



# Fashion recommendation based on style and social events

Federico Becattini<sup>1</sup> · Lavinia De Divitiis<sup>1</sup> · Claudio Baecchi<sup>1</sup> · Alberto Del Bimbo<sup>1</sup>

Received: 1 August 2022 / Revised: 3 March 2023 / Accepted: 6 April 2023 /

Published online: 17 May 2023

© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2023

## Abstract

Fashion recommendation is often declined as the task of finding complementary items given a query garment or retrieving outfits that are suitable for a given user. In this work we address the problem by adding an additional semantic layer based on the style of the proposed dressing. We model style according to two important aspects: the mood and the emotion concealed behind color combination patterns and the appropriateness of the retrieved garments for a given type of social event. To address the former we rely on Shigenobu Kobayashi's color image scale, which associated emotional patterns and moods to color triples. The latter instead is analyzed by extracting garments from images of social events. Overall, we integrate in a state of the art garment recommendation framework a style classifier and an event classifier in order to condition recommendation on a given query.

**Keywords** Style · Social events · Garment recommendation · Fashion

## 1 Introduction

Fashion is a way to express personal style and to communicate emotional states that reflect personality or mood. At the same time, certain social events follow a dress code which can be strictly required or implicitly followed by participants. The rules of such social requirements however are not written and can be hard to comply with, especially for first-time attendants. Even formal events, where dressing options are more constrained, might leave

---

✉ Federico Becattini  
federico.becattini@unifi.it

Lavinia De Divitiis  
lavinia.dedivitiis@unifi.it

Claudio Baecchi  
claudio.baecchi@unifi.it

Alberto Del Bimbo  
alberto.delbimbo@unifi.it

<sup>1</sup> University of Florence, Media Integration and Communication Center (MICC), Viale Morgagni 65, Florence, Italy

space for attendants to decline an outfit according to specific preferences or personalities. Even at ceremonies such as weddings, one can adhere to the most formal requirements or play with details to give a casual or sporty touch the outfit and thus to his or her appearance. Such customization of an outfit is a mean to communicate to others a certain intent, to establish a friendly or formal environment and even to be provocative or remark a social status. Such moods and tastes are conveyed by both the shape and color of any fashion garment and may be subjected to trends and personal interpretations, yet color combinations and outfit making appear to be grounded on general clothing rules that are well established within the society.

Patterns behind color-combinations have been extensively studied, in particular with reference to different styles one may want to communicate. The seminal work by Shigenobu Kobayashi [18] introduced a scale to express emotions or attitudes based on color combinations. According to Kobayashi, just considering triplets of colors, one can identify a wide variety of lifestyles which can then be expressed by personal spaces such as interiors or offices or personal items such as outfits. On the other hand, deriving precise rules for specific events is harder since it can vary in time and follow different trends in different communities of people.

In this work we propose to follow a data driven approach to learn to model these implicit rules, both following color-combination styles and social event requirements. For color-combinations we rely on the Kobayashi's color image scale, whereas for social events we extract from online available data information about garment appearance and the respective social event. We intend to develop a complete recommendation system capable of taking into account the conveyed emotion and the compliance to societal standards for different types of events. We believe that this can come to the aid of people that seek help in shop assistants, either physically in shops or virtually on an online marketplace.

Since different modalities can convey different emotions or be suitable for different events, we also analyze the capabilities of a recommendation system to take into account also diversity. In general, diversity is an important factor to consider while tackling any information retrieval task. As suggested in [31], this also applies to benchmark datasets and the way they are exploited to produce any recommendation. The introduction of stylish aspects based on visual cues such as color combinations is also a step towards this direction.

A preliminary version of our approach was described in [9]. The system presented in this work differs substantially from [9] in several ways: (i) instead of focusing only on Kobayashi's color image scale, we present a novel analysis concerning social events by gathering the *Fashion4Events* dataset comprising approximately 400k garment images with social event labels; (ii) using the collected dataset, we train an outfit based event classifier; (iii) we integrate both the color classifier and the event classifier in our recommendation system in order to provide a more fine-grained filtering of the recommendations. We also provide an improved analysis of the state of the art.

The paper is organized as follows. In Section 2 we first provide an overview on the state of the art for fashion recommendation. An overview of our presented method is presented in Section 3. Here we introduce: (i) an outfit emotion classifier based on color combinations, capable of mapping a generic outfit onto Kobayashi's color scale; (ii) a garment based social event classifier, used to infer the event category suitable for a given outfit; (iii) the integration of such modules with a state of the art garment recommendation system. The three aspects are then detailed respectively in Sections 4, 5 and 6. The results of our experiments are reported in Section 7. Finally, we draw conclusions in Section 8

## 2 Related work

Fashion recommendation must pursue the goals of recommending pertinent items [2, 8, 27, 29], providing a set of different dressing modalities [8, 10, 25] and possibly complying with some user query or desired criteria [9, 17].

Several methods have focused on proposing garments in order to complement a given query item [10, 29, 30]. In particular, [28] focused on complementary clothing matching, starting from a top garment and recommending a bottom item. The proposed method devised a compatibility modeling scheme with attentive knowledge distillation also exploiting a teacher-student network scheme. The approach has then been improved by studying a personalized compatibility modeling, leveraging both general and subjective aesthetic preferences with a personalized compatibility modeling scheme named GP-BPR [29]. Still exploiting Bayesian Personalized Ranking (BPR), [20] used multiple autoencoder neural networks to leverage the multi-modalities of fashion items and their inter-compatibility. Following up on this line of research, PAI-BPR [26] proposed an attribute-wise interpretable compatibility scheme with personal preference modelling.

Top-bottom recommendation has also been addressed exploiting Memory Augmented Neural Networks (MANN) [8–10]. This type of architecture exploits an external memory [13, 21–24, 33] by pairing different clothing items and training a memory writing controller to store a non-redundant subset of samples. This is then used to retrieve a ranked list of suitable bottoms to complement a given top. External memories have also been used to store disentangled features for separate attributes [8, 17]. In our work we rely on a memory based approach, namely GR-MANN [10], building on top of it a style-based and event-based recommendation system. Architecture-wise, our approach shares several similarities, since it still retains its key-value structure to pair top and bottom garments. However, we introduce the novel event classifier filter and we extend the work of [9], which introduced the Kobayashi-style filter.

A similar, yet more complex task is the one of generating a whole outfit from a given garment seed. One of the first approaches [14] modeled sequences of suggestions by jointly learning a visual-semantic embedding and training a bidirectional LSTM (Bi-LSTM) model to sequentially predict the next item conditioned on previous ones. *Vasileva et al.* [30] instead learned pairwise embeddings to obtain separate representation for different pairs of fashion categories. Recent approaches tend to represent an outfit as a graph, linking fashion items among themselves if they are compatible with each other [5, 7, 34].

As for interpreting and diagnosing the proposed outfit compatibility and suggestions, [32] learns type-specified pairwise similarities between items and uses the backpropagation gradients to diagnose incompatible factors. Towards interpretable and customized fashion outfit compositions, [11] train a partitioned embedding network to favor interpretability of intermediate representations.

Finally, an interesting emerging topic is the one of understanding user reactions in order to integrate implicit user feedback into an iterative recommendation system, e.g. looking at body movements [4] or facial expressions [3].

## 3 Overview

In this paper we propose a system for compatible outfit recommendation, meaning that given a garment of a specific category (e.g., a top), we are able to propose suitable garments of a complementary category (e.g., bottoms) in order to compose an outfit. Our system pays

attention to three important aspects. First, *style-based recommendations* take into account the desired style and emotions the user would like to express. We base the recommendations on Kobayashi's Color Image Scale (CIS) [18] by training a style classifier, which we integrate into our system to condition garment proposals. Second, *event-based recommendations* focus on garments that can complement the given input and at the same time can be suitable for attending a certain kind of social event, such as a work meeting, a hike or a wedding. To this end, we train an event classifier that, given an outfit, provides a probability distribution over different event categories. Similarly to the style-based recommendation, we integrate the event classifier into our system to condition the recommendations. Third, we focus on a fundamental, yet often neglected, aspect for recommendation systems, i.e. *recommendation variety*. We rely on a model specifically designed to recommend diverse garments with few redundancies and repetitions. Rather than proposing several variations of the same outfit/style, we aim at proposing different modalities that the user can choose. To this end we introduce an entropy-based evaluation to quantify such variety.

In the following we are going to provide detailed explanations of the style and event classifiers and we are then going to illustrate how these can be integrated into a recommendation system thought for recommending variegate outfits. We carry out this study taking as reference GR-MANN [10], a recent state of the art garment recommendation system based on the usage of Memory Augmented Neural Networks.

## 4 Style-based outfit recommendation

We refer to the task of Style-based Outfit Recommendation as the task of recommending fashion items to complement an outfit, conditioned by a given style. Styles can convey moods or emotions and with this in mind we derived an interpretation of color patterns from Kobayashi's Color Image Scale [18]. We first train a style classifier which can then be combined with a recommendation system.

### 4.1 Kobayashi's color image scale

Originally presented in the early 90s', Kobayashi's Color Image Scale (CIS) [18] connected images and colors from a psychological viewpoint, investigating associations between colors and their underlying semantics. The study carried out by Kobayashi involved a decade of color-based psychophysical experiments, asking humans to annotate triplets of colors with an adjective describing perceived emotions, moods or styles. The research eventually identified a list of 1170 triplets made of 130 unique colors and associated with 180 adjectives referred to as color images. These labels have been grouped into 15 patterns representing selected terms in fashion and lifestyle and identifying clusters in a color space spanned by two orthogonal warm-cool and short-hard axes (Fig. 1).

### 4.2 Outfit style classifier

We exploit a CNN model to infer the style of an outfit. In order to analyze outfits starting from a top and a bottom, we use a concatenation of the two images depicting the two garments. Styles are instead identified by Kobayashi's patterns in the Color Image Scale. We consider each pattern as a semantic description for the style of an outfit (e.g., *casual*, *elegant*, *dandy*). Such styles indicate the feelings that an individual may want to communicate rather than describing outfit characteristics such as shape.





Fig. 1 Kobayashi’s color image scale [18] applied to outfits

In order to train such a model, we collected a set of ground truth annotations following a semi-automatic procedure. Outfit images are preprocessed by removing the background and color-quantizing foreground pixels to the palette of 130 tonalities used in CIS. We then take the three most frequent colors and compare them to the Kobayashi’s triplets using an euclidean distance in order to find the closest style characterizing the outfit:

$$d^* = \min_{p,j} \sqrt{\|P(c_o, p) - c_j\|_2} \tag{1}$$

where  $P(c, p)$  is the  $p$ -th permutation of colors in the  $c$  triplet,  $c_o$  is the triplet for outfit  $o$  and  $c_j$  the  $j$ -th of the 1170 triplets identified by Kobayashi. We retain only outfits with a clear style, i.e. if  $d^* < \theta$ . The resulting category is then mapped to one of Kobayashi’s 15 style patterns. Finally, we asked human annotators to validate the final labeling, discarding or correcting erroneous assignments.

The procedure yielded a dataset of 77390 labeled outfits for training and 10516 for testing. All outfits are taken from the IQON3000 dataset [29]. In our experiments we considered all styles except the *casual* pattern, for which enough samples are not present in the dataset, yielding to a total of 14 style categories (Fig. 2). In order to recognize outfit styles, we



**Fig. 2** Sample distribution for the training (left) and test (right) sets of the collected dataset. Categories correspond to Kobayashi's color patterns

trained a ResNet18 [16] classifier that takes as input a concatenation of the top and bottom images. We trained the network for 50 epochs using an Adam optimizer with a learning rate of 0.001.

## 5 Event-based outfit recommendation

Similarly to style-based recommendation, event-based outfit recommendation aims at proposing complementary items to generate an outfit which is suitable to attend a given type of social event. Since there is no general written rule one must adhere to when it comes to dressings for social events, we first analyze images that appear in photos taken at such events. This allows us to create a dataset of garment images paired with social event labels. Such dataset, which we refer to as *Fashion4Events* can then be used to train an event classifier capable of inferring a probability distribution over social event categories from a garment image (Fig. 3).

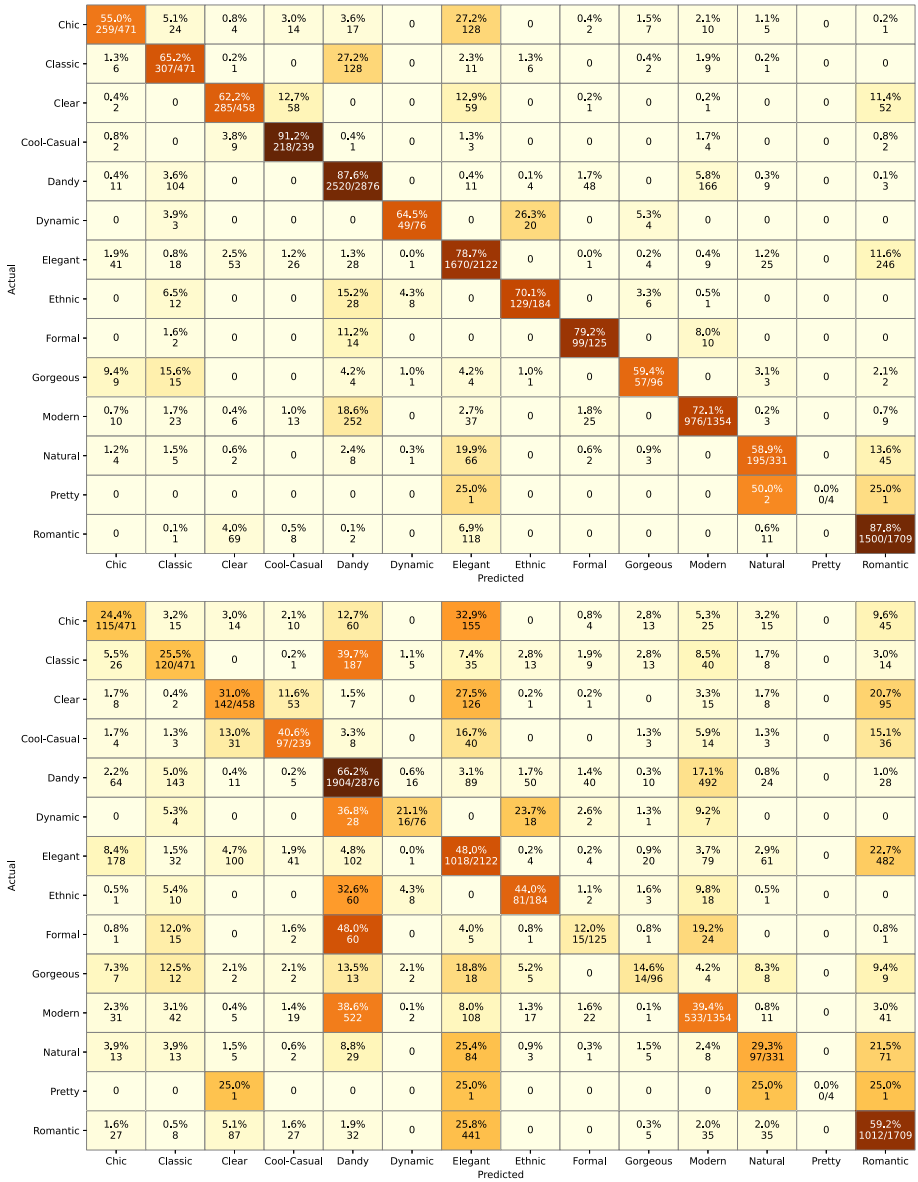
### 5.1 Dataset creation

To collect *Fashion4Events*, a dataset of garment images paired with social event labels, we exploited two different sources of data: the DeepFashion2 dataset [12] and the USED dataset [1].

DeepFashion2 is a dataset that proposed an unified benchmark for clothes detection, segmentation, retrieval and landmark prediction. It contains approximately 491K images of clothes, divided into train (391K), validation (34K) and test (67K) belonging to 13 different classes. Garments exhibit large variations in style, pose, scale, color, occlusion and viewing angle.

We use DeepFashion2 to train a garment detection and segmentation network by fine-tuning a Mask-RCNN model [15]. The model has been pre-trained on COCO [19] using a ResNet50 backbone.


Once we have a model capable of detecting and segmenting individual garments in images, we apply it on the USED dataset [1]. USED is a dataset comprising 525K images of people attending 14 different types of social events. Images have been downloaded from Flickr and event types have been selected among the most common ones on social media.



**Fig. 3** Confusion matrix for the outfit style classifier. Top: our ResNet18-based style classifier. Bottom: nearest neighbor

The categories are the following: concert, graduation, meeting, mountain-trip, picnic, sea-holiday, ski-holiday, wedding, conference, exhibition, fashion, protest, sport and theater-dance. Images represent both indoor and outdoor scenes and the dataset accounts for a variable number of people in each image.

By applying Mask-RCNN on the images from USED, we obtain a set of detected garments paired with pixelwise segmentations and a social event label. We retain only




**Fig. 4** Garments are detected and segmented using Mask-RCNN. For each detected garment, a new image is created. Event labels associated with the original images are associated with the detected garments

detections with a confidence score higher than 0.8 in order to remove noise. For each detected garment, we generate a separate image by taking the segmented clothing item on a black background, as shown in Fig. 4. The final set of detections yields a dataset of 317,174 images for training and 94,592 for testing. In Fig. 5 we show the distribution of samples over the 14 event categories in the Fashion4Events dataset.

**5.2 Social event style classification**

Thanks to the Fashion4Events dataset, which pairs segmented garment images and social event labels, we can train a garment-based event classifier. The aim of such a classifier is to provide a probability distribution of social event categories based on the style of a given garment or outfit. We trained a ResNet18 [16] model taking outfit images as inputs. As for outfit styles, we feed outfits to the model as horizontally concatenated top and bottom images. We started from a pre-trained model on ImageNet by removing the original classification head and adding two fully connected layers with respectively 512 and 14 neurons. The Adam optimizer was used to train the model, with a learning rate of 0.0005 and a weight decay of 0.0001 for 32 epochs in total.



**Fig. 5** Sample distribution for the training (left) and test (right) sets of the collected Fashion4Events dataset. Categories correspond to social events



## 6 Style-based and event-based outfit recommendation

We integrate our outfit style classifier and outfit event classifier in the state of the art garment recommendation system GR-MANN [10]. Note that in principle the classifiers could be applied to any recommendation system capable of generating an outfit composed of a top and a bottom garment.

GR-MANN is a neural network based on a persistent external memory in which non-redundant samples are stored to guide recommendations. The samples in memory express different modalities to combine tops and bottoms. At inference time, the query top is presented to the model and encoded with a convolutional encoder. Its feature is then used as key to retrieve similar tops in memory and access the correspondent stored bottoms.

We extend GR-MANN by using the outfit style and event classifiers to check whether the recommended outfits comply to a certain criterion requested by the user. Once a bottom is proposed, it is concatenated to the input top and fed to the classifiers. We perform different fusion methods, in order to characterize the behaviour of the system, as detailed in Section 7.

The choice of exploiting GR-MANN stems from the fact that its memory is populated by a controller, trained to store only relevant samples. Classifying such samples based on style and event type will therefore allow us to quantify the diversity that GR-MANN strives to achieve in its recommendations.

## 7 Experiments

We demonstrate our method on the IQON3000 dataset [29], performing several evaluations. At first we discuss the accuracy of the style classifier and the event classifier and then we combine them with GR-MANN to analyze their capability to recommend a diverse set of bottom garments covering multiple styles. In order to evaluate the performance of the event classifier we rely on the event-garment dataset based on USED, obtained as explained in Section 5.1.

### 7.1 Style classifier evaluation

In Fig. 3 the confusion matrix for the style classifier is shown. As a reference, a Nearest Neighbor (NN) baseline is also reported. Here we simply extract features from a standard ResNet, pretrained on ImageNet, and transfer the style category from the closest sample in the training set onto a given test outfit. Our model achieves an accuracy of 78.56% while the NN baseline only 49.11%, highlighting how the task is not as straightforward as comparing visual features.

Interestingly, most of the errors committed by the model tend to confuse similar categories, such as *elegant* and *chic* or *classic* and *dandy*.

### 7.2 Event classifier evaluation

Overall, the event classifier obtained a 49.9% accuracy on the test set of the garment-event dataset (Section 5.1). As can be seen in the confusion matrix in Fig. 6, the most difficult event class appears to be *Conference*, which is often confused with *Meeting*. This is easily

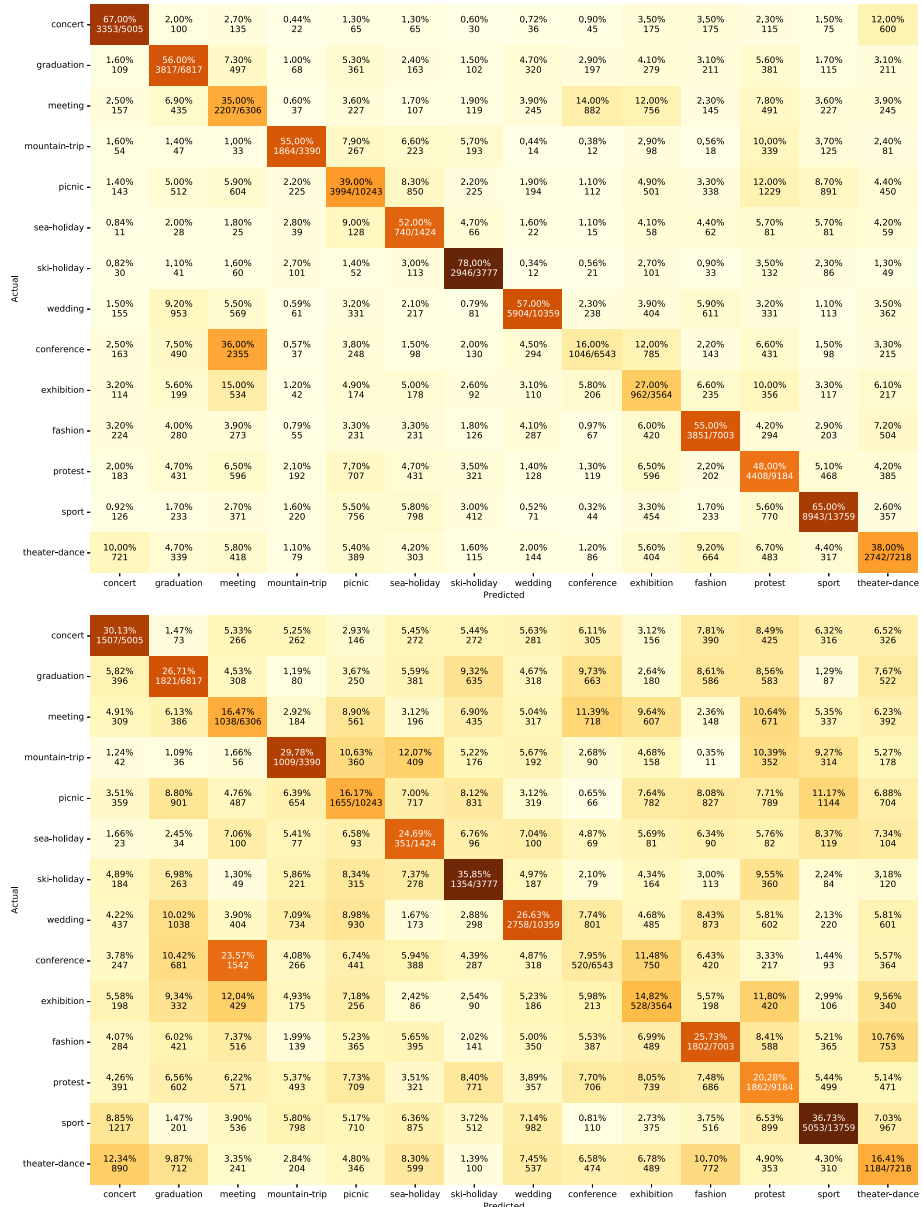


Fig. 6 Confusion matrix for the event classifier. Top: our ResNet18-based style classifier. Bottom: nearest neighbor

explainable given that these two social events are work related and thus have similar dress-codes. In the Nearest Neighbor baseline, the accuracy drops to 23.45% and the confusion among related classes is more present, such as *Theatre-dance* and *Concert* or *Exhibition* and *Meeting*.

### 7.3 Style-based and event-based outfit recommendation evaluation

We now assess the capabilities of the GR-MANN recommendation system with reference to styles and social events. First, we measure how the recommendation system is able to suggest garments that comply with the ground truth style and event category, without adding any prior knowledge to the model. We report Accuracy, measuring if at least one of the recommended outfits adheres with the requested category, and mAP, which takes the ranking of the correct bottoms into account.

In Table 1 we compare the results for style-based recommendation obtained by the model against a baseline in which style categories are drawn at random. It can be seen that, both for Accuracy and for mAP, the results improve considerably and that the model is able to provide at least an outfit with the desired style most of the times even with only 5 recommendations.

As for event-based recommendation, we report a similar analysis in Table 2. Also in this case, the model generates a set of recommendations comprising a suitable bottom to comply with the social event of interest.

Additionally, following the evaluation protocol of [10], we also measure accuracy color-wise, category-wise and combining both together. However, we filter the output of GR-MANN in order to provide a ranked list of bottoms with the correct category using the style classifier or the event classifier. Therefore, in this experiment we are relaxing the formulation of the task, assuming that the desired style is known a-priori. The rationale behind this evaluation is to see if the proposed garments share similar visual traits with the ground truth when performing a category-conditioned (either style-based or event-based) recommendation. In Tables 3 and 4 we report results for both accuracy and mAP. Interestingly, the two approaches obtain similar results, with slightly lower accuracies for the event-based filtering strategy. In particular, both approaches are capable of retrieving garments with either the same color or the same category as the ground truth in approximately 80% of the cases, even when limiting the result list to the top-5 elements. We also report similar results for a combined approach, i.e. using both the style and event classifiers to filter out results, as shown in Table 5.

### 7.4 Recommendation diversity evaluation

As studied in [31], diversity is an important aspect of information retrieval systems. To this end, we also perform an evaluation of the entropy of the proposed labels to establish the variation degree of our proposals. This evaluation was first proposed in [6] to perform an unsupervised assessment of a generic classifier with unsupervised data. Given a probability distribution  $X = \{x_1, x_2, \dots, x_n\}$  over  $N$  different classes, we can compute the Shannon entropy  $H$  for the probability vector  $X$  as  $H(X) = -\sum_{i=1}^N x_i \log(x_i)$ . The entropy will be 0 when all samples are labeled with the same class, and will increase as more information and diversity are introduced in the predictions. Ideally we would like to stay as close as possible to the entropy of any random label permutations, but preserving good recommendation results. Results, shown in Tables 6 and 7, show that our method is able to maintain a reasonable amount of entropy in the predictions while performing significantly better than random, as shown in Tables 1 and 2.

Interestingly, the entropy for Kobayashi's style categories is higher than the one relative to event categories. We attribute this to the distribution of garments in IQON3000, for which certain categories such as ski-holiday and wedding are underrepresented.

**Table 1** Accuracy and mAP obtained by GR-MANN [10] on the IQON3000 [29] dataset

Num Items	5	10	20	30	40	50	60
Accuracy	73.42	84.17	91.47	94.22	95.73	96.70	97.39
Random Acc.	36.50	53.12	80.37	91.14	96.51	99.02	99.89
mAP	48.17	46.35	43.46	41.78	40.70	39.96	39.40
Random mAP	18.14	17.21	17.04	16.59	14.39	13.41	11.46

The metrics are computed in order to retrieve an outfit with the same style of the ground truth

**Table 2** Accuracy and mAP obtained by GR-MANN [10] on the IQON3000 [29] dataset

Num Items	5	10	20	30	40	50	60
Accuracy	75.83	88.50	94.96	95.06	96.72	97.24	97.70
Random Acc.	37.12	53.03	79.87	91.21	96.73	98.99	99.16
mAP	43.05	40.00	36.66	34.25	33.07	32.30	31.42
Random mAP	18.31	16.87	17.34	16.44	14.12	13.36	11.07

The metrics are computed in order to retrieve an outfit with the same social event category of the ground truth

**Table 3** Accuracy and mAP varying the number or retrieved items

Num Items	5	10	20	30	40	50	60
Cat × Col Acc.	57.66	59.22	61.92	64.37	66.42	68.36	69.97
Category Accuracy	83.92	84.81	86.27	87.60	88.68	89.67	90.39
Color Accuracy	81.41	82.73	84.71	86.37	87.69	88.82	89.77
mAP	18.50	18.48	18.37	18.23	18.11	17.98	17.88

All proposals are filtered by the style classifier in order to share the desired one

**Table 4** Accuracy and mAP varying the number or retrieved items

Num Items	5	10	20	30	40	50	60
Cat × Col Acc.	52.00	53.70	55.46	57.92	60.86	62.74	64.68
Category Accuracy	86.14	86.36	88.24	88.78	89.98	90.76	90.40
Color Accuracy	75.92	77.70	79.96	81.58	83.78	84.74	86.46
mAP	16.02	15.92	15.47	15.33	15.57	15.42	14.98

All proposals are filtered by the event classifier in order to share the desired category

**Table 5** Accuracy and mAP varying the number of retrieved items

Num Items	5	10	20	30	40	50	60
Cat × Col Acc.	49.86	50.91	52.51	55.04	58.72	60.37	63.12
Category Accuracy	82.79	83.22	84.16	85.40	86.83	87.91	89.77
Color Accuracy	75.64	77.31	79.62	80.94	83.17	84.11	85.38
mAP	15.44	14.76	14.60	14.28	14.21	14.16	13.83

All proposals are filtered by both the style classifier and the event classifier in order to share the desired categories

**Table 6** Entropy of the recommendations with reference to outfit styles

Method \ Num Items	5	10	20	30	40	50	60
Random	1.419	1.905	2.255	2.391	2.458	2.499	2.523
Style-Entropy	0.849	1.061	1.152	1.265	1.267	1.288	1.317

A sufficiently high entropy indicates variety in the proposed outfits

**Table 7** Entropy of the recommendations with reference to outfit event categories

Method \ Num Items	5	10	20	30	40	50	60
Random	1.353	1.839	2.189	2.325	2.392	2.433	2.457
Event-Entropy	0.671	0.819	0.908	0.941	0.958	0.965	0.970

A sufficiently high entropy indicates variety in the proposed outfits

## 8 Conclusions

In this paper we presented an approach to take into account a style based filtering and an event based filtering for fashion recommendation. Styles can convey a color-based mood according to Kobayashi's Color Image Scale or can reflect dress codes for specific social event categories. We leveraged the work of Kobayashi to train a style classifier that we used to filter the results of a memory network based garment recommender. Similarly, we trained a garment-based event classifier to be combined with the recommender by exploiting a garment detector and Fashion4Events, an image dataset of annotated social events. Experiments show that our system is able to generalise on color styles and social events and that the recommendation system is able to propose a variety of outfit styles compatible with the query garment.

**Acknowledgments** This work was partially supported by the Italian MIUR within PRIN 2017, Project Grant 20172BH297: I-MALL - improving the customer experience in stores by intelligent computer vision.

**Data Availability** The datasets generated and analysed during the current study are available in the Fashion4Events repository: <https://github.com/fedebecat/Fashion4Events>.

## Declarations

**Conflict of Interests** The authors declare that they have no conflict of interest.

## References

1. Ahmad K, Conci N, Boato G, De Natale FG (2016) Used: a large-scale social event detection dataset. In: Proceedings of the 7th International conference on multimedia systems, pp 1–6
2. Baldrati A, Bertini M, Uricchio T, Del Bimbo A (2021) Conditioned image retrieval for fashion using contrastive learning and clip-based features. In: ACM Multimedia Asia, pp 1–5
3. Becattini F, Song X, Baccchi C, Fang S-T, Ferrari C, Nie L, Del Bimbo A (2021) Plm-ipe: a pixel-landmark mutual enhanced framework for implicit preference estimation. In: ACM Multimedia Asia, pp 1–5
4. Bigi W, Baccchi C, Del Bimbo A (2020) Automatic interest recognition from posture and behaviour. In: Proceedings of the 28th ACM International conference on multimedia, pp 2472–2480
5. Cucurull G, Taslakian P, Vazquez D (2019) Context-aware visual compatibility prediction. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition, pp. 12617–12626
6. Cuffaro G, Becattini F, Baccchi C, Seidenari L, Del Bimbo A (2016) Segmentation free object discovery in video. In: European conference on computer vision, pp 25–31. Springer
7. Cui Z, Li Z, Wu S, Zhang X, Wang L (2019) Dressing as a whole: Outfit compatibility learning based on node-wise graph neural networks. In: The Web Conference 2019 - Proceedings of the World Wide Web Conference, WWW 2019, pp. 307–317. <https://doi.org/10.1145/3308558.3313444>
8. De Divitiis L, Becattini F, Baccchi C, Bimbo AD (2022) Disentangling features for fashion recommendation. ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM)
9. De Divitiis L, Becattini F, Baccchi C, Del Bimbo A (2021) Style-based outfit recommendation. In: 2021 International Conference on Content-Based Multimedia Indexing (CBMI), pp. 1–4. IEEE
10. Divitiis LD, Becattini F, Baccchi C, Bimbo AD (2020) Garment recommendation with memory augmented neural networks. In: Pattern recognition. ICPR International workshops and challenges - virtual event, January 10–15, 2021, proceedings, Part II. Lecture notes in computer science, vol 12662, pp 282–295. Springer. <https://doi.org/10.1007/978-3-030-68790-8.23>
11. Feng Z, Yu Z, Yang Y, Jing Y, Jiang J, Song M (2018) Interpretable partitioned embedding for customized fashion outfit composition. arXiv:1806.04845

12. Ge Y, Zhang R, Wang X, Tang X, Luo P (2019) Deepfashion2: a versatile benchmark for detection, pose estimation, segmentation and re-identification of clothing images. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition, pp 5337–5345
13. Graves A, Wayne G, Danihelka I (2014) Neural turing machines. arXiv:1410.5401
14. Han X, Wu Z, Jiang Y-G, Davis LS (2017) Learning fashion compatibility with bidirectional lstms. In: Proceedings of the 25th ACM International conference on multimedia, pp 1078–1086
15. He K, Gkioxari G, Dollár P, Girshick R (2017) Mask r-cnn. In: Proceedings of the IEEE International conference on computer vision, pp 2961–2969
16. He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on computer vision and pattern recognition, pp 770–778
17. Hou Y, Vig E, Donoser M, Bazzani L (2021) Learning attribute-driven disentangled representations for interactive fashion retrieval. In: Proceedings of the IEEE/CVF International conference on computer vision, pp 12147–12157
18. Kobayashi S (2009) Color image scale. [http://www.ncd-ri.co.jp/english/main\\_0104.html](http://www.ncd-ri.co.jp/english/main_0104.html)
19. Lin T-Y, Maire M, Belongie S, Hays J, Perona P, Ramanan D, Dollár P, Zitnick CL (2014) Microsoft coco: Common objects in context. In: Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13, pp 740–755. Springer
20. Liu J, Song X, Chen Z, Ma J (2019) Neural fashion experts: I know how to make the complementary clothing matching. *Neurocomputing* 359:249–263
21. Marchetti F, Becattini F, Seidenari L, Del Bimbo A (2020) Mantra: Memory augmented networks for multiple trajectory prediction. In: Proceedings of the IEEE Conference on computer vision and pattern recognition
22. Marchetti F, Becattini F, Seidenari L, Del Bimbo A (2020) Multiple trajectory prediction of moving agents with memory augmented networks. *IEEE Trans Pattern Anal Mach Intell*
23. Marchetti F, Becattini F, Seidenari L, Del Bimbo A (2022) Smemo.: Social memory for trajectory forecasting. arXiv:2203.12446
24. Marchetti F, Becattini F, Seidenari L, Del Bimbo A (2023) Explainable sparse attention for memory-based trajectory predictors. In: Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part V, pp 543–560. Springer
25. Sá J, Queiroz Marinho V, Magalhães AR, Lacerda T, Goncalves D (2022) Diversity vs relevance: a practical multi-objective study in luxury fashion recommendations. In: Proceedings of the 45th International ACM SIGIR Conference on research and development in information retrieval, pp 2405–2409
26. Sagar D, Garg J, Kansal P, Bhalla S, Shah RR, Yu Y (2020) Pai-bpr: Personalized outfit recommendation scheme with attribute-wise interpretability. In: 2020 IEEE 6th International Conference on Multimedia Big Data (BigMM), pp 221–230. IEEE
27. Sarkar R, Bodla N, Vasileva M, Lin Y-L, Beniwal A, Lu A, Medioni G (2022) Outfittransformer: Outfit representations for fashion recommendation. In: Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition, pp 2263–2267
28. Song X, Feng F, Han X, Yang X, Liu W, Nie L (2018) Neural compatibility modeling with attentive knowledge distillation. In: The 41st international ACM SIGIR conference on research & development in information retrieval, pp 5–14
29. Song X, Han X, Li Y, Chen J, Xu X-S, Nie L (2019) Gp-bpr: Personalized compatibility modeling for clothing matching. In: Proceedings of the 27th ACM International conference on multimedia, pp 320–328
30. Vasileva MI, Plummer BA, Dusad K, Rajpal S, Kumar R, Forsyth D (2018) Learning type-aware embeddings for fashion compatibility. In: Proceedings of the European Conference on Computer Vision (ECCV), pp 390–405
31. Wang JZ, Boujema N, Del Bimbo A, Geman D, Hauptmann AG, Tesić J (2006) Diversity in multimedia information retrieval research. In: Proceedings of the 8th ACM International Workshop on Multimedia Information Retrieval. MIR '06, pp 5–12. Association for Computing Machinery. <https://doi.org/10.1145/1178677.1178681>
32. Wang X, Wu B, Zhong Y (2019) Outfit compatibility prediction and diagnosis with multi-layered comparison network. In: Proceedings of the 27th ACM International conference on multimedia, pp 329–337
33. Weston J, Chopra S, Bordes A (2014) Memory networks. arXiv:1410.3916
34. Yang X, Du X, Wang M (2020) Learning to match on graph for fashion compatibility modeling. In: Proceedings of the AAAI Conference on artificial intelligence 34(01), 287–294. <https://doi.org/10.1609/aaai.v34i01.5362>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.