

# PAI-BPR: Personalized Outfit Recommendation Scheme with Attribute-wise Interpretability

Dikshant Sagar\*, Jatin Garg\*, Prarthana Kansal\*, Sejal Bhalla\*, Rajiv Ratn Shah  
IIT Delhi, India  
{dikshant17338, jatin17345, prarthana17357, sejal17100, rajivrtn}@iiitd.ac.in

Yi Yu  
NII Japan  
yiyu@nii.ac.jp

**Abstract**—Fashion is an important part of human experience. Events such as interviews, meetings, marriages, etc. are often based on clothing styles. The rise in fashion industry and its effect on social influencing have made outfit compatibility a need. Thus, it necessitates an outfit compatibility model to aid people in clothing recommendation. However, due to the highly subjective nature of compatibility, it is necessary to account for personalization. Our paper devises an attribute-wise interpretable compatibility scheme with personal preference modelling which captures user-item interaction along with general item-item interaction. Our work solves the problem of interpretability in clothing matching by locating the discordant and harmonious attributes between fashion items. Extensive experiment results on IQON3000, a publicly available real-world dataset, verify the effectiveness of the proposed model.

**Keywords**-fashion analysis, personalized compatibility modeling, multimodal, interpretability

## I. INTRODUCTION

The rise in fashion industry and its effect on social influencing has made outfit compatibility a need. Modern consumers choose to wear harmonious outfits (see Figure 1), and thus, it has gained increasing research attention. However, due to the highly subjective nature of compatibility, it is desirable to account for user’s personal preference in compatibility modelling. Despite its subjective nature, only a limited amount of work has been done to personalise the fashion recommendation process. For instance, most existing efforts such as the fashionability prediction [1], [2], clothing retrieval [3], [4], [5], and compatibility modeling [6], [7], [8] focus on the numerical compatibility modeling between fashion items with advanced neural networks, and hence suffer from poor interpretation. Some recent works have focused on user input in the form of image or text to achieve personalization [4], [9]. Some have used Computer Vision algorithms to develop a real-time fashion recommendation system consistent with the latest fashion trends [10]. The algorithm analyzes low-level features such as shape of cloth and recommends a suitable match from existing database which is either the user’s wardrobe or websites from internet.

Fashion recommendation systems hence entail two main requirements: 1) the *compatibility* of items in the recom-



Figure 1. Example of fashionable outfits from Fashionista.<sup>1</sup>

mended fashion outfits, 2) the *personalization* in the recommendation process. While compatibility is a measurement of how harmonious a set of items is, personalization takes user’s personal choices into account. Furthermore, in the real-world clothing matching scenarios, users may not only want to know whether the given fashion items are compatible or not, but also expect to understand the underlying reason which leads to discordance in an outfit. This leads to another significant requirement of a recommendation system, *i.e.*, *interpretability* of the model. The recent proliferation of online fashion communities like IQON<sup>2</sup> or Chictopia<sup>3</sup>, has enabled us to identify fashion choices particular to a given user. Additionally, with a suitable representation of these fashion items, it is possible to map the interactions between semantically meaningful fashion elements, hereafter referred to as *attributes*. For instance, we can infer that a user prefers, say {*pear-shaped top*, *A-line skirt*} over {*pear-shaped top*, *midi skirt*} from her past outfit preferences (see Figure 3). Apart from enhancing the recommendation, this additional information helps the user comprehend the role of attribute interactions, pinpoint the discordant attributes and harmonize the outfit accordingly.

In order to learn effective representations of fashion items, several existing works deploy Deep Learning based

<sup>1</sup><https://fashionista.com/>

<sup>2</sup><http://www.iqon.jp/>

<sup>3</sup><http://www.chictopia.com/>

\* These authors contributed equally.

methods to define and measure compatibility between them [11], [12]. However, the “uninterpretability” of Deep Neural Networks has always been a limiting factor for applications requiring explanations of the learned representation [13]. The approach fails to explicitly map each dimension of the learned representation to the intuitive semantic aspects of fashion items. Although a few researchers have taken a step in enhancing the interpretability by modeling the outfit compatibility at an attribute-level (e.g., color and texture) [14], their result suffers from the lack of attributes. For example, Wang et al. [15] limit their analysis to low-level features (such as color and shape) and high-level features (such as category and style) without exploring them further and pinpointing the exact type of the feature (e.g., for the *shape* feature, the types are *hourglass-shaped*, *pear-shaped*, etc). In some other works like PAICM [16], the issue of interpretability is addressed to a great extent but at the cost of personalization, yielding a poor performance.

To comprehensively tackle the aforementioned issues, we aim to devise an outfit recommendation scheme termed PAI-BPR (Personalized Attribute-wise Interpretable - BPR), as shown in Figure 2, which models personalized compatibility and explains the underlying reasons for the incompatibility of a given set of items. We follow an attribute-wise interpretable scheme to explain compatibility since attributes are the most intuitive semantic cues to characterize fashion items [16]. Without loss of generality, we focus on the problem of matching a top and a bottom. To address the same, the scheme learns the semantic attribute representation of the fashion items through deep neural networks for general compatibility modeling. To reduce the variability of the results with different datasets, the attribute classification models are pre-trained on a dataset which is richly annotated with attributes, see Section IV-A2 for more details. On the other hand, the personal preference modeling focuses on extracting the latent preference factor based on the multimodal data (textual description and image) of fashion items and hence captures the user-item interactions comprehensively. Ultimately, the model jointly regularizes the general compatibility and personal preference modeling by utilizing the Bayesian Personalized Ranking (BPR) framework [17].

We summarise our main contributions as follows:

- To the best of our knowledge, this is the first attempt to comprehensively accomplish the automatic outfit recommendation task by answering the two essential questions of compatibility determination and discordant component identification in accordance with user preference.
- Considering that both textual description and the image associated with a fashion item contribute significant information about user preferences, we build an attribute-wise interpretable personal preference modeling scheme by integrating the multimodal data of fashion items.
- Extensive experiments conducted on a real-world dataset verify the effectiveness of the proposed model, showing a clear advantage of using a fine-grained attribute representation for fashion item images.

The remainder of this paper is structured as follows. Section II briefly reviews the related work. In Section III, we detail the proposed model. We present the experimental results and analyses in Section IV, followed by our concluding remarks and future work in Section V.

## II. RELATED WORK

### A. Fashion Analyses

In recent years, the huge economic value of the fashion market has motivated research in the fashion domain, such as the fashion trends prediction [18], fashionability prediction [1], [2], clothing retrieval [3], [4], clothing parsing [19], attribute learning [16], [20] and compatibility modeling [6], [7]. For example, McAuley et al. [8] proposed a general framework to extract visual features which model the human visual preference for a given pair of objects based on the Amazon real-world co-purchase dataset while Song et al. [11] investigated the problem of complementary fashion item matching with a multimodal dataset collected from Polyvore<sup>4</sup>. While most works focus on learning the compatibility between a pair of items, Stile [21] stood out as an end-to-end intelligent fashion consultant system which generates stylish outfits consisting of a top, bottom, footwear and accessories, for a given item. To further boost the performance of fashion item recommendation, Zhou et al. [22] proposed a novel approach sensitive to the current fashion trends.

### B. Personalized Fashion Recommendation

A limited amount of work has been done to personalise the fashion recommendation process. Some recent works have focused on user input in the form of image or text to achieve personalization [4], [9]. Wang et al. [23] incorporated user factor to some extent by modelling behavior of similar interest individuals and analysing retail transaction data to finally improve fashion recommendations. In the recent years, the focus of the research has shifted to explore the breadth of factors which makes a recommendation personalised. Hidayati et al. [24] extracted fashion knowledge from social networking websites like Instagram to suggest items according to the body shape of a person. A similar work [25] took features like gender, sex, body shape, etc of an individual with the occasion associated with dressing to recommend fashion items. However, the state-of-the-art system for personalised clothing matching which takes the most crucial factor of a user’s personal preferences into account, is GP-BPR [26], which models

<sup>4</sup>Polyvore has been acquired by the global fashion platform Ssense in 2018

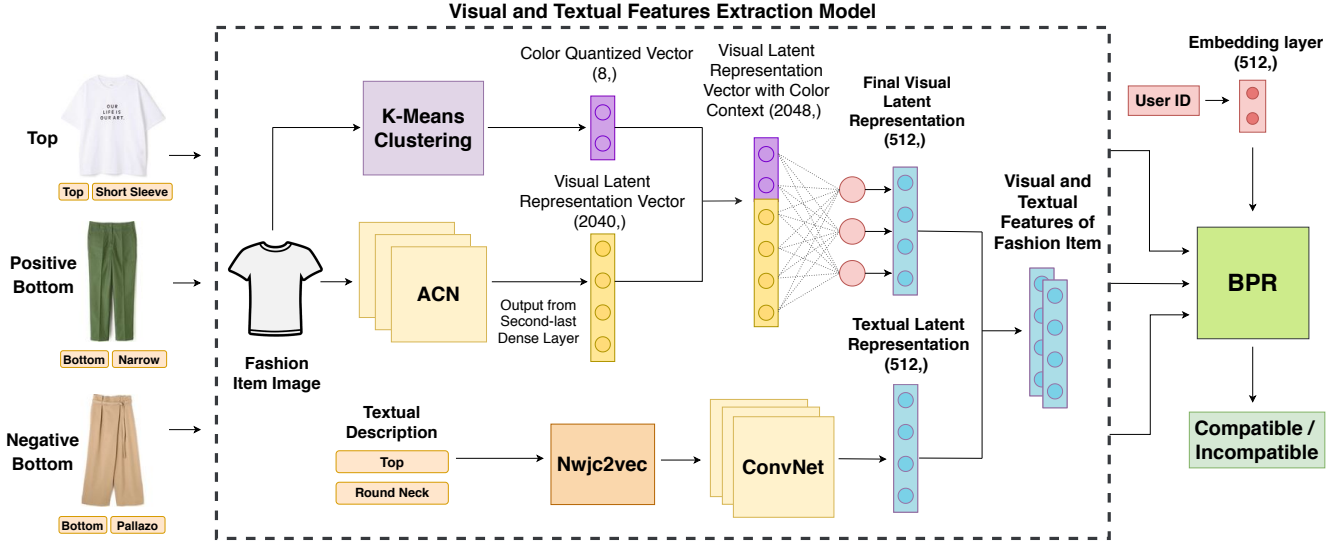


Figure 2. Feature extraction and training schematic. We obtain the semantic attribute representations for each fashion item via the pre-trained attribute classification network (ACN). The textual latent embedding is learned with Nwjc2vec followed by a Text-CNN. The visual and contextual latent embeddings are integrated with the user preferences by the BPR framework.

visual and contextual features of clothing items along with personal preferences. We attempt to achieve personalization by taking inspiration from this work and supplementing it with interpretable results which increases the performance of our model while establishing the credibility of our work.

### C. Methodology Review

Rising interest in the fashion domain has led to the exploration of a wide variety of implementations of recommendation systems. Chao et al. [10] proposed a real-time fashion recommendation system consistent with the latest trends using computer vision algorithms. The algorithm analyzes low-level features such as shape of the cloth a person is wearing and recommends a suitable match from existing database i.e. either the user’s wardrobe or various websites from internet. Adding another dimension of occasion specificity, certain works recommend suitable outfits to a user according to various attributes of clothing, person and the occasion using neural networks like Bi-LSTM and CNN [25], [27]. To study finer and semantic details of the cloth such as the fabric, texture etc., Al-Halah et al. [18] used CNN representations. They predicted and forecast future trends using images and consumer data. Another CNN implementation to mitigate the problem of two-piece clothing matching compatible with the latest fashion trends was proposed by Zhou et al. [22]. They combined the perception and reasoning models, and built two parallel CNNs to enable the system to capture the clothing features; one for the upper-body clothes and the other for the lower-body clothes. A completely new aspect of fashion recommendation was interpreted and studied by Piazza et al. [28] using Factorization Machines (FM). They investigated the impact of users’ emotions on the fashion

products and integrated the user’s mood and emotions in the model to make the recommendation system suited to the user’s preferences. Overall, most existing efforts focus on the latent user-item interactions to tackle the personalized recommendation problems. Beyond that, in this work, we aim to fulfil the task of personalized clothing matching by taking both user-item and item-item interactions into account, along with addressing interpretability.

## III. METHODOLOGY

In this section, we first formally describe the research problem and its formulation; then detail the proposed model and approach (see Figure 2).

### A. Problem Formulation

Assume we have a set of users  $U = \{u_1, u_2, \dots, u_M\}$ , a set of tops  $T = \{t_1, t_2, \dots, t_{N_t}\}$  and a set of bottoms  $B = \{b_1, b_2, \dots, b_{N_b}\}$ , where  $M$ ,  $N_t$  and  $N_b$  denote the total numbers of users, tops and bottoms, respectively. Each user  $u_m$  is associated with a set of historically composed top-bottom pairs  $O_m = \{(t_{i_1^m}, b_{j_1^m}), (t_{i_2^m}, b_{j_2^m}), \dots, (t_{i_{N_m}^m}, b_{j_{N_m}^m})\}$ , where  $i_k^m \in [1, 2, \dots, N_t]$  and  $j_k^m \in [1, 2, \dots, N_b]$  refer to the index of the top and bottom. Further, each item  $t_i(b_j)$  is associated with an image having a clear background, a textual description and structured category labels. In this work, we characterize each fashion item with a set of attributes (e.g., the color and category)  $A = \{a_q\}_{q=1}^Q$ , where  $a_q$  is the  $q$ -th attribute and  $Q$  is the total number of attributes. Each attribute  $a_q$  is associated with a set of elements representing its possible values  $E_q = \{e_q^1, e_q^2, \dots, e_q^{M_q}\}$ , where  $e_q^i$  refers to the  $i$ -th element and  $M_q$  is the total

number of elements corresponding to  $a_q$ . For simplicity, all  $E_q$  are combined in order and hence a unified set of attribute elements  $E = \cup_{q=1}^Q E_q = \{e_1, e_2, \dots, e_M\}$  is derived, where  $M = \sum_{q=1}^Q M_q$ .

In this work, we aim to tackle the essential task of modeling interpretable compatibility between fashion items for outfit recommendation by taking the user factor into account. Without loss of generality, we particularly investigate the following problem, “which bottom would be preferred by the user to match the given top”. Let  $p_{ij}^m$  denotes the preference of the user  $u_m$  towards the bottom  $b_j$  for the top  $t_i$ . Based on this metric, we can generate a personalized ranking list of bottoms  $b_j$ ’s for a given top  $t_i$  and hence solve the practical problem of personalized clothing matching. This ranking list is further used to infer compatible and incompatible pairs of attributes which can help us understand the underlying reason for discordance in an outfit.

In order to accurately measure  $p_{ij}^m$ , we devise a personalized compatibility modeling network  $F$ , which is capable of incorporating the user’s personal preferences into the compatibility modeling between fashion items as follows,

$$p_{ij}^m = F(t_i, b_j, u_m | \Theta_F) \quad (1)$$

where  $\Theta_F$  refers to the model parameters which are to be learned.

### B. Semantic Attribute Learning

Attributes play a pivotal role in both characterizing fashion items and interpreting the matching results. However, most of the existing benchmark datasets pertaining to outfit compatibility or recommendation lack attribute ground truth for fashion items. The task of acquiring accurate fine-grained attribute representations for benchmark datasets poses a challenge for us. We tackle this by deploying a modified Alexnet [29] architecture for each attribute learning task. While we argue that well pre-trained attribute classification networks is the most straightforward way to obtain the interpretable representations for a fashion item, we should not neglect the category information as an essential attribute as well. Therefore, after passing the image  $I$  of a fashion item into separate classification networks for each attribute ( $h_q$  for  $q^{th}$  attribute), we obtain the final semantic representation as follows:

$$f_i = [f_i^1; f_i^2; \dots; f_i^{Q-1}; f_i^c] \quad (2)$$

where  $f_i^q = h_q(I | \theta_q)$ ,  $\theta_q$  being the network parameter of  $h_q$ , and  $f_i^c$  stands for the one-hot semantic attribute representation derived from the category labels for the  $i^{th}$  item.

### C. PAI-BPR

To achieve personalized clothing matching, it is essential to account for both the item-item compatibility and the user-item preference. Hence, we measure the user preference

Table I  
ATTRIBUTE CATEGORIES AND EXAMPLES OF CORRESPONDING ATTRIBUTE ELEMENTS IN DEEPFASHION [30].

| Attributes              | Attribute Elements               |
|-------------------------|----------------------------------|
| Texture of clothes      | checked, striped, marble print   |
| Style of clothes        | athletic, bold, biker, chic      |
| Fabric of clothes       | classic cotton, netted, woven    |
| Shape of clothes        | a-line, oversized, swing         |
| Part details of clothes | back cutout, bell-sleeve, pocket |

towards a bottom for a given top by modeling both the general compatibility and the personal preferences of the user. Formally, we have,

$$p_{ij}^m = \mu \cdot s_{ij} + (1 - \mu) \cdot c_{mj} \quad (3)$$

$$s_{ij} = G(t_i, b_j | \Theta_G) \quad (4)$$

$$c_{mj} = P(u_m, b_j | \Theta_P) \quad (5)$$

where  $G$  and  $P$  correspond to the general compatibility modeling and personal preference modeling networks, respectively with  $\Theta_G$  and  $\Theta_P$  as the corresponding model parameters. Likewise,  $s_{ij}$  denotes the general compatibility between the top  $t_i$  and bottom  $b_j$ , and  $c_{mj}$  represents the personal preference of user  $u_m$  towards the bottom  $b_j$ .  $\mu$  is a non-negative tradeoff parameter to control the relative importance of both components.

The general item-item compatibility between fashion items entails complex interactions between attributes. These non-linear interactions are learned by a multi-layer perceptron (MLP) which performs excellently in various representation learning tasks [31], [32], [33]. We also leverage both the textual and visual modalities in the form of item description and image respectively, which coherently characterize a fashion item. On the other hand, for the personal preference modeling of a bottom, we resort to the matrix factorization framework, which has shown great success in personalized recommendation tasks [34], [35], [36], [37]. Inspired by Song et al. [26], the matrix factorisation method is extended to incorporate the latent content-based preference factors. The underlying philosophy is that the user-item interaction matrix can be decomposed into the latent user and item factors, whose inner product encodes the interaction scores between a user and an item. Further, since a user’s preference for an item is highly affected by the visual characteristics and the contextual features, like the brand and material, it is important to consider the content-based factors as well.

In order to model the implicit interaction among users and fashion items (i.e., tops and bottoms), the BPR framework is adopted since it has proven to be powerful in the pairwise preference modeling [38], [39], [40]. A training set,  $D$  is constructed in the following manner to train the BPR:

$$D := \{(m, i, j, k) | u_m \in U \cap (t_i, b_j) \in O_m \cap b_k \in B \setminus b_j\},$$

where the quadruplet  $(m, i, j, k)$  indicates that the user  $u_m$  prefers bottom  $b_j$  over  $b_k$  to match the given top  $t_i$ .

Table II  
PERFORMANCE OF ATTRIBUTE REPRESENTATION LEARNING IN TERMS OF AREA UNDER THE ROC CURVE.

| Attributes              | Tops | Bottoms |
|-------------------------|------|---------|
| Texture of clothes      | 0.77 | 0.75    |
| Style of clothes        | 0.81 | 0.77    |
| Fabric of clothes       | 0.86 | 0.82    |
| Shape of clothes        | 0.83 | 0.70    |
| Part Details of clothes | 0.84 | 0.75    |

#### IV. EXPERIMENT

To evaluate the proposed method, we conducted extensive experiments on the real-world dataset IQON3000 [26] by answering the following research questions:

- Does the proposed model outperform the state-of-the-art methods?
- How does the proposed model perform in the application of the personalized complementary fashion item retrieval?
- What are various approaches to understand the predictions and establish interpretability of proposed model?

##### A. Experimental Settings

1) *Dataset*: To evaluate our model, we adopted the real-world dataset IQON3000 [26] consisting of 217,806 unique outfits created by 3,568 users. Each fashion item is associated with a visual image, relevant categories and the description. Each of these images was resized to a smaller vector of size (100, 100, 3), in order to make the training process computationally efficient. All the results reported in Table III were evaluated using these resized images. Additionally, to train the attribute classification networks and obtain the semantic attribute representations of fashion items, we utilized an auxiliary benchmark dataset DeepFashion [30], comprising of 299,221 fashion items, each of which is labeled by 5 attributes with a total of 1,000 attribute elements.

2) *Attribute Representation Learning*: For learning the semantic attribute representation of fashion items, a modified AlexNet [29] architecture with an additional Dense Layer of size = 2040 before the final fully connected layer, was used. The output of this additional dense layer (size = (2040, 1)) was extracted and used as a representative vector of the same size. Apart from that, since the auxiliary dataset lacked the annotations for color, we quantized the image to a color vector of size (8, 1) using K-Means Clustering. This color vector was appended to the the vector obtained from the neural network to have a resultant attribute representation vector of size (2048, 1). The dataset was randomly split into training set (80%) and testing set (20%), and the

Table III  
PERFORMANCE COMPARISON AMONG DIFFERENT APPROACHES IN TERMS OF AREA UNDER THE ROC CURVE.

| Baselines/Proposed | Approach       | AUC           |
|--------------------|----------------|---------------|
| Baselines          | POP-T          | 0.6042        |
|                    | POP-U          | 0.5951        |
|                    | RAND           | 0.5014        |
|                    | Bi-LSTM [47]   | 0.6611        |
|                    | BPR-DAE [11]   | 0.6912        |
|                    | BPR-MF [38]    | 0.7867        |
|                    | VBPR [48]      | 0.8088        |
|                    | TBPR [26]      | 0.8102        |
|                    | VTBPR [26]     | 0.8194        |
|                    | GP-BPR [26]    | 0.8321        |
| Proposed           | PAI-BPR-V      | 0.8413        |
|                    | PAI-BPR-T      | 0.8432        |
|                    | <b>PAI-BPR</b> | <b>0.8502</b> |

‘binary cross-entropy’ loss was chosen due to multi-label hot encoded target type to train the Attribute Classification Network. The area under the ROC curve (AUC) was adopted as the metric to evaluate the performance of the attribute representation learning. Table II shows the performance of several attribute examples and the corresponding attribute elements.

3) *Contextual Representation*: In this work, we consider the category metadata and title description as the contextual information of the fashion item. First of all, we used the Japanese morphological analyzer Kuromoji<sup>5</sup> to tokenize the extracted text. Further, due to the the compelling success of the CNNs in various natural language processing tasks [41], [42], [43], we choose them for obtaining effective contextual representations, instead of the traditional linguistic features [44], [45]. Each word in the contextual description was represented in the form of 300-dimensional vector provided by the Japanese Word2vec called Nwjc2vec, in the search mode which is created from NINJAL Web Japanese Corpus [46]. Overall, each contextual description was first represented as a concatenated word vector, where each row represents one constituent word. Next, a single channel CNN consisting of a convolutional layer on top of the concatenated word vectors and a max pooling layer, was deployed. More precisely, we used four kernels with sizes of 2, 3, 4, and 5, respectively and for each kernel, we had 100 feature maps. After adding rectified linear unit (ReLU) as the activation function followed by another dense layer, we ultimately obtained a 512-D contextual representation vector for each fashion item.

##### B. Model Comparison

To evaluate the proposed model, the following state-of-the-art methods were chosen as baselines.

<sup>5</sup><http://www.atilika.org/http://www.atilika.org/>

| Query Top                           |  | Ranked Bottoms for given Query Top |   |   |   |   |   |
|-------------------------------------|--|------------------------------------|---|---|---|---|---|
| Some Historically Preferred Outfits |  | 1                                  | 2 | 3 | 4 | 5 | 6 |
| User 1                              |  |                                    |   |   |   |   |   |
| User 2                              |  |                                    |   |   |   |   |   |

Figure 3. Illustration of the ranking results of a given set of bottoms for a given top for two different users with varying historical preferences.

- **POP-T:** The “popularity” of the bottom was used as a measure of its compatibility with a top. It is defined as the number of outfits that the bottom appeared in the training set.
- **POP-U:** For this baseline, the popularity of the bottom was defined as the number of users who used this bottom as a part of an outfit in the training set.
- **RAND:** The compatibility scores between positive and negative pairs were assigned randomly.
- **Bi-LSTM:** The bidirectional LSTM method which sequentially models the outfit compatibility by predicting the next item conditioned on the previous ones [47], was modified to deal with a pair of top and bottom.
- **BPR-DAE:** In this baseline the content-based neural scheme introduced by Song et al. [11], that uses a dual autoencoder network to jointly model the coherent relation between different modalities of fashion items and the implicit preference among items was used.
- **BPR-MF:** This baseline captures the latent user-item relations by the MF method in the pairwise ranking method introduced by Rendle et al. [38]
- **VBPR:** The VBPR proposed by McAuley et al. [48] was adopted to exploit the visual data of fashion items with the factorization method to recommend an item for the user.
- **TBPR:** VBPR was modified by replacing the visual signals with the textual description of fashion items [26].
- **VTBPR:** This is an extended version of VBPR which includes context factor to comprehensively characterize the user’s preference from both the visual and contextual perspectives [26].
- **GP-BPR:** The current state-of-the-art in personalized compatibility modeling for outfit recommendation [26].

The performance comparison among different approaches can be seen in Table III. We can further infer the following from the table:

- 1) The content-based scheme performs better than a sequential model in modeling the general compatibility since BPR-DAE shows superiority over Bi-LSTM.
- 2) Both the modalities, i.e., contextual and visual, are important in modeling the personal preference because VTBPR outperforms VBPR, TBPR and BPR-MF.
- 3) GP-BPR achieves better performance than all the other baselines that focus on either the general compatibility modeling or personal preference modeling. This confirms that it is essential to incorporate both the general item-item compatibility and user-item preference in the context of personalized clothing matching.
- 4) Lastly, our model outperforms all other approaches, validating the advantage of a fine-grained semantic attribute representation for a fashion item.

We considered multimodal data since it provides a more effective representation of the information [49]. To validate this, we evaluate the significance of each modality in our model by comparing PAI-BPR with its two derivatives: PAI-BPR-V and PAI-BPR-T, where the visual and contextual modality of fashion items were individually explored, respectively. The performance comparison of these three models can also be seen in Table III. It can be observed that PAI-BPR outperforms both its derivatives, indicating that the visual and contextual signals complement each other and both contribute to the personalized compatibility modeling. Moreover, a comparison between the textual and visual variants of both VBPR and PAI-BPR shows that the textual variants perform better. This may be due to multiple reasons. Firstly, the key features of fashion items like the pattern and style can be effectively summarized in








Figure 5. Pairwise Attribute Correlation among attributes of selected tops and bottoms for a given user.

the number of samples in which the attribute was present and the importance value of that attribute. Darker nodes signify higher importance values, i.e. more important features for a particular user. This helps us extract and rank the most preferred attributes in a top and a bottom for a given user.

For analysing the underlying reason behind the ranking of bottoms for a given top as shown in Figure 3, we find the correlation between every attribute corresponding to the top and every attribute present in the set of all bottoms. Figure 5 shows the correlation between the 6 most prominent attributes in the top and the 3 most prominent attributes in each bottom. Since some bottoms had overlapping attributes, we had a total of 13 distinct attributes present in the set of bottoms. The bottom attributes are present on the X-axis while Y-axis signifies top’s attributes. The correlation between attribute pairs is in line with the ranking results. To find out the overall correlation of a bottom attribute, we calculate the sum of correlation between the bottom attribute and all the attributes in the top. For example, the attributes *Faux* and *Skater*, which are present in the top-ranked bottom for User 1, have the maximum overall correlation with top attributes. Also, even though the third-ranked bottom has a highly correlated attribute *Striped*, its other attributes like *Maxi* and *A-line* contribute to its low rank. Lastly, the two lowest ranked bottoms have the least correlated attributes like *Polka dot*, *Maxi*, *A-line*, *Print* and *Zippered*. To the best of our knowledge, this is the first time such analyses have been done to draw these inferences. These can be helpful in creating much better recommendation systems for fashion-centric portals, which would maximize recommendation hits and subsequently the revenue.

## V. CONCLUSION AND FUTURE WORK

In this work, we present an outfit recommendation scheme towards interpretable personalized clothing matching, called PAI-BPR, which models the compatibility between fashion items based on both general aesthetics and personal preferences of a user. We deploy a fine-grained attribute representation model to extract the visual features of an item to overcome the limitation of deep learning techniques in providing explainable results. Apart from visual features, we also incorporate the textual description and categorical information associated with an item since both the modalities are found to deliver valuable information, and together increase the overall performance of the model. In future, we plan to explore multiple ways of multimodal fusion in order to maximise the information gain. Nevertheless, extensive experiments have been conducted on a real-world dataset and our proposed model outperforms the state-of-the-art, demonstrating the effectiveness of PAI-BPR. Finally, with the help of our attribute representation model, we are successfully able to identify the prominent attribute-pairs in an outfit which make it compatible or incompatible. However, PAI-BPR currently works only for a pair of top and bottom and we would like to extend it to incorporate multiple items like footwear and accessories to comprise an entire outfit in the future.

## ACKNOWLEDGMENT

Rajiv Ratn Shah is partly supported by the Infosys Center for AI and the Center of Design and New Media at IIIT Delhi.



## REFERENCES

- [1] Y. Li, L. Cao, J. Zhu, and J. Luo, "Mining Fashion Outfit Composition Using an End-to-End Deep Learning Approach on Set Data," *IEEE Transactions on Multimedia*, vol. 19, no. 8, Aug. 2017.
- [2] E. Simo-Serra, S. Fidler, F. Moreno-Noguer, and R. Urtasun, "Neuroaesthetics in fashion: Modeling the perception of fashionability," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2015.
- [3] D. J. Hu, R. Hall, and J. Attenberg, "Style in the long tail: discovering unique interests with latent variable models in large scale social E-commerce," in *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '14*, 2014.
- [4] S. Liu, T. V. Nguyen, J. Feng, M. Wang, and S. Yan, "Hi, magic closet, tell me what to wear!" in *Proceedings of the 20th ACM international conference on Multimedia - MM '12*, 2012.
- [5] Si Liu, Zheng Song, Guangcan Liu, Changsheng Xu, Hanqing Lu, and Shuicheng Yan, "Street-to-shop: Cross-scenario clothing retrieval via parts alignment and auxiliary set," in *2012 IEEE Conference on Computer Vision and Pattern Recognition*, Jun. 2012.
- [6] X. Han, Z. Wu, Y.-G. Jiang, and L. S. Davis, "Learning fashion compatibility with bidirectional lstms," in *Proceedings of the 25th ACM International Conference on Multimedia*, 2017.
- [7] X. Song, F. Feng, X. Han, X. Yang, W. Liu, and L. Nie, "Neural compatibility modeling with attentive knowledge distillation," in *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. Association for Computing Machinery, 2018.
- [8] J. McAuley, C. Targett, Q. Shi, and A. van den Hengel, "Image-based recommendations on styles and substitutes," in *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2015.
- [9] Y. Hu, X. Yi, and L. S. Davis, "Collaborative Fashion Recommendation: A Functional Tensor Factorization Approach," in *Proceedings of the 23rd ACM international conference on Multimedia - MM '15*, 2015.
- [10] X. Chao, M. J. Huiskes, T. Gritti, and C. Ciuhu, "A framework for robust feature selection for real-time fashion style recommendation," in *Proceedings of the 1st international workshop on Interactive multimedia for consumer electronics - IMCE '09*, 2009.
- [11] X. Song, F. Feng, J. Liu, Z. Li, L. Nie, and J. Ma, "NeuroStylist: Neural Compatibility Modeling for Clothing Matching," in *Proceedings of the 2017 ACM on Multimedia Conference - MM '17*, 2017.
- [12] Y. Li, L. Cao, J. Zhu, and J. Luo, "Mining fashion outfit composition using an end-to-end deep learning approach on set data," *IEEE Transactions on Multimedia*, vol. PP, 08 2016.
- [13] F. Fan, J. Xiong, and G. Wang, "On interpretability of artificial neural networks," 2020.
- [14] Z. Feng, Z. Yu, Y. Yang, Y. Jing, J. Jiang, and M. Song, "Interpretable Partitioned Embedding for Customized Multi-item Fashion Outfit Composition," in *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*, Jun. 2018.
- [15] X. Wang, B. Wu, and Y. Zhong, "Outfit compatibility prediction and diagnosis with multi-layered comparison network," in *Proceedings of the 27th ACM International Conference on Multimedia*, 2019.
- [16] X. Han, X. Song, J. Yin, Y. Wang, and L. Nie, "Prototype-guided Attribute-wise Interpretable Scheme for Clothing Matching," in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval - SIGIR'19*, 2019.
- [17] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme, "BPR: Bayesian Personalized Ranking from Implicit Feedback," 2009.
- [18] Z. Al-Halah, R. Stiefelwagen, and K. Grauman, "Fashion Forward: Forecasting Visual Style in Fashion," in *2017 IEEE International Conference on Computer Vision (ICCV)*, Oct. 2017.
- [19] 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 Transactions on Multimedia*, vol. 18, no. 6, Jun. 2016.
- [20] J. Huang, R. Feris, Q. Chen, and S. Yan, "Cross-Domain Image Retrieval with a Dual Attribute-Aware Ranking Network," in *2015 IEEE International Conference on Computer Vision (ICCV)*, Dec. 2015.
- [21] Y. Jiang, Q. Xu, X. Cao, and Q. Huang, "Who to Ask: An Intelligent Fashion Consultant," in *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval - ICMR '18*, 2018.
- [22] Z. Zhou, X. Di, W. Zhou, and L. Zhang, "Fashion Sensitive Clothing Recommendation Using Hierarchical Collocation Model," in *2018 ACM Multimedia Conference on Multimedia Conference - MM '18*, 2018.
- [23] P. Wang, J. Guo, and Y. Lan, "Modeling Retail Transaction Data for Personalized Shopping Recommendation," in *Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management - CIKM '14*, 2014.
- [24] S. C. Hidayati, C.-C. Hsu, Y.-T. Chang, K.-L. Hua, J. Fu, and W.-H. Cheng, "What Dress Fits Me Best?: Fashion Recommendation on the Clothing Style for Personal Body Shape," in *2018 ACM Multimedia Conference on Multimedia Conference - MM '18*, 2018.
- [25] Y. Ma, X. Yang, L. Liao, Y. Cao, and T.-S. Chua, "Who, Where, and What to Wear?: Extracting Fashion Knowledge from Social Media," in *Proceedings of the 27th ACM International Conference on Multimedia - MM '19*, 2019.
- [26] X. Song, X. Han, Y. Li, J. Chen, X.-S. Xu, and L. Nie, "GP-BPR: Personalized Compatibility Modeling for Clothing Matching," in *Proceedings of the 27th ACM International Conference on Multimedia - MM '19*, 2019.

- [27] D. Verma, K. Gulati, and R. R. Shah, "Addressing the Cold-Start Problem in Outfit Recommendation Using Visual Preference Modelling," *arXiv:2008.01437 [cs]*, Aug. 2020, arXiv: 2008.01437. [Online]. Available: <http://arxiv.org/abs/2008.01437>
- [28] A. Piazza, P. Kröckel, and F. Bodendorf, "Emotions and fashion recommendations: evaluating the predictive power of affective information for the prediction of fashion product preferences in cold-start scenarios," in *Proceedings of the International Conference on Web Intelligence - WI '17*, 2017.
- [29] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 1*, 2012.
- [30] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang, "DeepFashion: Powering Robust Clothes Recognition and Retrieval with Rich Annotations," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jun. 2016.
- [31] Z. Wang, X. Bai, M. Ye, and S. Satoh, "Incremental Deep Hidden Attribute Learning," in *2018 ACM Multimedia Conference on Multimedia Conference - MM '18*, 2018.
- [32] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs," *arXiv:1606.00915 [cs]*, May 2017, arXiv: 1606.00915.
- [33] M. Liu, X. Wang, L. Nie, Q. Tian, B. Chen, and T.-S. Chua, "Cross-modal Moment Localization in Videos," in *2018 ACM Multimedia Conference on Multimedia Conference - MM '18*, 2018.
- [34] J. Bobadilla, R. Bojorque, A. Hernando Esteban, and R. Hurtado, "Recommender Systems Clustering Using Bayesian Non Negative Matrix Factorization," *IEEE Access*, vol. 6, 2018, conference Name: IEEE Access.
- [35] D. Kim, C. Park, J. Oh, S. Lee, and H. Yu, "Convolutional Matrix Factorization for Document Context-Aware Recommendation," in *Proceedings of the 10th ACM Conference on Recommender Systems*, Sep. 2016.
- [36] C. Packer, J. McAuley, and A. Ramisa, "Visually-Aware Personalized Recommendation using Interpretable Image Representations," *arXiv:1806.09820 [cs]*, Aug. 2018, arXiv: 1806.09820.
- [37] Y. Koren, "Collaborative filtering with temporal dynamics," in *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining - KDD '09*, 2009.
- [38] D. Cao, L. Nie, X. He, X. Wei, S. Zhu, and T.-S. Chua, "Embedding Factorization Models for Jointly Recommending Items and User Generated Lists," in *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, Aug. 2017.
- [39] S. Liu, Q. Wu, and C. Miao, "Personalized recommendation considering secondary implicit feedback," in *Proceedings of the 3rd IEEE International Conference on Agents (ICA'18)*, July 2018.
- [40] B. Loni, R. Pagano, M. Larson, and A. Hanjalic, "Bayesian Personalized Ranking with Multi-Channel User Feedback," in *Proceedings of the 10th ACM Conference on Recommender Systems*, Sep. 2016.
- [41] Z. Hu, X. Ma, Z. Liu, E. Hovy, and E. Xing, "Harnessing deep neural networks with logic rules," in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Aug. 2016.
- [42] Y. Kim, "Convolutional neural networks for sentence classification," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Oct. 2014.
- [43] A. Severyn and A. Moschitti, "Twitter sentiment analysis with deep convolutional neural networks," in *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2015.
- [44] X. Song, L. Nie, L. Zhang, M. Akbari, and T.-S. Chua, "Multiple social network learning and its application in volunteerism tendency prediction," in *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2015.
- [45] X. Song, L. Nie, L. Zhang, M. Liu, and T.-S. Chua, "Interest inference via structure-constrained multi-source multi-task learning," in *Proceedings of the 24th International Conference on Artificial Intelligence*, 2015.
- [46] H. Shinnou, M. Asahara, K. Komiya, and M. Sasaki, "nwjc2vec: Word Embedding Data Constructed from NIN-JAL Web Japanese Corpus," *Journal of Natural Language Processing*, vol. 24, no. 5, 2017.
- [47] X. Han, Z. Wu, Y.-G. Jiang, and L. S. Davis, "Learning Fashion Compatibility with Bidirectional LSTMs," in *Proceedings of the 2017 ACM on Multimedia Conference - MM '17*, 2017.
- [48] R. He and J. McAuley, "VBPR: Visual Bayesian Personalized Ranking from Implicit Feedback," *arXiv:1510.01784 [cs]*, Oct. 2015, arXiv: 1510.01784.
- [49] R. Shah and R. Zimmermann, *Multimodal analysis of user-generated multimedia content*. Springer, 2017.
- [50] L. Jiang, S.-I. Yu, D. Meng, Y. Yang, T. Mitamura, and A. G. Hauptmann, "Fast and Accurate Content-based Semantic Search in 100M Internet Videos," in *Proceedings of the 23rd ACM international conference on Multimedia - MM '15*, 2015.
- [51] J. Yao, Y. Wang, Y. Zhang, J. Sun, and J. Zhou, "Joint latent dirichlet allocation for social tags," *IEEE Transactions on Multimedia*, vol. 20, no. 1, Jan 2018.
- [52] H. Yin, H. Chen, X. Sun, H. Wang, Y. Wang, and Q. V. Hung Nguyen, "Sptf: A scalable probabilistic tensor factorization model for semantic-aware behavior prediction," in *2017 IEEE International Conference on Data Mining (ICDM)*, Nov 2017.