extracted by a pose detector.

3. Semantic attention model for fashion images classification

Our task is to predict a category C and an attribute vector A for each image. Category classification is posed as a

multi-class problem, thus C satisfies $0 \leq C \leq n_c - 1$, where n_c is the total number of classes. At the attributes level we solve a multi-label problem with $A = (a_1, \dots, a_{n_a})$, $a_i \in \{0, 1\}$, where n_a is the total number of attributes and $a_i = 1$ indicates that the image has the i -th attribute.




Figure 2. Examples of pose detections and respective heatmaps, used during training, with the relevant joints to classify the clothing items in the images.

Current frameworks incorporate attention modules to augment off-the-shelf classification Neural Nets, both in an unsupervised manner [13, 11] or with supervision [4]. In our case, this mechanism acts on a feature combination scheme used in [12] and is supervised by the pose of the human model wearing the clothing item. The pose is extracted using the off-the-shelf pose detector *OpenPose* [1]. Given the joint positions provided by the pose detector, we create a heatmap that localizes the clothing item in the image. In particular, these ground-truth heatmaps, used in training, are generated by placing 2D Gaussian filters on the joint locations and connections as shown in Fig. 2.

The proposed VSAM model architecture is depicted in Fig. 3. Specifically, we regularize the backbone network (VGG-16 [9]) with ground-truth heatmaps with the relevant joints obtained by *OpenPose* (see Fig. 2). This regularization is performed on the feature combination from convolutional blocks 2, 3 and 4 that mix different image resolutions, which are then regularized by the ground-truth heatmap. We also use a spatial attention module (a concatenation of max and average pooling across the channel axis), similarly to [11], to highlight informative regions before the regularization. The convolutional layer after the spatial attention module has a sigmoid activation so that the feature values are in the same interval as the heatmaps values (i.e. $[0, 1]$).

The model is trained by minimizing the cross entropy loss at the category level and the focal loss [3] at the attributes level. The latter has been shown to be effective in preventing excessive negative examples from hindering model training due to the large imbalance of negative/positive examples that arise from annotation sparsity. Finally, for the attention regularizer we use the pixel-wise L2-norm difference between the estimated (\hat{S}) and the ground-truth (S) heatmaps $\sum_u^w \sum_v^h \|\hat{S}(u, v) - S(u, v)\|_2$.

Our model has two hyperparameters, one weights the contribution of the attention branch, and the other, the heatmap fidelity parameter, multiplies the loss term of the pixel-wise difference between \hat{S} and S .

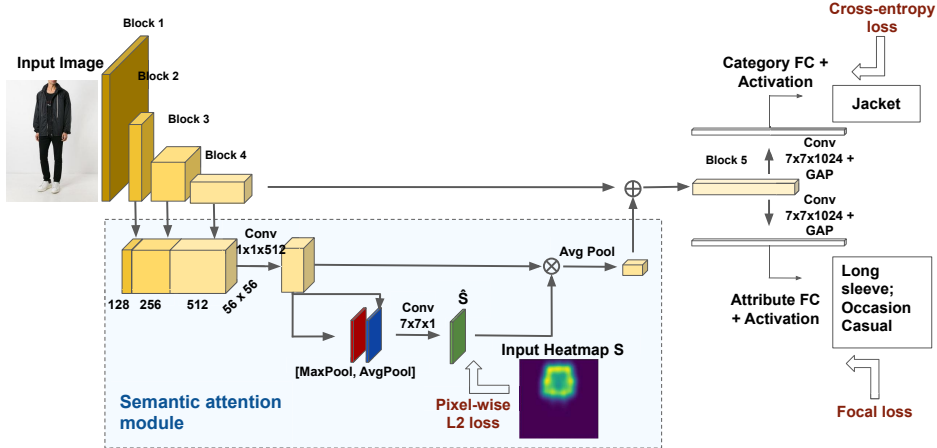


Figure 3. Proposed VSAM architecture with the semantic attention mechanism regularization.

4. Experiments and Results Analysis

4.1. Experimental Setting

Datasets: We use the same (in-house) high-quality dataset from [7], which has approximately 245k images with $n_c = 17$ and $n_a = 53$ and an average number of 1.2 attribute annotations per product. Despite manually curated, some inconsistencies and missing labels naturally arise in this context, leading to a weakly annotated dataset. We train the proposed model (in Fig. 3) on front model images to which we apply OpenPose, and we used a 75% / 25% train/test split ratio.

To provide results on a common benchmark, we also experiment with the widely used DeepFashion dataset [6], containing approximately 289k images, $n_c = 50$ and $n_a = 1000$. Nonetheless, note that, unlike the methods in [6, 10, 5], ours does not use any kind of landmark annotations.

Compared methods: For the in-house dataset, we compare our model with the state-of-the-art model from [7], which had a ResNet-50 as backbone CNN. We also compare our model with [6] that introduced DeepFashion dataset and with the current best-performing model from [5].

Metrics: We focus on the multi-label problem (attributes level) for both datasets. For the in-house dataset we report precision, recall and F1-score at top-k (P@k, R@k, F1@k, where k is the number of ground-truth labels of each product), as well as average precision (AP). For the DeepFashion dataset we follow the same evaluation settings from [6, 5] and report top-k recall for attribute prediction. For all these metrics, the larger value, the better the performance.

All models were run in a single-shot manner (end-to-end) under equal conditions, i.e., for 40 epochs with the same optimizer (Adam with initial learning rate 10^{-4} and decay 10^{-5}), batch size 64 of images resized to 224×224 , and augmentation ratio 0.33. When applied, we used the

Method	Metric			
	P@k	R@k	F1@k	AP
Model from [7]	73.02	75.26	73.33	69.17
VSAM + FL	80.70	81.56	80.63	75.44

Table 1. Experimental results for multi-label classification.

focal loss standard parameters ($\gamma = 2, \alpha = 0.6$).

4.2. Quantitative results

In-house dataset: For categories and considering precision, recall and F1-score, the performance of the proposed method is on par with the performance of the model from [7] that was already very high.

More importantly, Table 1 reports the multi-label performance results at the attributes level for the same dataset. As shown, our VSAM-FL semantic attention and focal loss model outperforms the baseline by a large margin for all metrics.

Robustness to wrong annotations: In the dataset from [7], the *Longsleeved* attribute is not commonly used for Coats and Jackets retrieval, thus these labels are set to zero ($a_{GT_Longsleeved} = 0$) to save annotation effort. To study the impact of positive annotations, we modify 25% of the ground-truth annotations of this attribute to the correct value 1. In Fig. 4 we observe a larger shift of the score’s mass of the proposed method towards higher values compared with the VGG (our model without semantic attention). For example, for a decision threshold of 0.5 the recall of the model with semantic attention is 84% whereas for the VGG is 52% (for 0.4 the recall for the proposed model is 13% higher than for the VGG, and for 0.6 is 25% higher). This suggests that the semantic attention module can have a significant impact in model robustness to wrong annotations which are recurrent in this industry.

Method \ Metric	Texture		Fabric		Shape		Part	
	top-3	top-5	top-3	top-5	top-3	top-5	top-3	top-5
Model from [7]	44.39	53.91	31.82	41.70	39.88	50.51	31.11	40.76
FashionNet [6] [†]	37.46	49.52	39.30	49.84	39.47	48.59	44.13	54.02
Liu et al. [5] [†]	56.30	65.82	43.05	53.64	58.75	67.80	46.47	57.39
VSAM + FL	56.28	65.45	41.73	52.01	55.69	65.40	43.20	53.95

Table 2. Experimental results for attribute classification for the DeepFashion dataset.[‡]

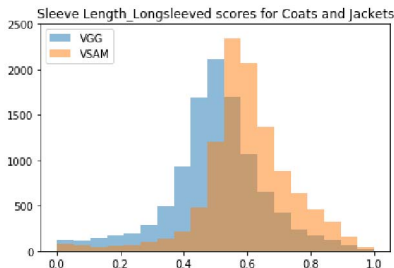


Figure 4. Histograms of the scores for 75% missing label attribute *Longsleeved* for Coats and Jackets categories

DeepFashion: Table 2 reports the results on the DeepFashion dataset. Despite the simpler machinery (approx. 25M parameters vs. 76M of [6]) and less information, as it does not use landmarks, our model achieves nearly the same results as the top performing from the state of the art.

4.3. Qualitative results

The attention regularized model hallucinated good attention heatmaps (correct pose) for test images (feature maps appear to focus on relevant locations), and visualization techniques as saliency maps or CAMs [8] highlight more meaningful regions than for the VGG, as shown in Fig. 1 and in the appendix.

5. Conclusions

Creating complete and consistent datasets seems an unrealistic task, since fully labeling fashion attributes and, specially, spatial landmarks with human annotators would require a tremendous amount of effort and be extremely expensive. As a consequence, taking advantage of *a priori* cues to look for these details in the images appears as a promising strategy. In this work we introduced a fashion classification model, VSAM, whose attention is guided by the pose of the human wearing the clothing items. In spite of its much lower complexity, VSAM outperformed, by a large margin, a previous model in the challenging multi-label scenario for a fashion e-commerce platform dataset (without landmark annotations), and performed on par with

[†] method uses landmarks

[‡] we used the code kindly made available by [5]. We did not include results for Style attributes because we were not able to reproduce them.

the state-of-the-art methods for the DeepFashion dataset that benefit from landmark annotations. Furthermore, the proposed model was robust to wrong annotations and provided more meaningful visualizations and interpretability. The encouraging results suggest that additional gains are attainable by learning attribute specific attention maps.

References

- [1] Z. Cao, T. Simon, S. Wei, and Y. Sheikh. Realtime multi-person 2D pose estimation using part affinity fields. In *CVPR*, 2017.
- [2] P. Gutierrez, P. Sondag, P. Butkovic, M. Lacy, J. Berges, F. Bertrand, and A. Knudson. Deep learning for automated tagging of fashion images. In *ECCV 2018 Workshops*, 2019.
- [3] T. Lin, P. Goyal, R. Girshick, K. He, and P. Dollr. Focal loss for dense object detection. In *ICCV*, 2017.
- [4] D. Linsley, D. Shiebler, S. Eberhardt, and T. Serre. Learning what and where to attend. In *ICLR*, 2019.
- [5] J. Liu and H. Lu. Deep fashion analysis with feature map up-sampling and landmark-driven attention. *ECCV 2018 Workshops*, 2019.
- [6] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang. DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations. In *CVPR*, 2016.
- [7] B. Quintino Ferreira, L. Baía, J. Faria, and R. Sousa. A Unified Model with Structured Output for Fashion Images Classification. In *KDD'18 Workshop on AI for Fashion*, 2018.
- [8] R. Selvaraju, A. Das, R. Vedantam, M. Cogswell, D. Parikh, and D. Batra. Grad-CAM: Why did you say that? Visual Explanations from Deep Networks via Gradient-based Localization. In *ICCV*, 2017.
- [9] K. Simonyan and A. Zisserman. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *ICLR*, 2015.
- [10] W. Wang, Y. Xu, J. Shen, and S.-C. Zhu. Attentive Fashion Grammar Network for Fashion Landmark Detection and Clothing Category Classification. In *CVPR*, 2018.
- [11] S. Woo, J. Park, J. Lee, and I. S. Kweon. CBAM: Convolutional Block Attention Module. In *ECCV*, 2018.
- [12] S. Zhang, G. Wu, J. Costeira, and J. Moura. Understanding Traffic Density from Large-Scale Web Camera Data. In *CVPR*, 2017.
- [13] F. Zhu, H. Li, W. Ouyang, N. Yu, and X. Wang. Learning Spatial Regularization with Image-level Supervisions for Multi-label Image Classification. In *CVPR*, 2017.