

## Addressing the Cold-Start Problem in Outfit Recommendation Using Visual Preference Modelling

Dhruv Verma\*, Kshitij Gulati\* and Rajiv Ratn Shah  
 Indraprastha Institute of Information Technology, Delhi, India  
 {dhruv17046, kshitij17162 and rajivrtn}@iiitd.ac.in

**Abstract**—With the global transformation of the fashion industry and a rise in the demand for fashion items worldwide, the need for an effectual fashion recommendation has never been more. Despite various cutting-edge solutions proposed in the past for personalising fashion recommendation, the technology is still limited by its poor performance on new entities, i.e. the cold-start problem. In this paper, we attempt to address the cold-start problem for new users, by leveraging a novel *visual preference modelling* approach on a small set of input images. We demonstrate the use of our approach with *feature-weighted clustering* to personalise occasion-oriented outfit recommendation. Quantitatively, our results show that the proposed visual preference modelling approach outperforms state of the art in terms of clothing attribute prediction. Qualitatively, through a pilot study, we demonstrate the efficacy of our system to provide diverse and personalised recommendations in cold-start scenarios.

**Keywords**—personalised outfit recommendation, cold-start problem, visual preference modelling, feature-weighted clustering

### I. INTRODUCTION

With the growing influence of online social media, we have a strong desire to remain connected while sharing almost every aspect of our lives, including delectable food, places we visit, to the outfits we wear. Since an individual’s dressing style speaks volumes about his/her personality, it has given increased importance to the appropriateness of an outfit in a particular context. We always try to portray a better version of ourselves and in this case fashion knowledge not only helps us look better, but it speaks to the physiological needs and caters to the demands of social events and activities. With this increased consciousness, e-commerce for fashion is booming, but consumers still face issues while selecting fashion outfits that are appropriate in context and suit their liking.

In the past, various studies have attempted to tackle different fashion-related issues like fashion retrieval [1], [2], clothing recognition and classification [3], [4], outfit recommendation [5], [6], etc. When we talk about outfit recommendation, it cannot be generic, since user preference is inherently subjective. User preference is a reflection of age, occupation, culture, place of living, etc. Therefore, personalisation is essential because it ensures outfit recommendation is in line with the users’ personal fashion taste and incorporates users’ likes and dislikes from a variety of perspectives. Personalised fashion recommendation has gained a lot of attention in recent work [6]–[9]. However, most of them fail to address the cold-start problem [10] for new users (see Figure 1) and solving this issue forms the basis for our approach.

Fashion knowledge can best be represented in terms of person, occasion and clothing [11]. Presently, various fashion datasets comprising DeepFashion [12], Polyvore [13], FashionVC [14] have been collected. However, none of these has adequate representation for user-context as well as occasion. Hence, we curate a

\*These authors contributed equally to this work.

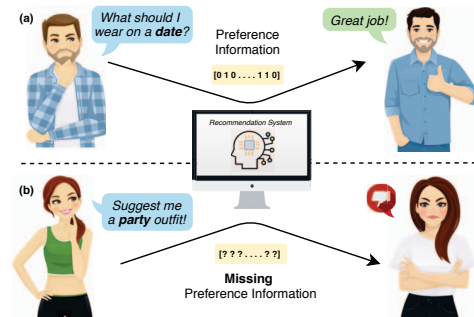


Figure 1. An illustration of the cold-start problem in outfit recommendation for new users. (a) represents a profiled user, whose preferences are known to the system, while (b) represents a new user.

high-quality dataset, constituting images scraped from Instagram and Pinterest, representing a diverse variety of occasion-oriented fashion knowledge. The optimal recommendation for a particular occasion scenario would essentially be suitable for that occasion while keeping the fashion elements or styles as per the user’s liking. For example, a man attending a conference would be recommended to wear a formal attire, say a suit. However, elements of the outfit like colour, texture, style, etc. can be modulated as per his preference. In this paper, we attempt to address the cold-start problem for new users by modelling their preference visually. By taking a few images of the users depicting their fashion taste as an input, we aim to extract rich fashion concepts and model these concepts as preferences to the recommender for cold-start scenarios.

We experiment with various pre-existing techniques to extract fashion concepts from the users’ images. Our model predicts clothing categories with an accuracy of 73.5%, F1 score of 0.72 and attributes with an accuracy of 89.1%, F1 score of 0.88. Our model outperforms state of the art in attribute classification and has comparable performance for category classification. We further evaluate our system by conducting a pilot study on ten subjects and find that most users find recommendations with visual preference modelling more relevant as opposed to recommendations without it. The main contributions of this work are as follows:

- We contribute a fashion knowledge dataset and demonstrate the efficacy of multi-task learning in terms of capturing fashion concepts from a wide variety of images.
- We propose a novel visual preference modelling approach with feature-weighted clustering for personalising outfit recommendation in cold-start scenarios<sup>1</sup>.

<sup>1</sup>Source code and dataset: <https://github.com/Dhruv-Verma/fashionist>.

## II. RELATED WORK

The main objective of this paper is to address the cold-start problem in outfit recommendation. Hence, we primarily discuss recent developments in fashion concept prediction and personalised outfit recommendation in this section.

### A. Fashion Concept Prediction

Fashion outfits contain various high-level concepts like style, season, occasion and environment, which can be inferred from design elements like shape, colour, cloth, texture, etc. [15]. Studies conducted in the past [16] suggest that low-level elements of fashion like clothing categories and fine-grained attributes can be translated to these high-level fashion concepts. To this end, various computer-vision based approaches have been explored for representing clothing elements from images. With the availability of rich fashion databases like DeepFashion [12], ModaNet [17], etc., deep neural network-based approaches have risen to new heights in terms of performance. Liu et al. [12] leveraged a branched CNN architecture, *FashionNet*, which learns clothing features by jointly predicting clothing attributes and landmarks. Ma et al. [11] proposed a contextualized fashion concept learning module using Bi-LSTMs integrated with a CNN backbone. Yan et al. [18] proposed *StyleNet* which used multi-task representation learning for creating clothing representations accommodating different fashion concepts like seasons, garments and styles. While some of these approaches were successful in capturing the dependencies and correlation between different fashion elements and concepts, their results could not show significant improvement in predicting fine-grained clothing attributes. Since multi-task learning has shown promising results in many applications [19], we are motivated to leverage the information-sharing capabilities of multi-task learning in the domain of fashion concept prediction while modelling the dependencies between clothing categories and attributes. With this approach, we observe a significant improvement in performance in comparison with existing work while keeping the computational complexity low. This is attributed to our elemental approach towards clothing representation along with dependence modelling.

### B. Cold-Start Problem in Outfit Recommendation

In recent work, various successful attempts towards personalising outfit recommendation have been made, viz; tensor factorization [7], visual compatibility modelling [9], [20], and multimodal attention-based networks [5], [6]. However, all of these works are limited in scope by their inability to address the cold-start problem for outfit recommendation. Only a small body of work has focussed on tackling this problem. Bracher et al. [21] proposed novel latent space representation for fashion items called *FashionDNA*. [22], [23] used visual representations incorporating categories and styles for clothing items. These methods attempted to address cold-start in terms of new items for recommendation but overlooked the new user aspect. Piazza et al. [22] evaluated the predictive power of users' affective information for mitigating the cold-start problem for new users. Their findings suggest that mood information positively influences the quality of predictions. However, this improvement was not significant. We, therefore, find an opportunity to address this problem via a practical approach. We exploit fashion concept prediction for inferring preferences from a small set of input images for a new user.

## III. DATASET

Social media platforms offer rich user-centric data with images of users across the world, depicting diverse real-world fashion knowledge with natural backgrounds. These images usually contain indications to various occasions such as conferences, weddings and parties. We contribute a dataset which comprises of 2893 high-quality images scraped from Instagram and Pinterest. We collect images for seven types of occasions, namely *Travel*, *Sports*, *Dating*, *Wedding*, *Party*, *Office* and *Prom*, by using the presently trending hashtags and keywords associated with these occasions. For example, we used *#sportsoutfit* for Sports and *#officewear* for Office. We ensure that images extracted are balanced in terms of gender and are easily distinguishable and diverse in terms of various fashion concepts.

After extracting these images, we perform manual filtering to rule out all possible noisy images which do not adequately represent fashion knowledge and have poor quality, low brightness, side angle shots and selfies. We use a pre-trained CNN-based model [24] for age and gender classification of each image. We define a set of categories and attributes to best represent the fashion knowledge in terms of the clothing aspects for each image, as shown in Figure 2. Each category defines a particular region of clothing (e.g. t-shirt: upper body, athletic pants: lower body) and each attribute further characterises a particular category (e.g. sleeve length, neckline: t-shirt, lower body length: athletic pants).

Two fashion experts annotate each image following our annotation scheme. For evaluating the reliability of annotations, an inter-annotator agreement is calculated using kappa statistics. In case of disagreement for a particular annotation, a third fashion expert breaks the tie. As a general guideline, a kappa value above 0.6 is considered adequate and above 0.8 is considered almost perfect [25]. We find the inter-annotator agreement for both tasks to be reasonably high with a kappa value of *0.81* for category annotation and a kappa value of *0.68* for attribute annotation. As a result, we achieve our fashion knowledge database which describes each image in terms of the *person* (Gender and Age), *clothing* (Categories and Attributes) and *occasion*.

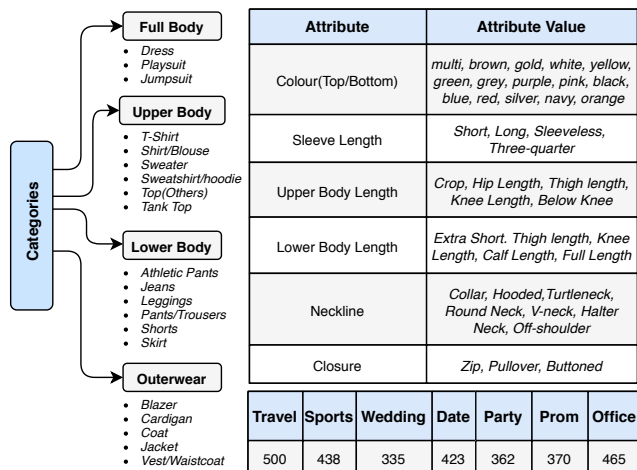


Figure 2. Our dataset, with 20 categories, 39 attributes across 7 occasions

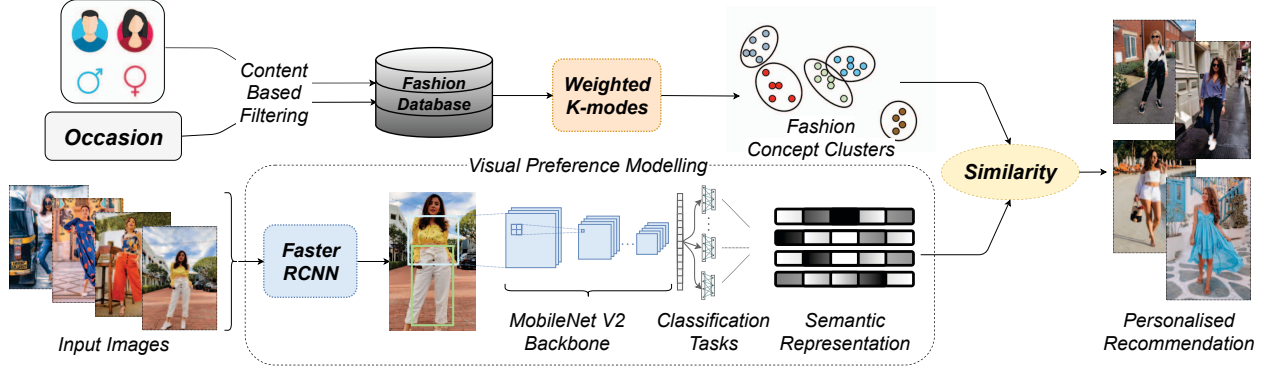


Figure 3. The proposed pipeline for personalised outfit recommendation in cold-start scenarios

#### IV. METHODOLOGY

Our objective is to demonstrate the use of fashion concept prediction to model the preferences of new users given a small set of images as input. We expect the users to provide images that depict a diverse range of fashion choices highlighting their likes and dislikes towards various fashion elements. The first portion of our work is focused on the preference modelling of visual fashion elements like clothing categories and attributes. In the second portion, we discuss our strategy for incorporating the visual preferences into outfit recommendation. (see Figure 3)

##### A. Visual Preference Modelling

Fashion outfits are a combination of different clothing items (e.g. tops and bottoms or full body dresses). To extract fashion concepts with efficacy, we first need to represent different fashion items in terms of regions of interest. We use Faster RCNN (with a ResNet-50 backbone) [26], which is a state of the art technique for real-time object detection. We fine-tuned the Faster RCNN (pre-trained on the MS-COCO dataset [27]) on the ModaNet dataset [17] to ensure that it only detects relevant fashion elements as regions of interest.

These regions of interest are unlabelled and cannot directly be used to represent visual fashion knowledge. Therefore, we need to transform them into a semantic space representation which can encode the visual fashion knowledge in terms of different categories and attributes. Since our objective is to recommend outfits from our curated dataset, the semantic space representation of the users' images should also incline with the labelled fashion knowledge in our dataset. Hence, for a given image  $I$ , we have a set of *clothing regions*  $\{r_1, r_2, \dots, r_i, \dots, r_n\}$ , the goal is to predict *category label*  $c_{r_i} \in \{c_1, c_2, \dots, c_m\}$  and a set of *attribute values*  $\{a_{r_i}^1, a_{r_i}^2, \dots, a_{r_i}^p\}$  where  $a_{r_i}^j \in$  *attribute types*  $\{a_1^j, a_2^j, \dots, a_x^j\}$  for each  $r_i$ . Hence, each region of interest is characterized by a clothing category and attribute values semantically similar to our labeled dataset. Thereafter, the representations of individual regions of interest are pooled together to form vector  $v_I$  for every image  $I$  in the given collection of users' images.

Since clothing categories and attribute values are not independent entities, there exist dependencies between them. In order to learn the hidden representations capable of accommodating individual features as well as dependencies, we make use of multi-task learning. We use a multi-task classifier with a MobileNet V2

[28] backbone. Different classification heads are responsible for optimizing the category and attribute prediction tasks. The minimization objective function can be described as in Equation 1.

$$L_{cumulative} = \lambda_1 L_{category} + \lambda_2 \sum L_{attribute} \quad (1)$$

$\lambda_1$  and  $\lambda_2$  are the weights for category and attribute losses respectively. The cumulative loss function ensures that the shared hidden representation incorporates both individual as well as dependency-based features. The MobileNet backbone is initially trained on the large DeepFashion [12] dataset to learn various high-level feature maps relevant for fashion domain tasks. Finally, this backbone with multiple classification heads is fine-tuned on our dataset.

##### B. Personalised Outfit Recommendation

After extracting the vector representation (as described in Section IV-A) for the given input images, our goal is to fuse this extracted preference knowledge into occasion-oriented outfit recommendation. Within the boundaries of a particular occasion, relevant outfits may represent various high-level fashion concepts like style, patterns, seasons, etc. For example, male outfits belonging to the occasion *Office* may form subgroups based on styles present, viz; a style subgroup may have a more formal appeal (e.g. suits) while another may depict a more casual style like a sweater with chinos. Similarly, the presence of outerwear and attribute values like sleeve length and neckline could lead to the formation of subgroups depicting different seasons.

After applying content-based filters for gender and occasion on our dataset, we retrieve the relevant outfit images. We need to extract the high-level concept subgroups from this set. Since the outfits in this set are labelled with appropriate category and attribute information, we use an unsupervised learning approach, a modified version of the *k-modes* clustering algorithm [29] for forming the concept clusters. The modifications include a weighted dissimilarity measure instead of the simple hamming distance (as described in [30]) and different weights  $\alpha$  and  $\beta$  for the category and attribute related features respectively (refer Equation 2). For identifying the optimal number of clusters to be formed for each subgroup, we use the metric *silhouette index* [31].

$$D_m(X, Y) = \alpha \sum_{j \in cat} \delta(x_j, y_j) + \beta \sum_{j \in attr} \delta(x_j, y_j) \quad (2)$$

After forming the concept clusters, a trivial approach is to draw an equal number of samples from each cluster for the final recommendation. While this may ensure greater diversity in the recommendation, it does not account for incorporating the users’ preferences in any way. Therefore, we weight each cluster by calculating the similarity between the semantic representations of outfits in the input images and the concept clusters (semantic representation of centroids used). The number of samples drawn for the recommendation is decided based on the weights assigned to each cluster. For example, if some images given by a female user represent her liking for sleeveless dresses (category: dress, sleeve length: sleeveless), the recommended results are likely to have a greater number of outfits reflecting these features. By doing so, we ensure the diversification of results in addition to the incorporation of users’ preferences in a cold-start scenario.

Table I  
PERFORMANCE COMPARISON OF DIFFERENT CNN ARCHITECTURES

Backbone (Configuration)	Task	Categories		Attributes	
		Acc.	F1	Acc.	F1
ResNet-18 (fine-tuned 1 block)	category	0.68	0.65	-	-
	attribute	-	-	0.73	0.71
	multi-task	0.69	0.68	0.87	0.86
ResNet-50 (fine-tuned 1 block)	category	0.65	0.64	-	-
	attribute	-	-	0.71	0.70
	multi-task	0.68	0.665	0.86	0.85
MobileNet-V2 (fine-tuned 3 blocks)	category	0.71	0.69	-	-
	attribute	-	-	0.76	0.74
	multi-task	<b>0.735</b>	<b>0.72</b>	<b>0.89</b>	<b>0.88</b>

## V. CLASSIFICATION RESULTS

For performing robust fashion concept prediction using multi-task learning, we experimented with different networks as backbones and fine-tuned them with different configurations. In this section, we briefly describe the performance of these network architectures and compare them with baselines (single classification task CNNs) and state of the art. We experimented with various CNN-based backbones like MobileNet V2, ResNet-18 and ResNet-50 for classification of clothing categories and attributes in both single and multi-task configurations. We initially train these networks on DeepFashion [12] and then on our dataset. We use the *cross-entropy* loss for each classification head, *adam* optimizer with a learning rate of 0.001, iterated for 100 epochs. We find that **multi-task classifier with MobileNet backbone** (with 3 blocks of fine-tuning,  $\lambda_1/\lambda_2 = 2$ ,  $\alpha/\beta = 2$ , calculated empirically) outperforms its alternatives, with an accuracy of **73.5%** and **89.1%** for category and average attribute classification respectively (refer Table I).

We also compare the performance of our model with state of the art architectures like FashionKE [11] and FashionNet [12]. Our model compares well these systems in terms of category classification and outperforms them for attribute classification (refer Table II). Our results demonstrate the value added by the information-sharing capabilities of multi-task learning.

## VI. PILOT STUDY AND EVALUATION

In order to evaluate our proposed method for personalised outfit recommendation, we conducted a brief pilot study with ten subjects

Table II  
PERFORMANCE COMPARISON WITH STATE OF THE ART ARCHITECTURES

Model	Accuracy	
	Category	Attributes
FashionNet [16]	67.33%	65.6%
FashionKE [11]	<b>73.95%</b>	69.59%
<b>MobileNet MTL (Our Method)</b>	73.5%	<b>89.1%</b>

(five male and five female, belonging to an age group of 18 - 35). We asked users to rate each recommendation on a scale of 10 in terms of *relevance* and *diversity*. We ask users to give a total of 10 images displaying their fashion preferences along with the occasion for which they wish to dress up. The system then recommends the user a set of 10 outfits for that particular occasion both with and without integrating their personal preferences. The users are not aware of whether the recommendations are inclusive of their preferences or not, to eliminate any bias while evaluating. The results of the pilot study (refer Table III) show that users rate the recommendations highly in terms of diversity, affirming the success of the weighted clustering approach. We find that on an average, users find recommendations with visual preference modelling more relevant as opposed to recommendations without it. An important observation is a slight reduction in the diversity score with visual preference modelling. This can be attributed to the increased importance given to fashion concept clusters which are more similar to the users’ outfit images, leading to more number of outfits from the same cluster. Overall, the pilot study results confirm that we have been able to optimise outfit recommendation as well as address the cold-start problem with this approach.

Table III  
RESULTS OF THE PILOT STUDY (NORMALISED)

Subject	Without Preference		With Preference	
	Diversity	Relevance	Diversity	Relevance
Male	0.76	0.58	0.66	0.73
Female	0.85	0.62	0.74	0.81
<b>Average</b>	<b>0.81</b>	<b>0.60</b>	<b>0.70</b>	<b>0.77</b>

## VII. CONCLUSION AND FUTURE WORK

In this paper, we attempt to address the cold-start problem for new users in occasion-oriented outfit recommendation. We demonstrate the use of the proposed *Visual Preference Modelling* with the *Feature-Weighted Clustering* method for generating personalised recommendations using a small set of images as input. The results of our preliminary study affirm that our approach is beneficial for improving outfit recommendation in cold-start scenarios. To further assess the performance of our system, we plan to conduct a more extensive evaluation with a larger participant group. We also note that personalisation is significantly dependent on the quality of images provided by the user. The system can hereon be improved by instilling certain high-level fashion concepts like body shape, texture, weather, seasonal variation and cultural dependence of certain clothing styles which can help improve personalisation. Additionally, the use of multimodal information has shown significant improvement in performance as opposed to systems using a single modality [32]. We plan on investigating these avenues in the future.

## REFERENCES

- [1] L. Liao, X. He, B. Zhao, C.-W. Ngo, and T.-S. Chua, "Interpretable Multimodal Retrieval for Fashion Products," in *2018 ACM Multimedia Conference on Multimedia Conference - MM '18*, (Seoul, Republic of Korea), pp. 1571–1579, ACM Press, 2018.
- [2] M. R. Mustaffa, G. S. Wai, L. N. Abdullah, and N. A. Nasharudin, "Dress me up!: content-based clothing image retrieval," in *Proceedings of the 3rd International Conference on Cryptography, Security and Privacy - ICCSP '19*, (Kuala Lumpur, Malaysia), pp. 206–210, ACM Press, 2019.
- [3] Y. Kalantidis, L. Kennedy, and L.-J. Li, "Getting the look: clothing recognition and segmentation for automatic product suggestions in everyday photos," in *Proceedings of the 3rd ACM conference on International conference on multimedia retrieval - ICMR '13*, (Dallas, Texas, USA), p. 105, ACM Press, 2013.
- [4] S. C. Hidayati, W.-H. Cheng, and K.-L. Hua, "Clothing genre classification by exploiting the style elements," in *Proceedings of the 20th ACM international conference on Multimedia - MM '12*, (Nara, Japan), p. 1137, ACM Press, 2012.
- [5] Y. Lin, M. Moosaei, and H. Yang, "OutfitNet: Fashion Outfit Recommendation with Attention-Based Multiple Instance Learning," in *Proceedings of The Web Conference 2020*, (Taipei Taiwan), pp. 77–87, ACM, Apr. 2020.
- [6] X. Chen, H. Chen, H. Xu, Y. Zhang, Y. Cao, Z. Qin, and H. Zha, "Personalized Fashion Recommendation with Visual Explanations based on Multimodal Attention Network: Towards Visually Explainable Recommendation," in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, (Paris France), pp. 765–774, ACM, July 2019.
- [7] 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*, (Brisbane, Australia), pp. 129–138, ACM Press, 2015.
- [8] W. Chen, P. Huang, J. Xu, X. Guo, C. Guo, F. Sun, C. Li, A. Pfadler, H. Zhao, and B. Zhao, "POG: Personalized Outfit Generation for Fashion Recommendation at Alibaba iFashion," in *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, (Anchorage AK USA), pp. 2662–2670, ACM, July 2019.
- [9] R. Yin, K. Li, J. Lu, and G. Zhang, "Enhancing Fashion Recommendation with Visual Compatibility Relationship," in *The World Wide Web Conference on - WWW '19*, (San Francisco, CA, USA), pp. 3434–3440, ACM Press, 2019.
- [10] B. Lika, K. Kolomvatsos, and S. Hadjiefthymiades, "Facing the cold start problem in recommender systems," *Expert Systems with Applications*, vol. 41, pp. 2065–2073, Mar. 2014.
- [11] 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*, (Nice France), pp. 257–265, ACM, Oct. 2019.
- [12] 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)*, (Las Vegas, NV, USA), pp. 1096–1104, IEEE, June 2016.
- [13] 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*, (Mountain View, California, USA), pp. 1078–1086, ACM Press, 2017.
- [14] 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*, (Mountain View, California, USA), pp. 753–761, ACM Press, 2017.
- [15] R. Sorger and J. Udale, *The fundamentals of fashion design*. AVA academia, Lausanne: AVA Publishing, SA, 2nd edition ed., 2012.
- [16] K. Vaccaro, S. Shivakumar, Z. Ding, K. Karahalios, and R. Kumar, "The Elements of Fashion Style," in *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, (Tokyo Japan), pp. 777–785, ACM, Oct. 2016.
- [17] S. Zheng, F. Yang, M. H. Kiapour, and R. Piramuthu, "Modanet: A large-scale street fashion dataset with polygon annotations," in *Proceedings of the 26th ACM International Conference on Multimedia*, MM '18, (New York, NY, USA), p. 1670–1678, Association for Computing Machinery, 2018.
- [18] C. Yan, L. Zhou, and Y. Wan, "A Multi-Task Learning Model for Better Representation of Clothing Images," *IEEE Access*, vol. 7, pp. 34499–34507, 2019.
- [19] Y. Zhang and Q. Yang, "A Survey on Multi-Task Learning," *arXiv:1707.08114 [cs]*, July 2018. arXiv: 1707.08114.
- [20] D. Sagar, J. Garg, P. Kansal, S. Bhalla, R. R. Shah, and Y. Yu, "PAI-BPR: Personalized Outfit Recommendation Scheme with Attribute-wise Interpretability," *arXiv:2008.01780 [cs]*, Aug. 2020. arXiv: 2008.01780.
- [21] C. Bracher, S. Heinz, and R. Vollgraf, "Fashion DNA: Merging Content and Sales Data for Recommendation and Article Mapping," *arXiv:1609.02489 [cs]*, Sept. 2016. arXiv: 1609.02489.
- [22] 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*, (Leipzig Germany), pp. 1234–1240, ACM, Aug. 2017.
- [23] M. Sun, F. Li, J. Lee, K. Zhou, G. Lebanon, and H. Zha, "Learning multiple-question decision trees for cold-start recommendation," in *Proceedings of the sixth ACM international conference on Web search and data mining - WSDM '13*, (Rome, Italy), p. 445, ACM Press, 2013.
- [24] G. Levi and T. Hassner, "Age and gender classification using convolutional neural networks," in *2015 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, (Boston, MA, USA), pp. 34–42, IEEE, June 2015.

- [25] S. Nowak and R. Stefan, "üger. 2010. how reliable are annotations via crowdsourcing: A study about interannotator agreement for multi-label image annotation," in *Proceedings of the International Conference on Multimedia Information Retrieval*, pp. 557–566.
- [26] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, NIPS'15*, (Cambridge, MA, USA), p. 91–99, MIT Press, 2015.
- [27] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L. Zitnick, and P. Dollár, "Microsoft COCO: Common Objects in Context," *arXiv:1405.0312 [cs]*, Feb. 2015. arXiv: 1405.0312.
- [28] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications," *arXiv:1704.04861 [cs]*, Apr. 2017. arXiv: 1704.04861.
- [29] A. Chaturvedi, P. E. Green, and J. D. Carroll, "K-modes clustering," *Journal of classification*, vol. 18, no. 1, pp. 35–55, 2001.
- [30] S.Aranganayagi and K.Thangavel, "Improved K-Modes For Categorical Clustering Using Weighted Dissimilarity Measure," Mar. 2009. Publisher: Zenodo.
- [31] S. Aranganayagi and K. Thangavel, "Clustering Categorical Data Using Silhouette Coefficient as a Relocating Measure," in *International Conference on Computational Intelligence and Multimedia Applications (ICCIMA 2007)*, (Sivakasi, Tamil Nadu, India), pp. 13–17, IEEE, Dec. 2007.
- [32] R. Shah and R. Zimmermann, *Multimodal analysis of user-generated multimedia content*. Springer, 2017.