

RESEARCH ARTICLE

Attribute-Guided Multi-Level Attention Network for Fine-Grained Fashion Retrieval

LING XIAO^{ID}, (Member, IEEE), AND TOSHIHIKO YAMASAKI^{ID}, (Member, IEEE)

Department of Information and Communication Engineering, Graduate School of Information Science and Technology, The University of Tokyo, Tokyo 113-8656, Japan

Corresponding author: Ling Xiao (ling@cvm.t.u-tokyo.ac.jp)

This paper is supported by Next Generation Artificial Intelligence Research Center of The University of Tokyo. It is also partially financially supported by Institute for AI and Beyond of The University of Tokyo.

ABSTRACT Fine-grained fashion retrieval searches for items that share a similar attribute with the query image. Most existing methods use a pre-trained feature extractor (e.g., ResNet 50) to capture image representations. However, a pre-trained feature backbone is typically trained for image classification and object detection, which are fundamentally different tasks from fine-grained fashion retrieval. Therefore, existing methods suffer from a feature gap problem when directly using the pre-trained backbone for fine-tuning. To solve this problem, we introduce an attribute-guided multi-level attention network (AG-MAN). Specifically, we first enhance the pre-trained feature extractor to capture multi-level image embedding, thereby enriching the low-level features within these representations. Then, we propose a classification scheme where images with the same attribute, albeit with different values, are categorized into the same class. This can further alleviate the feature gap problem by perturbing object-centric feature learning. Moreover, we propose an improved attribute-guided attention module for extracting more accurate attribute-specific representations. Our model consistently outperforms existing attention based methods when assessed on the FashionAI (62.8788% in MAP), DeepFashion (8.9804% in MAP), and Zappos50k datasets (93.32% in Prediction accuracy). Especially, ours improves the most typical ASENet_V2 model by 2.12%, 0.31%, and 0.78% points in FashionAI, DeepFashion, and Zappos50k datasets, respectively. The source code is available in <https://github.com/Dr-LingXiao/AG-MAN>.

INDEX TERMS Attention, fine-grained fashion retrieval, fashion, feature gap.

I. INTRODUCTION

Fashion retrieval is a crucial research topic [1], [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16], [17], especially for fashion recommendations [18], [19], [20], [21], [22]. It aims to learn the overall similarity among different fashion items [6], [23], [24]. Fine-grained fashion retrieval is an emerging topic in the fashion retrieval field. It focuses on learning fine-grained similarity rather than overall similarity, because the similarity between items varies depending on the specific attribute considered in the task. For example, two shirts with a crew neck and a v-neck, respectively, are considered similar in a fashion retrieval task, while they will be treated as different items when doing

fine-grained fashion retrieval under the guidance of neckline design. Moreover, the attribute can be very diverse, thus this research task is more challenging than the general fashion retrieval.

This research topic, while challenging, holds significant importance in the field. Achieving high-quality fine-grained fashion retrieval is crucial not only for meeting the varied and diverse needs of users in the real world but also plays a pivotal role in fashion copyright protection, as highlighted by Martin et al. [25]. This capability is instrumental in identifying items with copied or duplicated designs, which hinges on the nuanced ability to discern intricate similarities among fashion items. The process involves analyzing detailed attributes of clothing, such as texture, pattern, and shape, at a granular level, enabling the detection of subtle differences and similarities that are not apparent at a glance. As the

The associate editor coordinating the review of this manuscript and approving it for publication was Mingbo Zhao^{ID}.

fashion industry continues to grow and evolve, with an ever-increasing volume of designs and the rapid proliferation of knock-offs, the demand for advanced fine-grained similarity analysis technologies becomes more pressing. These technologies not only support the creative integrity of designers by safeguarding their original designs but also enhance consumer experience by providing more accurate and personalized fashion recommendations. Therefore, developing and refining these technologies is essential for advancing both copyright protection efforts and the overall user experience in fashion e-commerce platforms.

Current approaches for fine-grained fashion retrieval are mainly the conditional similarity network [26] and attention networks [27], [28], [29]. The most recent contrastive learning method [30] and clustering method [31] use the attention network [27] as their base network.

However, existing state-of-the-art (SOTA) methods have two problems. (1) They rely on pre-trained Convolutional Neural Network (CNN) backbones that were initially trained for image classification on the ImageNet dataset [32] to extract image representations. The captured image embedding is then processed to calculate the fine-grained similarities. This leads to a feature gap problem due to the distinct nature of the image classification task and fine-grained fashion retrieval task. (2) Existing work adopts high level features for fine-grained fashion similarity learning. Considering the diversity of the attributes, the neglect of low level features will degrade the model performance, especially for some attributes that care about small texture difference.

To address the above-mentioned issue, we present an attribute-guided multi-level attention network (AG-MAN) to improve retrieval accuracy in fine-grained fashion similarity learning. The proposed AG-MAN can extract more discriminative image features. Specifically, we firstly enhance the pre-trained CNN backbone to increase the low-level features contained in image representations. Then, when fine-tuning the pre-trained CNN for extracting image features, we suggest incorporating a classification loss that groups images sharing the same attribute but differing in subclasses into a common category. This can further alleviate the feature gap problem by perturbing for object-centric feature learning. Once improved image representations are attained, we introduce an improved attribute-guided attention module to derive more accurate attribute-specific representations. The proposed AG-MAN consistently outperforms existing attention networks over three datasets. In brief, this paper offers the following key contributions:

- 1) This paper identifies the feature gap problem in SOTA fine-grained fashion retrieval methods and proposes solutions to address it.
- 2) We propose an improved attribute-guided attention module, named attribute-guided attention (AGA).
- 3) The proposed attribute-guided multi-level attention network, AG-MAN, consistently outperforms existing methods across three datasets.

II. RELATED WORK

A. FASHION RETRIEVAL

As fashion recommendations become a trending research topic, fashion retrieval also attracts more attention [5], [23], [24], [33], [34], [35], [36], [37], [38], [39]. Ji et al. [33] leveraged detailed label information from e-commerce websites to improve cross-domain retrieval. This aided the attention mechanism in pinpointing clothing within complex scenes, facilitating in-shop retrieval based on the extracted clothing features. Han et al. [24] employed a bi-directional long short-term memory (Bi-LSTM) network to recommend an item for the given incomplete outfit in the fashion compatibility task. Subsequently, the advancement of the attention mechanism bolstered the models' capacity to understand context and witnessed widespread integration into computer systems [40], [41], [42], [43] and recommendation systems [44], [45], [46]. Wang et al. [47] used a self-guided attention module to autonomously capture clothing features without relying on external landmark annotations. Nevertheless, they primarily concentrated on the broader resemblance of attire, falling short of the requirements for fine-grained retrieval. To address this issue, several fine-grained fashion retrieval methods have been proposed [26], [27], [28], [29]. In [26], they learned an embedding space and used a fixed mask to select relevant dimensions of the embedding space based on a specified attribute. The masked embedding was adopted for final loss calculation. Different from it, Ma et al. [27] proposed to acquire distinct embedding spaces tailored to specific attributes, rather than relying on a single and unified embedding space. To clarify, they evaluated the nuanced similarity within each corresponding feature space. Wan et al. [28] presented a similar network architecture. They believed that the concurrent acquisition and integration of spatial and channel attention features are effective, with the exception of the sequential learning employed within the attention modules. Yan et al. [29] employed iterative learning strategies to this task. Xiao et al. [30] employed a contrastive learning method for solving the multi-view problem [30]. Jiao et al. [31] incorporated different supervisions at both instance and cluster level for more accurate fine-grained fashion retrieval. However, they often overlooked the feature gap issue.

B. ATTENTION MECHANISM

The attention mechanism witnessed a significant traction and has shown remarkable effectiveness in a variety of research domains, encompassing fields such as computer vision [40], [48], [49], [50], [51] and natural language processing [52], [53], [54]. At its core, the attention mechanism can be understood as a method that allocates more weights towards the most relevant parts of the input data.

In the fashion retrieval field, attention mechanisms are also very important because using a fixed region of interest for a specified attribute across all of the images is not optimal due to the diverse nature of fashion images, which often feature complex backgrounds, pose variations, and other factors. Consequently, attention mechanism is commonly

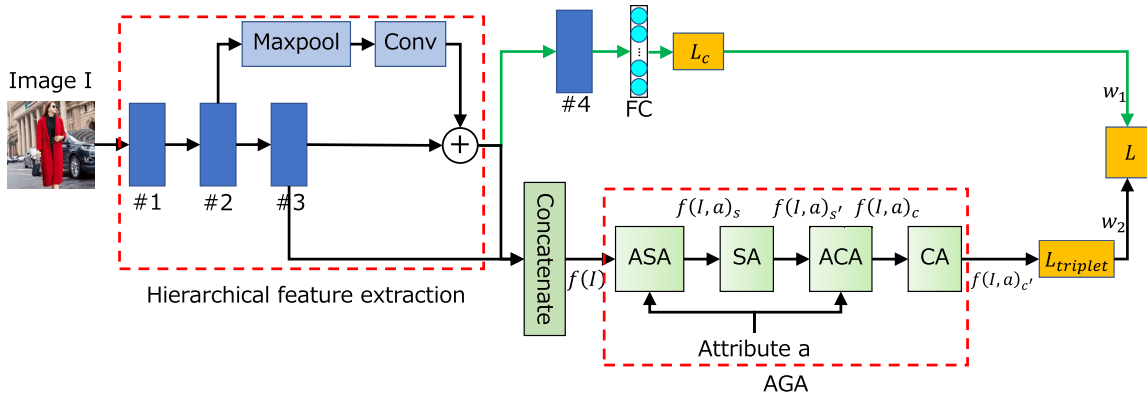


FIGURE 1. Model architecture. #1, #2, #3, and #4 are four blocks in ResNet50 backbone.

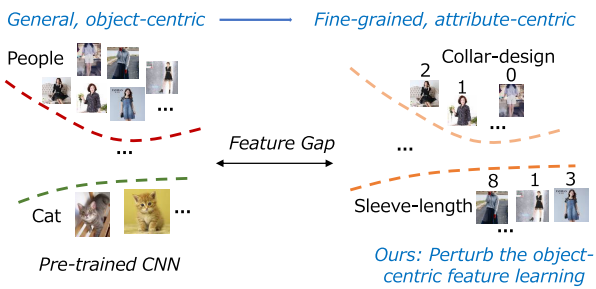


FIGURE 2. The proposed classification branch perturbs the object-centric feature learning in the fine-tuning process by grouping images with the same attribute but different sub-classes into the same class. The number above the image denotes the sub-class under each attribute.

used in the fashion domain [7], [27], [33], [47], [55], [56]. For instance, Ak et al. [7] leveraged prior knowledge of garment construction to pinpoint the precise components of clothing. Nevertheless, their methodology was confined to upper-body clothing, restricting its broader applicability. Wang et al. [47] proposed incorporating channel attention mechanisms into a fully convolutional network. Ji et al. [33] employed attributes to enhance attention modeling. In [27], they introduced two attribute-guided attention mechanisms that used a specific attribute as auxiliary input alongside the given image. The suggested attention modules have successfully discerned patterns associated with the specified attribute under the directive of that attribute. Wan et al. [28] improved existing attribute attention modules by fusing the spatial and channel attention information. Yan et al. [29] enhanced attribute localization accuracy through iterative refinement of the attention module. In contrast to existing methods, we theoretically analyze the feature gap problem in current fine-grained fashion retrieval methods and introduce a more powerful attention network for solving the problem and achieving more accurate fine-grained fashion retrieval.

III. PROPOSED METHOD

A. OVERVIEW

The primary motivation of this research is to tackle the feature gap issue that currently challenges SOTA attention

networks in fine-grained fashion retrieval. This problem arises from the common practice of directly employing pre-trained CNN backbones to extract image representations, which are subsequently processed with attention modules. Our underlying theoretical foundation stems from the notion that leveraging a feature embedding for classification may hinder the learning process for the given task, mainly because it overlooks essential lower-level image features. Therefore, it is imperative to exercise caution when employing pre-trained CNNs, especially when the new research task diverges significantly from traditional image classification.

Figure 1 depicts the architecture of our proposed network. In this research topic, the query attribute can be related to the local information (e.g., neckline-design, texture-related, and lapel-design, etc.) and global information (e.g., style-related, pant-related, and shape-related, etc.). Therefore, capturing multi-level image features is crucial rather than solely focusing on designing complex attention modules. Thus, we first enhance the pre-trained CNN backbone to capture hierarchical image features. We also propose to calculate the logit distribution difference between images with different attribute labels by adding one classification loss L_c . This can disrupt the object-centric feature learning within a pre-trained CNN backbone during the fine-tuning process, thus contributing to the mitigation of the feature gap issue. To extract more accurate attribute-specific embeddings, we introduce an improved attribute-guided attention (AGA) module.

Concretely, for a given query item (I, a) , with I denoting the image and a indicating the attribute, the attribute attended feature vector is $f(I, a) \in \mathbb{R}^c$. This feature vector lies in a continuous vector space \mathbb{R}^c and encapsulates the distinctive features of the specified attribute within the image. Here, c represents the dimensionality of the feature vector. For a pair of fashion images, (I, I') , we compute the embedding similarity between their respective attribute-specific feature vectors, $(f(I, a), f(I', a))$, to quantify their similarity with respect to the same attribute. Furthermore, we extend this similarity assessment across multiple attributes by aggregating the similarity scores for each individual attribute.

Specifically, we calculate the similarity score for each attribute separately and then sum these scores to obtain an overall measure of fine-grained similarity between the two images across all attributes. It is essential to emphasize that each attribute-specific feature vector resides in its respective attribute-specific embedding space. If there are a total of n attributes of interest, we learn n separate embedding spaces simultaneously to capture the distinct characteristics of each attribute within the image.

B. HIERARCHICAL FEATURE EXTRACTION

The concept of hierarchical features has been widely applied across various research domains. Yet, there is a lack of research focused on designing an effective hierarchical feature extraction module specifically for fine-grained fashion retrieval. In this paper, we propose to develop a hierarchical feature extraction module tailored for fine-grained fashion retrieval. We use the widely adopted pre-trained ResNet50 [57] as the backbone network for a fair comparison with existing models [26], [27], [28]. Because fine-grained fashion retrieval needs to deal with multiple attributes, and it cares about different level features, integrating more multi-level features into the image features is essential. To achieve this, we begin by performing an addition operation to combine the $F_2(\cdot)$ (output of intermediate block 2 of ResNet50) and $F_3(\cdot)$, resulting in a fused output denoted as $F_{2,3}(\cdot)$. Next, we concatenate $F_{2,3}(\cdot)$ with $F_3(\cdot)$ and feed this concatenated feature representation into the proposed attribute-guided attention module. This can enhance the multi-level features contained in the image representations that will be processed with the attribute-guided attention module.

Moreover, the pre-trained CNN is designed to learn some object-centric features, which are not what we need in fine-grained fashion retrieval. Thus, we propose adding a classification loss to group images with the same attribute but different sub-classes into the same class (Figure 2). This can boost the learning of more diverse image features. Concretely, the $F_{2,3}(\cdot)$ is passed into block 4 of ResNet50 [57] and a fully connected (FC) layer to calculate a classification loss of different attributes.

C. AGA

The proposed AGA comprises four parts: attribute-guided spatial attention (ASA), spatial attention (SA), attribute-guided channel attention (ACA), and channel attention (CA). It is an improvement of the attribute-aware attention module proposed in [27]. To encode the attribute information, we use a one-hot vector denoted as $a \in \{0, 1\}^n$, where $n \in N$ represents the total number of distinct attributes.

1) ASA

In this research task, learning feature representations for specific, related regions is essential. For instance, to capture attribute-specific features for the neck design, the region

surrounding the neck holds greater significance compared to others regions. Moreover, since fashion images can exhibit substantial variations, such as in poses and scales, relying on a fixed region for a specific attribute across all images is suboptimal. Thus, this paper introduces an ASA module, a mechanism that dynamically focuses on specific areas of an input image based on the relevance with a particular attribute, as illustrated in Figure 3a. To elaborate, with h and w representing the height and width dimensions, we first align the dimensionality of the image feature embedding $f(I) \in \mathbb{R}^{c \times h \times w}$ and the attribute embedding $p(a) \in \mathbb{R}^{c' \times h \times w}$. Formally, we obtain the transformed image feature representation $p(I) \in \mathbb{R}^{c' \times h \times w}$ by applying a convolutional layer $\text{Conv}_{c'}$ comprising $c' \times 1 \times 1$ convolution kernels to the $f(I) \in \mathbb{R}^{c \times h \times w}$, given as

$$p(I) = \tanh(\text{Conv}_{c'}(I)). \quad (1)$$

To obtain attribute representations, we initially transform the query attribute into a vector of dimensionality c' , which is achieved via an embedding layer and an FC layer. Subsequently, we apply self-attention to the attribute embedding. Therefore, the transformed attribute $p(a)$ belongs to the space $\mathbb{R}^{c' \times h \times w}$, where c' denoting the number of channels, and

$$p(a) = \sigma(W_a a) * W_a a, \quad (2)$$

where $W_a \in \mathbb{R}^{c' \times n}$ represents the transformation matrix for the embedding and FC layers, n represents the number of attributes, and σ indicates a sigmoid function. Finally, the spatially attended representation of the image is given by

$$f(I, a)_s = f(I) * W_1 s(p(I) \cdot p(a)), \quad (3)$$

where s denotes a softmax function and W_1 indicates a learned weight matrix. It's important to highlight that we use the sigmoid function for optimizing the attention weights in most cases, a common practice in many attention methods, as described in [40] and [58]. In Eq. 3, the softmax is adopted to further improve the model's ability in distinguishing the differences between features of different attributes.

2) SA

After obtaining $f(I, a)_s$, we first calculate the average and max channel pooling outputs of $f(I, a)_s$, denoted as $\text{Avg}(I_s)$ and $\text{Max}(I_s)$, and concatenate them in the channel dimension to retain more information. They are then passed into a convolution layer and a sigmoid layer to extract an attention weight for the $f(I, a)_s$. SA aims to further improve the positioning ability of $f(I, a)_s$. The processed $f(I, a)_s$ is denoted as $f(I, a)_{s'}$. The detailed operation is given below:

$$W_{SA} = \sigma(\text{Conv}_1([\text{Avg}(I_s), \text{Max}(I_s)])), \quad (4a)$$

$$f(I, a)_{s'} = f(I, a)_s * W_{SA}, \quad (4b)$$

where Conv_1 represents a convolutional layer consisting of a single 1×1 convolution kernel and $[\cdot, \cdot]$ denotes a concatenation operation.

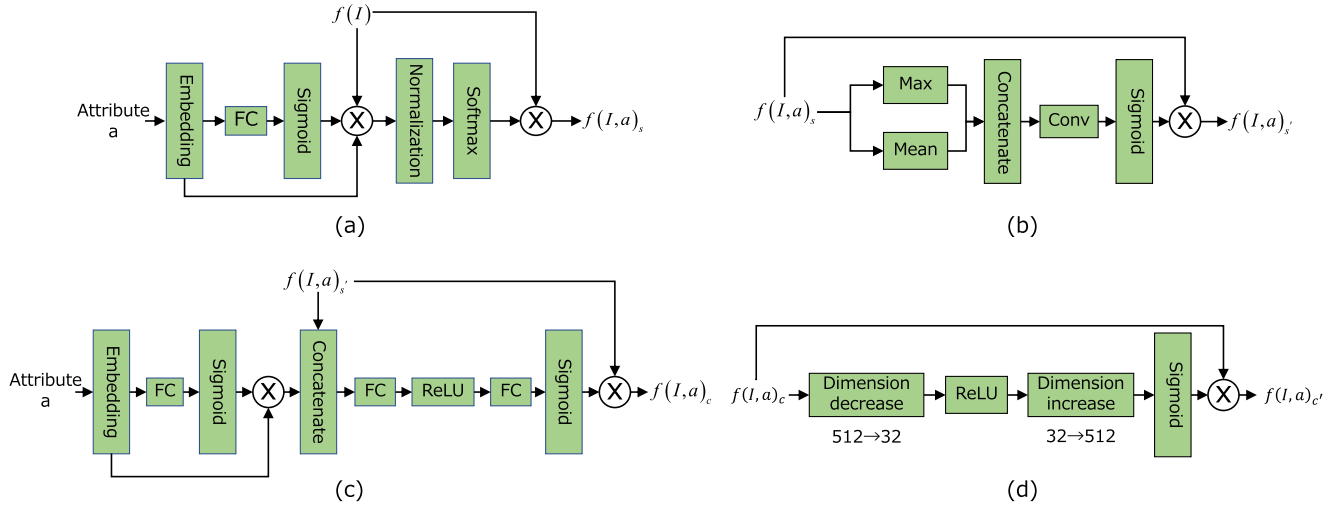


FIGURE 3. Details of (a) ASA, (b) SA, (c) ACA, and (d) CA in the proposed AGA module. The ASA and SA enhance the related region localization while the ACA and CA improve the ability to distinguish between different attributes within the same region.

The proposed SA is very straightforward. We concatenate the $Avg(I_s)$ and $Max(I_s)$ for more diverse information (Figure 3b).

3) ACA

While the ASA and SA dynamically concentrate on particular regions within the image, it's important to note that these identical regions might simultaneously relate to several attributes simultaneously. For instance, attributes such as lapel design and neck design are both associated with the area surrounding the neck. To address this issue, we introduce an ACA mechanism applied to the spatially attended embedding, denoted as $f(I, a)_s'$. The ACA is designed to select dimensions within the spatially attended feature vector that are relevant to the given attribute. In more detail, we start by transforming the attribute into an embedding vector with an embedding layer and a FC layer. This embedding vector has the same dimensionality as $f(I, a)_s'$, and is expressed as

$$q(a) = \sigma(W_{a'} a) * (W_{a'} a), \quad (5)$$

where $W_{a'} \in \mathbb{R}^{c' \times n}$ is a transformation matrix.

Then, the attribute representation $q(a)$ and spatially attended image embedding $f(I, a)_s'$ are combined by an element-wise addition and the resultant feature is passed through two subsequent FC layers to compute the channel attention weights. We also adopt a learned weight matrix W_2 to increase the adaptability. The final output is

$$W_{ACA} = \sigma(W_{c_2}(r(W_{c_1}([q(a), f(I, a)_s']))))), \quad (6a)$$

$$f(I, a)_c = f(I, a)_s' * W_2 * W_{ACA}, \quad (6b)$$

where W_{c_1} and W_{c_2} indicate the weight matrices of two FC layers, respectively, and r represents a ReLU activation function.

4) CA

To enhance the utilization of pertinent channel information derived from the obtained $f(I, a)_c$, we implement a simple dimensionality reduction, succeeded by an expansion operation, aimed at extracting channel-specific attention. This process yields $f(I, a)_c'$. For a more detailed description, please refer to Figure 3d.

D. MODEL LEARNING

The model learning process contains two losses: attribute classification loss and triplet ranking loss. First, the attribute classification loss is computed with

$$L_c = -w_a [y_a \cdot \log \sigma(x_a) + (1 - y_a) \cdot \log(1 - \sigma(x_a))]. \quad (7)$$

Then, with the triplet input (I, I^p, I^n) , we aim to minimize the distance between the embeddings of images with the same specific sub-class, denoted as D_{I, I^p} , and maximize the distance between the embeddings of images with different sub-classes, denoted as D_{I, I^n} . More formally, we define the triplet ranking loss as

$$D_{I, I^p} = \frac{f(I, a)_c' \cdot f(I^p, a)_c'}{\|f(I, a)_c'\| \|f(I^p, a)_c'\|}, \quad (8a)$$

$$D_{I, I^n} = \frac{f(I, a)_c' \cdot f(I^n, a)_c'}{\|f(I, a)_c'\| \|f(I^n, a)_c'\|}, \quad (8b)$$

$$L_{\text{triplet}} = \max\{0, m + D_{I, I^p} - D_{I, I^n}\}, \quad (8c)$$

where m represents the margin, empirically set to be 0.2.

The final loss is a dynamic weighted loss that combines the attribute classification loss L_c and the triplet ranking loss L_{triplet} in an adaptive manner. The primary goal of defining weighted loss is to control the training process and ensure that both losses contribute effectively to learning. Initially, the L_c and L_{triplet} are set to have the same weight coefficient of 0.5. Then they will change following an exponential

TABLE 1. Details of used datasets. They provide a rich set of attributes and their corresponding sub-classes to describe fashion items in detail.

Dataset	Attribute	Sub-classes	Total
FashionAI	skirt-length	short length, knee length, midi length, etc.	6
	sleeve-length	sleeveless, cup sleeves, short sleeves, etc.	9
	coat-length	high waist length, regular length, long length, etc.	8
	pant-length	short pant, mid-length, 3/4 length, etc.	6
	collar-design	shirt collar, peter pan, puritan collar, etc.	5
	lapel-design	notched, collarless, shawl collar, etc.	5
	neckline-design	strapless neck, deep V neckline, straight neck, etc.	11
	neck-design	invisible, turtle neck, ruffle semi-high collar, etc.	5
DeepFashion	texture-related	abstract, animal, bandana, etc.	156
	fabric-related	acid, applique, bead, etc.	218
	shape-related	a-line, ankle, asymmetric, etc.	180
	part-related	arrow collar, back bow, batwing, etc.	216
	style-related	americana, art, athletic, etc.	230
Zappos50k	category	shoes, boots, sandals, etc.	4
	gender	women, men, girls, etc.	4
	heel height	1-4in, 5in&over, flat, etc.	7
	closure	buckle, pull on, slip on, etc.	19




FIGURE 4. The conceptual structure of compared models and ours. The CSN is a conditional similarity network while the the ASENet_v2 and AttnFashion are attention networks in this research field. Ours is also an attention network.

function. To preventing the weight coefficient from becoming too large, which could lead to convergence issues, we add the w_0 and w_1 into the weighted loss. Thus, the final loss becomes

$$L = 0.5 * e^{w_0} L_c + w_0 + 0.5 * e^{w_1} L_{triplet} + w_1, \quad (9)$$

where w_0 and w_1 are learned weight parameters that are initialized as zero.

E. MATERIAL

A summary of the datasets used in this paper are given in Table 1.

1) FASHIONAI

FashionAI is a large-scale fashion dataset with hierarchical attribute annotations for fashion understanding. We use this dataset because of its high-quality attribute annotations. Since the full version of FashionAI has not been publicly released, we use an early version as in [27]. It consists of 180,335 apparel images, where each image is annotated with a fine-grained attribute. There are eight attributes, and each attribute is associated with a list of sub-classes. For instance, the attribute neckline design has 11 corresponding sub-classes such as round neckline and v neckline. We randomly split images into three sets with an 8:1:1 ratio, resulting in 144k/18k/18k images for training/validation/testing. Besides,

TABLE 2. Results on the FashionAI dataset: The best result is bolded, and the second best is underlined. Our model consistently surpasses its counterparts across all attribute types in terms of performance.

Methods	MAP for each attribute (Fashion AI)								MAP ↑
	Skirt-length	Sleeve-length	Coat-length	Pant-length	Collar-design	Lapel-design	Neckline-design	Neck-design	
CSN [26]	59.5852	43.9943	42.8203	62.8021	67.5715	48.1033	41.7291	55.2886	50.8043
ASENet_V2 [27]	<u>64.5731</u>	<u>54.9630</u>	<u>51.7645</u>	<u>64.4973</u>	<u>71.9256</u>	<u>66.7201</u>	<u>60.2929</u>	<u>60.8307</u>	<u>60.7551</u>
AttnFashion [28]	61.9057	45.8526	45.8116	60.6013	70.4638	62.7362	52.0484	59.5861	55.3646
Ours	65.5674	55.0840	55.2060	68.9536	73.7867	68.1758	63.1985	62.0777	62.8788

TABLE 3. Experimental results on DeepFashion dataset. Our model performs the best.

Methods	MAP for each attribute (DeepFashion)					MAP ↑
	Texture-related	Fabric-related	Shape-related	Pant-related	Style-related	
CSN [26]	14.4546	6.5021	11.2555	4.7751	3.4673	8.0747
ASENet_V2 [27]	<u>15.5168</u>	<u>7.1855</u>	<u>11.5066</u>	<u>5.5220</u>	<u>3.6514</u>	<u>8.6704</u>
AttnFashion [28]	11.8423	5.7248	8.6115	4.0904	3.2329	6.8060
Ours	15.6624	7.2746	12.3805	5.8935	3.7555	8.9804

Query images Query attributes

Top-10 retrieval clothes from the datasets

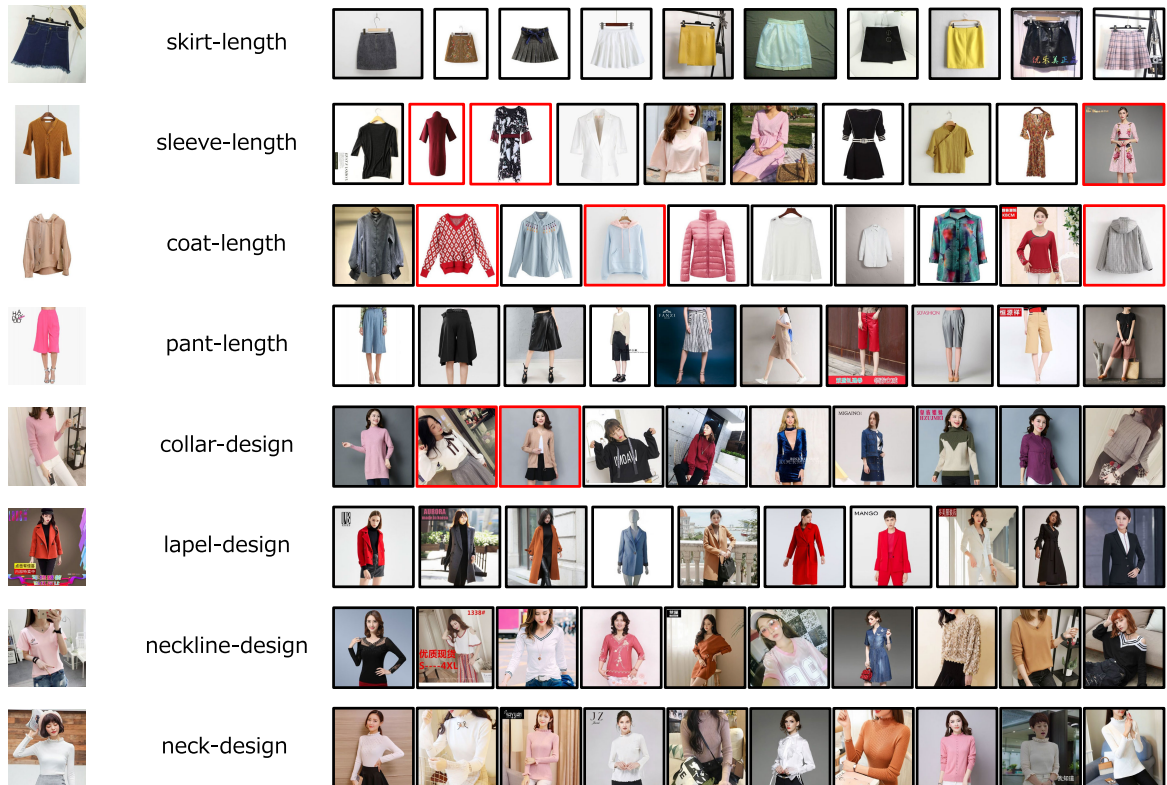


FIGURE 5. Top-10 retrieval examples from the FashionAI dataset, where the images with a black bounding box exhibit the same sub-classes with the query image, while images with a red bounding box have different sub-classes compared to the query image. The majority of the search results within the top 10 generated by our model are precise and correct.

for every epoch, we construct 100k triplets from the training set for model training. Concretely, for a triplet with respect to a specific attribute, we randomly sample two images of

the same corresponding sub-classes as the relevant pair and an image with different sub-classes as the irrelevant one. For validation or test set, 3,600 images are randomly picked

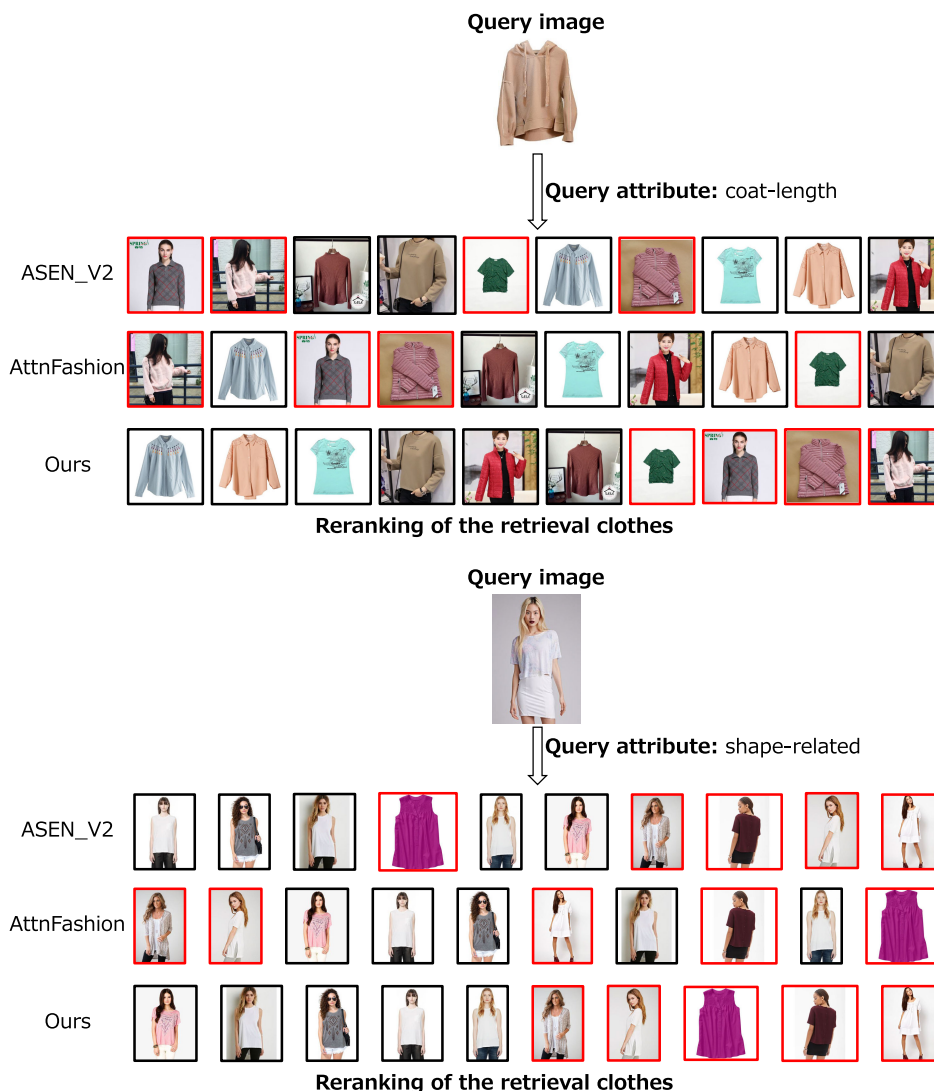


FIGURE 6. The model ranks the candidates based on their similarity to the query image in terms of the specified attribute. The proposed model effectively distinguishes between candidates that share the same sub-classes as the query image (indicated by a black bounding box) and those that differ (marked with a red bounding box). In contrast, other models struggle to achieve this precise ranking, often placing candidates with shared sub-classes lower in the list than they should.

as query images. The remaining images, annotated with the same attribute, are used as the candidate images for retrieval.

2) DEEPFASHION

DeepFashion [2] is a large dataset that consists of four benchmarks for tasks in clothing: category and attribute prediction, in-shop clothes retrieval, fashion landmark detection, and consumer-to-shop clothes retrieval. In our experiments, we use the category and attribute prediction benchmark for the attribute-specific retrieval task.

The category and attribute prediction benchmark contains 289,222 images with six attributes and 1,050 sub-classes, and each image is annotated with several attributes. We randomly split the images into training, validation, and test sets by an

8:1:1 ratio and construct 100k triplets for training. For both the validation and test sets, images are split into query and candidate images at a 1:4 ratio.

3) ZAPPOS50K

Zappos50k [59] is a large shoe dataset consisting of 50,025 images collected from the online shoe and clothing retailer, Zappos.com. For ease of cross-paper comparison, we use the identical split provided by [26]. Specifically, we use 70%/10%/20% images for training/validation/testing. Each image is tagged with four attributes: the type of the shoes, the suggested gender of the shoes, the height of the shoes' heels, and the closing mechanism of the shoes. For each attribute, 200k training, 20k validation, and 40k testing triplets are sampled for model training and evaluation.

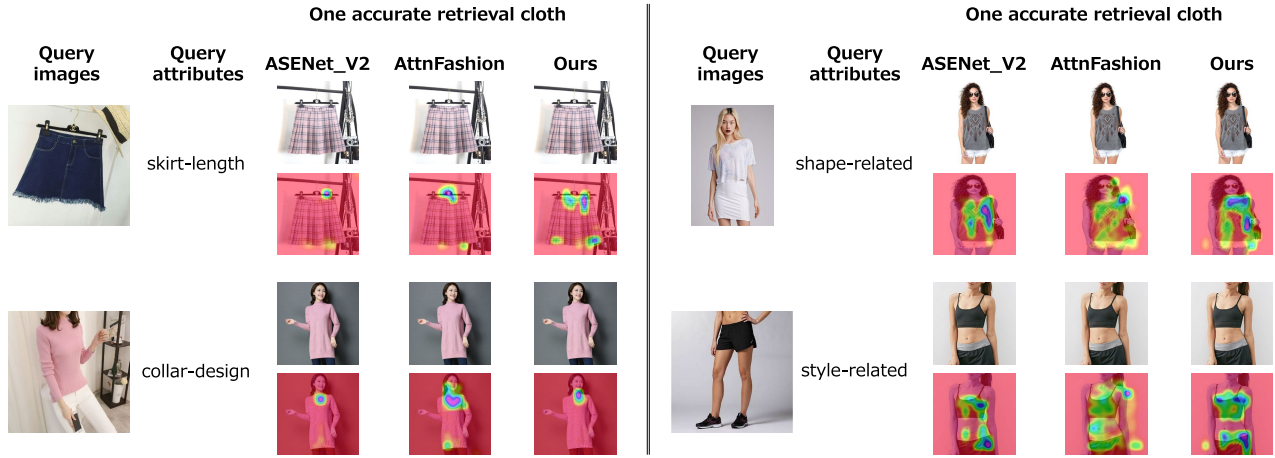


FIGURE 7. Visualization of the spatial attention based on a specified query attribute, depicted above the original retrieved image.

IV. EXPERIMENTAL RESULTS

A. EXPERIMENTAL SETTINGS

To assess the feasibility of the proposed model, we conduct evaluations on the attribute-specific fashion retrieval (ASFR) and triplet relation prediction (TR) tasks. Details of the definition of these two tasks and the experiments are given below.

1) ASFR TASK

When given a fashion image along with a specified attribute, the objective is to identify fashion images from the dataset that share the same attribute as the given reference image. We employ two datasets for this task: the FashionAI dataset [27] and the DeepFashion dataset [2]. Specifically, 1) The FashionAI dataset encompasses eight attributes, each featuring multiple values. These attributes can be categorized into design-related and length-related attributes. 2) In contrast, the DeepFashion dataset comprises five attributes, with approximately 200 values per attribute. These 13 attributes collectively pertain to distinct regions within images and encompass a wide spectrum of image features, ranging from low-level characteristics to high-level image comprehension. Therefore, experiments on these two datasets allow us to assess the proposed model's real-world applicability and effectiveness. Our evaluation metric is the Mean Average Precision (MAP), which is a widely recognized performance measure employed in various retrieval-related tasks [27].

2) TR TASK

With a triplet of $\{(I, I^p, I^n)|a\}$ as input, it aims to predict whether the relevance degree of (I, I^p) surpasses that of (I, I^n) for the specified attribute a . The dataset used in this task is the Zappos50k dataset [59]. We measure performance using prediction accuracy.

The batch size, embedding size, learning rate, learning rate decay step size, and decay rate are 16, 1024, 0.0001, 3, and 0.9, respectively. The training epoch number is 50. The

TABLE 4. Triplet relation prediction on Zappos50k.

Methods	Zappos50k	
	Average loss ↓	Prediction Accuracy(%) ↑
CSN [26]	<u>0.0393</u>	<u>93.14</u>
ASENet_V2 [27]	0.0430	92.54
AttnFashion [28]	0.0664	91.37
Ours	0.0390	93.32

Figure 4 illustrates the conceptual frameworks of the models under comparison, as well as our own model.

B. RESULTS

1) EXPERIMENTAL RESULTS ON THE ASFR TASK

We compare the proposed model with typical attention networks for fine-grained fashion retrieval task to show that our ways to improve the attribute-aware attention module and alleviate feature gap problem are effective. We do not compare our work with Yan et al. [29], due to the absence of released source code, preventing us from re-implementing their approach and replicating the reported results. Tables 2 and 3 present an overview of the comparison models' performance on the FashionAI and DeepFashion datasets, respectively. The tables display the results for various attribute types. It is evident that the proposed network consistently exhibits better performance than other models on both FashionAI and DeepFashion datasets.

Figure 5 presents the top-10 retrieval examples generated by the proposed model when query cloth and attribute are given. The accurate results in the top-10 retrieval examples serve as compelling evidence of the AG-MAN model's efficacy.

We also show some ranking results. Given a query image and 10 candidate clothes, the model evaluates and ranks the candidates based on their similarity to the query image with respect to a specified attribute of interest. This

TABLE 5. Ablation studies on the FashionAI dataset.

Methods	AGA	Hierarchical feature	Classification loss	MAP for each attribute (Fashion AI)								MAP ↑
				Skirt-length	Sleeve-length	Coat-length	Pant-length	Collar-design	Lapel-design	Neckline-design	Neck-design	
Ours	✓	-	-	66.2107	55.3089	52.6219	68.3924	71.2480	66.8022	62.1120	61.9079	61.9679
	✓	✓	-	65.1353	57.0038	53.4465	68.0217	72.2718	67.4483	62.6834	61.9855	62.4804
	✓	-	✓	65.8052	55.4035	54.3290	68.0434	72.4414	65.3819	62.6148	62.5654	62.2951
	✓	✓	✓	65.5674	55.0840	55.2060	68.9536	73.7867	68.1758	63.1985	62.0777	62.8788

TABLE 6. Ablation studies on the DeepFashion dataset.

Methods	AGA	Hierarchical feature	Classification loss	MAP for each attribute (DeepFashion)					MAP ↑
				Texture-related	Fabric-related	Shape-related	Pant-related	Style-related	
Ours	✓	-	-	15.2951	7.1808	12.2927	5.6646	3.6250	8.8069
	✓	✓	-	15.0432	7.2492	12.2757	5.8363	3.8243	8.8406
	✓	-	✓	15.2251	7.4224	12.5469	6.0820	4.0149	9.0512
	✓	✓	✓	15.6624	7.2746	12.3805	5.8935	3.7555	8.9804

TABLE 7. Ablation studies on the proposed AGA module (FashionAI dataset).

Methods	ASA	SA	ACA	CA	MAP for each attribute (Fashion AI)								MAP ↑
					Skirt-length	Sleeve-length	Coat-length	Pant-length	Collar-design	Lapel-design	Neckline-design	Neck-design	
Our AGA	-	✓	✓	✓	59.3651	43.3860	43.5791	62.3133	66.8545	54.7183	40.1253	54.1803	50.7971
	✓	-	✓	✓	65.7707	53.4347	52.7763	66.7755	72.3489	67.9203	62.6388	61.2233	61.6719
	✓	✓	-	✓	64.7258	55.3104	52.6036	68.3390	71.8768	65.7550	60.6618	62.1818	61.4976
	✓	✓	✓	-	65.6344	54.6032	52.5263	65.9612	71.7295	67.6182	61.2087	61.8379	61.4086
	✓	✓	✓	✓	65.5674	55.0840	55.2060	68.9536	73.7867	68.1758	63.1985	62.0777	62.8788

process involves extracting and comparing detailed feature representations of both the query image and each candidate garment to determine their match level in terms of the query attribute. The candidates are then ordered from most to least similar, facilitating the identification of garments that best match the user's search criteria. As we can see from Figure 6, even when other models fail, the proposed method can produce an accurate ranking. It embeds the labeled similar items closer than the labeled different ones.

To gain deeper insights into our proposed model, we visualize the attention maps learned by the proposed ASA. As depicted in Figure 7, these attention maps exhibit notable characteristics. They tend to exhibit higher response values in regions relevant to the specified attribute, while diminishing responses in regions deemed irrelevant. This can demonstrate the ASA's capability to discern which areas are more pertinent for a specific attribute. It's worth noting that attention maps for attributes related to item length exhibit greater complexity compared to design-related attributes. In the case of length-related attributes, multiple regions often exhibit heightened responses. However, our model excels in accurately locating the starting and ending points of a fashion item, enabling it to make informed speculations about its length. Even in scenarios involving the most intricate global attributes, our model consistently outperforms in pinpointing the relevant regions.

2) EXPERIMENTAL RESULTS ON THE TR TASK

Table 4 presents the results for triplet relation prediction on the Zappos50k dataset. Out of the four methods for fine-grained fashion retrieval, our proposed model demonstrates



FIGURE 8. Visualization examples for triplet relation prediction. The 'Close' column is the predicted closer item, while the 'Far' column denotes the predicted farther item.

superior performance compared to other methods. Figure 8 shows some triplet relation prediction examples. As we can see, the proposed method can accurately predict the relation of a triplet input while other models fail.

C. ABLATION STUDIES

Ablation studies were carried out to evaluate the performance changes of the proposed AG-MAN when removing one of the following: hierarchical feature and classification branch. As we can see from Tables 5 and 6, both the hierarchical feature and classification branch can improve the model's performance on the two datasets considered in this paper. The optimal performance is achieved when both components are retained in the model, further underscoring their importance.

We also conduct experiments to evaluate the performance changes when removing ASA, SA, ACA, or CA from the

TABLE 8. Ablation studies on the proposed AGA module (DeepFashion dataset).

Methods	ASA	SA	ACA	CA	MAP for each attribute (DeepFashion)					MAP ↑
					Texture-related	Fabric-related	Shape-related	Pant-related	Style-related	
Our AGA	-	✓	✓	✓	14.6402	6.3866	10.4550	4.7022	3.3938	7.8930
	✓	-	✓	✓	15.4613	7.0831	12.2323	5.4195	3.5369	8.7390
	✓	✓	-	✓	15.1934	6.8899	11.1476	5.4773	3.6428	8.4511
	✓	✓	✓	-	15.1324	7.1874	11.6915	5.6828	3.6776	8.6716
	✓	✓	✓	✓	15.6624	7.2746	12.3805	5.8935	3.7555	8.9804

TABLE 9. Comparison of the proposed AGA and exiting attribute-aware attention modules (FashionAI).

Methods	MAP for each attribute (Fashion AI)								MAP ↑
	Skirt-length	Sleeve-length	Coat-length	Pant-length	Collar-design	Lapel-design	Neckline-design	Neck-design	
ASENet_V2 [27]	<u>64.5731</u>	<u>54.9630</u>	<u>51.7645</u>	<u>64.4973</u>	71.9256	<u>66.7201</u>	<u>60.2929</u>	<u>60.8307</u>	<u>60.7551</u>
AttnFashion [28]	61.9057	45.8526	45.8116	60.6013	70.4638	62.7362	52.0484	59.5861	55.3646
Ours (only AGA)	66.2107	55.3089	52.6219	68.3924	<u>71.2480</u>	66.8022	62.1120	61.9079	61.9679

TABLE 10. Comparison of the AGA and other attribute-aware attention modules (DeepFashion).

Methods	MAP for each attribute (DeepFashion)					MAP ↑
	Texture-related	Fabric-related	Shape-related	Pant-related	Style-related	
ASENet_V2 [27]	15.5168	7.1855	<u>11.5066</u>	<u>5.5220</u>	3.6514	<u>8.6704</u>
AttnFashion [28]	11.8423	5.7248	8.6115	4.0904	3.2329	6.8060
Our AGA	<u>15.2951</u>	<u>7.1808</u>	12.2927	5.6646	<u>3.6250</u>	8.8069

proposed AGA module. The experimental results are shown in Tables 7 and 8. As we can see, when combining all of them, the best performance can be achieved.

D. OTHER EXPERIMENTS

To verify that our AGA is an improvement of the one in [27], and better than the attribute-aware attention modules proposed in other methods, we conduct experiments on both the FashionAI and DeepFashion datasets. As we can see from Tables 9 and 10, the proposed AGA is more effective compared with exiting attribute-specific attention modules. It shows that by adding the SA and CA module, they obtain a good coupling effect.

E. FAILURE ANALYSIS

Figure 9 shows some failure cases of our model. First, our model may incorrectly evaluate on some items that share very small differences with the query image, as seen in texture-related fashion retrieval in the figure. For certain attributes related to length, some fashion images are labeled with varying lengths that are so subtle that even human observers find it challenging to discern the differences accurately. Moreover, real-world datasets include pure item and try-on images that are taken from different viewpoints, making it difficult to distinguish the subtle difference in length between these two kinds of images. Finally, the images with occlusions also add some difficulty. In the future, we aim to address these issues mentioned above.

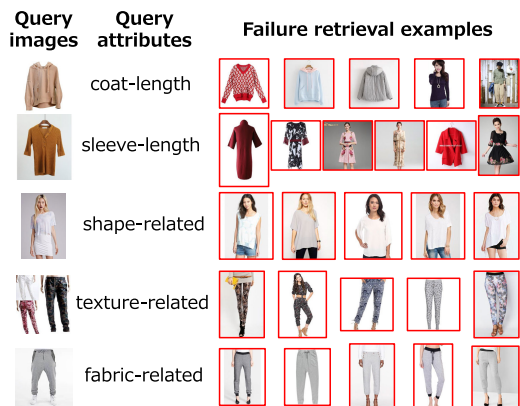


FIGURE 9. Some retrieval failures of the proposed method.

V. LIMITATIONS AND FUTURE WORKS

The proposed method is not scalable. When new attributes come, we need to train the whole model again. Scalable fine-grained fashion retrieval by designing novel continual learning and Large Language Model prompt techniques would be a promising research direction.

VI. CONCLUSION

This paper introduced a novel approach called the AG-MAN to address the feature gap problem and proposed a better attribute-aware attention module. Concretely, the AG-MAN modified the pre-trained CNN backbone to

capture multi-level feature representations for addressing the feature gap problem. We also proposed a method to disturb the object-centric feature learning to further alleviate this problem. Moreover, we proposed a more powerful attribute-guided attention module for extracting more discriminative representations. AG-MAN excelled in achieving precise attribute positioning and extracting highly distinctive features, while being guided by the specified attribute. Extensive experiments with the FashionAI, DeepFashion, and Zappos50k datasets showed that our AG-MAN surpasses existing attention networks for fine-grained fashion retrieval. Our AG-MAN exhibits significant promise for application in various attribute-guided recognition tasks.

REFERENCES

- [1] J. Huang, R. Feris, Q. Chen, and S. Yan, "Cross-domain image retrieval with a dual attribute-aware ranking network," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 1062–1070.
- [2] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang, "DeepFashion: Powering robust clothes recognition and retrieval with rich annotations," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 1096–1104.
- [3] X. Liang, L. Lin, W. Yang, P. Luo, J. Huang, and S. Yan, "Clothes co-parsing via joint image segmentation and labeling with application to clothing retrieval," *IEEE Trans. Multimedia*, vol. 18, no. 6, pp. 1175–1186, Jun. 2016.
- [4] Z. Chen, Z. Xu, Y. Zhang, and X. Gu, "Query-free clothing retrieval via implicit relevance feedback," *IEEE Trans. Multimedia*, vol. 20, no. 8, pp. 2126–2137, Aug. 2018.
- [5] B. Gajic and R. Baldrich, "Cross-domain fashion image retrieval," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1950–19502.
- [6] X. Gu, Y. Wong, L. Shou, P. Peng, G. Chen, and M. S. Kankanhalli, "Multi-modal and multi-domain embedding learning for fashion retrieval and analysis," *IEEE Trans. Multimedia*, vol. 21, no. 6, pp. 1524–1537, Jun. 2019.
- [7] K. E. Ak, J. H. Lim, J. Y. Tham, and A. A. Kassim, "Efficient multi-attribute similarity learning towards attribute-based fashion search," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Mar. 2018, pp. 1671–1679.
- [8] Z. Kuang, Y. Gao, G. Li, P. Luo, Y. Chen, L. Lin, and W. Zhang, "Fashion retrieval via graph reasoning networks on a similarity pyramid," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 3066–3075.
- [9] P. Jing, S. Ye, L. Nie, J. Liu, and Y. Su, "Low-rank regularized multi-representation learning for fashion compatibility prediction," *IEEE Trans. Multimedia*, vol. 22, no. 6, pp. 1555–1566, Jun. 2020.
- [10] P. Jing, J. Zhang, L. Nie, S. Ye, J. Liu, and Y. Su, "Tripartite graph regularized latent low-rank representation for fashion compatibility prediction," *IEEE Trans. Multimedia*, vol. 24, pp. 1277–1287, 2022.
- [11] A. Liu, T. Zhang, D. Song, W. Li, and M. Zhou, "FRSFN: A semantic fusion network for practical fashion retrieval," *Multimedia Tools Appl.*, vol. 80, pp. 17169–17181, Nov. 2021.
- [12] A. Baldrati, M. Bertini, T. Uricchio, and A. Del Bimbo, "Conditioned image retrieval for fashion using contrastive learning and CLIP-based features," in *Proc. ACM Multimedia Asia*, Dec. 2021, pp. 1–5.
- [13] C. Ning, Y. Di, and L. Menglu, "Survey on clothing image retrieval with cross-domain," *Complex Intell. Syst.*, vol. 8, no. 6, pp. 5531–5544, Dec. 2022.
- [14] E. Dodds, J. Culpepper, and G. Srivastava, "Training and challenging models for text-guided fashion image retrieval," 2022, *arXiv:2204.11004*.
- [15] S. Goenka, Z. Zheng, A. Jaiswal, R. Chada, Y. Wu, V. Hedau, and P. Natarajan, "FashionVLP: Vision language transformer for fashion retrieval with feedback," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Chad, Jun. 2022, pp. 14085–14095.
- [16] L. Xiao and T. Yamasaki, "Sat: Self-adaptive training for fashion compatibility prediction," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Oct. 2022, pp. 2431–2435.
- [17] L. Liao, X. He, B. Zhao, C.-W. Ngo, and T.-S. Chua, "Interpretable multimodal retrieval for fashion products," in *Proc. 26th ACM Int. Conf. Multimedia*, Oct. 2018, pp. 1571–1579.
- [18] X. Zhang, J. Jia, K. Gao, Y. Zhang, D. Zhang, J. Li, and Q. Tian, "Trip outfits advisor: Location-oriented clothing recommendation," *IEEE Trans. Multimedia*, vol. 19, no. 11, pp. 2533–2544, Nov. 2017.
- [19] X. Liu, Y. Sun, Z. Liu, and D. Lin, "Learning diverse fashion collocation by neural graph filtering," *IEEE Trans. Multimedia*, vol. 23, pp. 2894–2901, 2021.
- [20] H. Zhan, J. Lin, K. E. Ak, B. Shi, L.-Y. Duan, and A. C. Kot, "A³-FKG: Attentive attribute-aware fashion knowledge graph for outfit preference prediction," *IEEE Trans. Multimedia*, vol. 24, pp. 819–831, 2022.
- [21] Y. Ding, Y. Ma, W. K. Wong, and T.-S. Chua, "Modeling instant user intent and content-level transition for sequential fashion recommendation," *IEEE Trans. Multimedia*, vol. 24, pp. 2687–2700, 2022.
- [22] Z. Lu, Y. Hu, C. Yu, Y. Jiang, Y. Chen, and B. Zeng, "Personalized fashion recommendation with discrete content-based tensor factorization," *IEEE Trans. Multimedia*, vol. 25, pp. 5053–5064, 2022.
- [23] B. Zhao, J. Feng, X. Wu, and S. Yan, "Memory-augmented attribute manipulation networks for interactive fashion search," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 6156–6164.
- [24] X. Han, Z. Wu, Y.-G. Jiang, and L. S. Davis, "Learning fashion compatibility with bidirectional LSTMs," in *Proc. 25th ACM Int. Conf. Multimedia*, Oct. 2017, pp. 1078–1086.
- [25] T. Martin, "Fashion law needs custom tailored protection for designs," *Univ. Baltimore Law Rev.*, vol. 48, no. 3, p. 453, 2019.
- [26] A. Veit, S. Belongie, and T. Karalestos, "Conditional similarity networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1781–1789.
- [27] Z. Ma, J. Dong, Z. Long, Y. Zhang, Y. He, H. Xue, and S. Ji, "Fine-grained fashion similarity learning by attribute-specific embedding network," in *Proc. AAAI Conf. Artif. Intell.*, vol. 34, no. 7, 2020, pp. 11741–11748.
- [28] Y. Wan, K. Yan, C. Yan, and B. Zhang, "Learning attribute-guided fashion similarity with spatial and channel attention," *J. Experim. Theor. Artif. Intell.*, vol. 34, no. 4, pp. 1–17, Aug. 2022.
- [29] C. Yan, K. Yan, Y. Zhang, Y. Wan, and D. Zhu, "Attribute-guided fashion image retrieval by iterative similarity learning," in *Proc. IEEE Int. Conf. Multimedia Expo. (ICME)*, Jul. 2022, pp. 1–6.
- [30] L. Xiao, X. Zhang, and T. Yamasaki, "Toward a more robust fine-grained fashion retrieval," in *Proc. IEEE 6th Int. Conf. Multimedia Inf. Process. Retr. (MIPR)*, Sep. 2023, pp. 1–4.
- [31] Y. Jiao, N. Xie, Y. Gao, C. Wang, and Y. Sun, "Fine-grained fashion representation learning by online deep clustering," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2022, pp. 19–35.
- [32] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A large-scale hierarchical image database," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2009, pp. 248–255.
- [33] X. Ji, W. Wang, M. Zhang, and Y. Yang, "Cross-domain image retrieval with attention modeling," in *Proc. 25th ACM Int. Conf. Multimedia*, Oct. 2017, pp. 1654–1662.
- [34] S. Ibrahim, N. van Noord, Z. Geradts, and M. Worring, "Deep metric learning for cross-domain fashion instance retrieval," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. Workshop (ICCVW)*, Oct. 2019, pp. 3165–3168.
- [35] V. Sharma, N. Murray, D. Larlus, M. S. Sarfraz, R. Stiefel, and G. Csurka, "Unsupervised meta-domain adaptation for fashion retrieval," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Jan. 2021, pp. 1347–1356.
- [36] X. Liu, J. Li, J. Wang, and Z. Liu, "MMFashion: An open-source toolbox for visual fashion analysis," in *Proc. 29th ACM Int. Conf. Multimedia*, Oct. 2021, pp. 3755–3758.
- [37] R. Naka, M. Katsurui, K. Yanagi, and R. Goto, "Fashion style-aware embeddings for clothing image retrieval," in *Proc. Int. Conf. Multimedia Retr.*, Jun. 2022, pp. 49–53.
- [38] S. M. Islam, S. Joardar, and A. A. Sekh, "DSSN: Dual shallow Siamese network for fashion image retrieval," *Multimedia Tools Appl.*, vol. 82, no. 11, pp. 16501–16517, May 2023.
- [39] D. Suzuki, G. Irie, and K. Aizawa, "Text-to-image fashion retrieval with fabric textures," in *Proc. ACM Int. Conf. Multimedia Retr.*, Jun. 2023, pp. 525–529.
- [40] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, "CBAM: Convolutional block attention module," in *Proc. Eur. Conf. Comput. Vis.*, Sep. 2018, pp. 3–19.

- [41] T. Qiao, J. Dong, and D. Xu, "Exploring human-like attention supervision in visual question answering," in *Proc. AAAI Conf. Artif. Intell. (AAAI)*, vol. 32, no. 1, 2018, pp. 7300–7307.
- [42] C. R. Chen, Q. Fan, and R. Panda, "CrossViT: Cross-attention multi-scale vision transformer for image classification," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 347–356.
- [43] A. Vaswani, P. Ramachandran, A. Srinivas, N. Parmar, B. Hechtman, and J. Shlens, "Scaling local self-attention for parameter efficient visual backbones," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 12889–12899.
- [44] G. Zhou, X. Zhu, C. Song, Y. Fan, H. Zhu, X. Ma, Y. Yan, J. Jin, H. Li, and K. Gai, "Deep interest network for click-through rate prediction," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2018, pp. 1059–1068.
- [45] Y. Tay, A. T. Luu, and S. C. Hui, "Multi-pointer co-attention networks for recommendation," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2018, pp. 2309–2318.
- [46] H. Liu, L. Guo, P. Li, P. Zhao, and X. Wu, "Collaborative filtering with a deep adversarial and attention network for cross-domain recommendation," *Inf. Sci.*, vol. 565, pp. 370–389, Jul. 2021.
- [47] Z. Wang, Y. Gu, Y. Zhang, J. Zhou, and X. Gu, "Clothing retrieval with visual attention model," in *Proc. IEEE Vis. Commun. Image Process. (VCIP)*, Dec. 2017, pp. 1–4.
- [48] H. Zhao, J. Jia, and V. Koltun, "Exploring self-attention for image recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 10073–10082.
- [49] Z. Niu, G. Zhong, and H. Yu, "A review on the attention mechanism of deep learning," *Neurocomputing*, vol. 452, pp. 48–62, Sep. 2021.
- [50] Y. Li, K. Guo, Y. Lu, and L. Liu, "Cropping and attention based approach for masked face recognition," *Appl. Intell.*, vol. 51, no. 5, pp. 3012–3025, May 2021.
- [51] Q. Hou, D. Zhou, and J. Feng, "Coordinate attention for efficient mobile network design," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 13708–13717.
- [52] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. Adv. Neural Inform. Process. Syst.*, 2017, pp. 5998–6008.
- [53] A. Galassi, M. Lippi, and P. Torroni, "Attention in natural language processing," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 32, no. 10, pp. 4291–4308, Oct. 2021.
- [54] E. Sood, S. Tannert, P. Müller, and A. Bulling, "Improving natural language processing tasks with human gaze-guided neural attention," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 6327–6341.
- [55] X. Han, Z. Wu, P. X. Huang, X. Zhang, M. Zhu, Y. Li, Y. Zhao, and L. S. Davis, "Automatic spatially-aware fashion concept discovery," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 1472–1480.
- [56] K. E. Ak, A. A. Kassim, J. H. Lim, and J. Y. Tham, "Learning attribute representations with localization for flexible fashion search," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7708–7717.
- [57] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [58] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7132–7141.
- [59] A. Yu and K. Grauman, "Fine-grained visual comparisons with local learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 192–199.



LING XIAO (Member, IEEE) received the Ph.D. degree from Huazhong University of Science and Technology, in 2020. She is currently a Project Assistant Professor with the Department of Information and Communication Engineering, Graduate School of Information Science and Technology, The University of Tokyo. She was a Visiting Scholar with The University of Queensland, from October 2018 to November 2019. Her current research interests include computer vision, self-supervised learning, recommendations, and artificial intelligence.



TOSHIHIKO YAMASAKI (Member, IEEE) received the Ph.D. degree from The University of Tokyo. He is currently a Professor with the Department of Information and Communication Engineering, Graduate School of Information Science and Technology, The University of Tokyo. His current research interests include attractiveness computing based on multimedia big data analysis, computer vision, pattern recognition, and machine learning.

• • •