

POG: Personalized Outfit Generation for Fashion Recommendation at Alibaba iFashion

Wen Chen, Pipei Huang*

Jiaming Xu, Xin Guo, Cheng Guo, Fei Sun, Chao Li, Andreas Pfadler, Huan Zhao, Binqiang Zhao

Alibaba Group, Beijing, China

{chenyu.cw, pipei.hpp, jiaming.xjm, gary.gx, mike.gc, ofey.sf, ruide.lc, andreaswernerrober, chuanfei.zh, binqiang.zhao}@alibaba-inc.com

ABSTRACT

Increasing demand for fashion recommendation raises a lot of challenges for online shopping platforms and fashion communities. In particular, there exist two requirements for fashion outfit recommendation: the *Compatibility* of the generated fashion outfits, and the *Personalization* in the recommendation process. In this paper, we demonstrate these two requirements can be satisfied via building a bridge between outfit generation and recommendation. Through large data analysis, we observe that people have similar tastes in individual items and outfits. Therefore, we propose a **Personalized Outfit Generation (POG)** model, which connects user preferences regarding individual items and outfits with Transformer architecture. Extensive offline and online experiments provide strong quantitative evidence that our method outperforms alternative methods regarding both compatibility and personalization metrics. Furthermore, we deploy POG on a platform named Dida in Alibaba to generate personalized outfits for the users of the online application iFashion.

This work represents a first step towards an industrial-scale fashion outfit generation and recommendation solution, which goes beyond generating outfits based on explicit queries, or merely recommending from existing outfit pools. As part of this work, we release a large-scale dataset consisting of 1.01 million outfits with rich context information, and 0.28 billion user click actions from 3.57 million users. To the best of our knowledge, this dataset is the largest, publicly available, fashion related dataset, and the first to provide user behaviors relating to both outfits and fashion items.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; Multimedia content creation;

KEYWORDS

Fashion Outfit Generation; Fashion Outfit Recommendation; Deep Learning; Transformer; Self-attention

*Pipei Huang is the Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD '19, August 4–8, 2019, Anchorage, AK, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6201-6/19/08...\$15.00

<https://doi.org/10.1145/3292500.3330652>

ACM Reference Format:

Wen Chen, Pipei Huang and Jiaming Xu, Xin Guo, Cheng Guo, Fei Sun, Chao Li, Andreas Pfadler, Huan Zhao, Binqiang Zhao. 2019. POG: Personalized Outfit Generation for Fashion Recommendation at Alibaba iFashion. In *The 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '19)*, August 4–8, 2019, Anchorage, AK, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3292500.3330652>

1 INTRODUCTION

Faced with seemingly abundant choices of ever changing styles, fashion outfit recommendation becomes more and more important for modern consumers and has thus attracted interest from the online retail industry. Fashion outfit is a set of fashion items, which appears both visually compatible and functionally irredundant [3] (see Figure 1 as an example). Compared to traditional item recommendation, fashion outfit recommendation involves a remarkably creative outfit generation process, which requires both innovation and characteristic. Therefore, this task is usually performed by fashion experts, and becomes popular in numerous online fashion communities, such as Lookbook¹ and Chictopia². In Taobao, the largest online consumer-to-consumer platform in China, a new application iFashion is created to support fashion outfit recommendation. Approximately 1.5 million content creators were actively supporting Taobao as of March 31, 2018³. Still, the great gap between limited human labor and ever-growing market demands a complementary or even substitution of manual work. To alleviate this problem, we aim to assist the generation and recommendation process in iFashion.

There exist two requirements in fashion outfit generation and recommendation: 1) the *Compatibility* of the generated fashion outfits, 2) the *Personalization* in the recommendation process. *Compatibility* is a measurement of how harmonious a set of items is. Early studies mainly focus on learning compatibility metric between pairwise items [17, 30], or predicting the popularity of an outfit [11]. Some recent works attempt to generate outfits by modeling an outfit as an ordered sequence of items [3, 18]. However, it is not reasonable because shuffling items in the outfit should make no difference on its compatibility. *Personalization* represents how the recommendations meet users' personal fashion tastes. In recent works, personalization is achieved relying on explicit input (e.g., image or text) provided by the user [4, 15, 18]. This type of fashion generation works more like search than recommendation since it needs explicit user queries. On the other hand, classic Collaborative

¹<https://lookbook.nu/>

²<http://chictopia.com/>

³https://www.alibabagroup.com/en/news/press_pdf/p180504.pdf



Figure 1: A sample of iFashion application in Taobao. We recommend fashion outfits (sets of fashion items which interact with each other) to users.

Table 1: The percentage of the same brand, category, style, and pattern of the items in users' clicked outfits appearing in users' latest clicked items.

| Property | brand | category | style | pattern |
|------------|-------|----------|-------|---------|
| Percentage | 56.9% | 81.3% | 73.0% | 53.9% |

Filtering (CF) methods [14, 22] in recommendation mostly focus on recommending individual items rather than outfits.

In this work, we collect 1.21 billion user clicks on 4.68 million items and 192 thousand outfits from 5.54 million users in iFashion. As shown in Table 1, 81.3% of the clicked outfits contain the items with the same categories which have been appeared in the latest 50 clicked items from the same user. We observe that the brand, style, and pattern of the items in users' clicked outfits also have high probabilities to appear in users' latest clicked items (we choose the latest 50 items for each user). Figure 2 illustrates three example users with their behaviors on items and outfits. All these results show that users tend to keep similar tastes in individual items and outfits. We find that although fashion outfit generation and recommendation have been studied intensively in recent years, existing works usually study these two requirements separately.

Therefore, we attempt to build the bridge between fashion outfit generation and recommendation in a real-world application with millions of users. More specifically, we generate personalized outfits by capturing users' interests and tastes from their historical interactions on fashion items. For the *Compatibility* requirement, we propose a **Fashion Outfit Model (FOM)** by learning the compatibilities between each item and all the other items within the outfit. Each item should have different weighted interactions to the other items in the outfit. Thus, we set up a masked item prediction task based on the self-attention mechanism [29], which masks one item at a time in the outfit, and predicts the masked item based on the context from other items in the outfit. For the *Personalization* requirement, by integrating user preference into the pre-trained FOM, we propose a **Personalized Outfit Generation (POG)** model, which can generate compatible and personalized

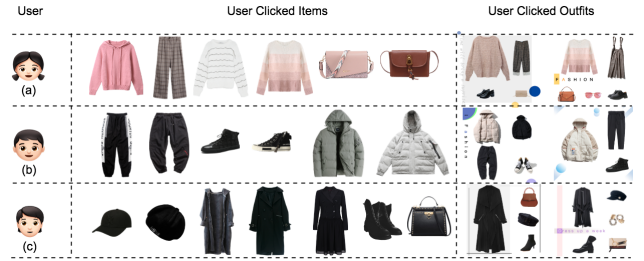


Figure 2: Illustration of user clicked items and outfits. User (a) is a young girl who likes clothes and outfits in light colors and sweet styles. User (b) is a college boy who clicks several winter outfits after clicking a lot of winter clothes. User (c) is probably an office lady who prefers OL style items as well as outfits.

outfits based on users' recent behaviors. Specifically, POG uses a Transformer encoder-decoder architecture [29] to model both signals from user preference and outfit compatibility. To the best of our knowledge, this is the first study to generate personalized outfits based on users' historical behaviors with an encoder-decoder framework. Last, we develop a platform named *Dida*, where POG has been deployed, to assist outfit generation and recommendation in large-scale online application iFashion.

The contributions of this work are summarized as follows:

- (1) We propose POG: an encoder-decoder model to generate personalized fashion outfits, which takes into account both outfit compatibility and user personalization. It makes fashion recommendations by generating personalized outfits based on users' recent behaviors.
- (2) We demonstrate that our model significantly outperforms other alternative methods through outfit compatibility experiments, including pushing the FITB (Fill In The Blanks) benchmark to 68.79% (5.98% relative improvement) and CP (Compatibility Prediction) benchmark to 86.32% (25.81% relative improvement).
- (3) We deploy POG on the real-world outfit generation platform *Dida*. Through extensive online experiments, we show that POG clearly outperforms the CF method by 70% increase in CTR (Click-Through-Rate) metric.
- (4) We release a dataset⁴ of 1.01 million outfits, 583 thousand fashion items associated with rich context information, and 0.28 billion user click actions from 3.57 million users.

2 RELATED WORK

Fashion is an important application domain of computer vision and multimedia. Much research effort has been made in this domain, focusing on fashion image retrieval [2, 32], clothing recognition [16], clothing parsing [12, 33], attribute learning [5, 15], outfit compatibility [3, 17, 27, 28, 30], and fashion recommendation [4, 8, 15]. The goal of this work is to compose personalized fashion outfits automatically based on user behaviors, we hence focus on the research areas of outfit generation and recommendation.

⁴<https://github.com/wenyuer/POG>

2.1 Fashion Outfit Generation

Methods for fashion outfit generation usually fall within one of two categories. Methods in the first category focus on calculating a pairwise compatibility metric [17, 23, 30]. McAuley et al. extract visual features to model human visual preference for a pair of items of the Amazon co-purchase dataset [17]. Siamese network [30] estimates pairwise compatibility based on co-occurrence in large-scale user behavior data. Methods belonging to the second category, such as presented in [11] and [3], are based on modeling a fashion outfit as a set or an ordered sequence. Li et al. deploy an end-to-end deep learning system which can classify a given outfit as popular or unpopular [11]. Han et al. train a bidirectional LSTM model to sequentially generate outfits [3]. These methods generally use simple pooling of item vectors to represent an outfit, or rely heavily on the order of the outfit items. We note that methods belonging to either category hardly consider all the interactions among the items in an outfit. Besides, it is unreasonable to regard an outfit as an ordered sequence, because shuffling items in the outfit should make no difference on its compatibility. We try to explicitly incorporate this into our modeling architecture by requiring that each item should have different interaction weights with respect to other items in one outfit. For example, a “red shirt” should have a higher interaction weight with “blue jeans”, but a smaller weight with a pair of “white gloves”.

2.2 Fashion Outfit Recommendation

Early works on recommendation typically use collaborative filtering to model users’ preferences based on their behavior histories [14, 22]. However, previous works mostly restrict attention to recommending individual items. There are a few approaches for recommending the whole fashion outfits. Si Liu et al. [15] propose an occasion-oriented clothing recommendation method based on attributes and categories. The work in [17] and [3] requires image or text queries to find complementary clothes. A functional tensor factorization approach is used to suggest sets of items to users in [4]. As these approaches all require user queries or user uploaded data as input, they are likely to be perceived as less user-friendly by a typical user. Moreover, we note that these approaches are rather impractical to implement in an efficient manner in a large-scale online recommender system.

2.3 Self-Attention and Transformer

Self-attention is an attention mechanism relating different positions of a single sequence [29], which has been used successfully in a variety of tasks [13, 19]. Transformer [29], a transduction model relying on self-attention, has been widely used and greatly improved the performance for language processing tasks [1, 21]. In recent language representation research, bidirectional Transformer encoder has exhibited the best performance (BERT [1]) when compared to left-to-right Transformer decoder (OpenAI GPT [21]) and bidirectional LSTM (ELMo [20]). Recently, multi-head self-attention is also introduced to model users’ behavior sequences for sequential recommendation [9, 25]. Different from them, in this paper, we adopt the self-attention to model the compatibility in fashion outfit generation.

Table 2: Statistics of the datasets.

| Dataset | #Outfits | #Users | #Items |
|-------------|-----------|-----------|-----------|
| Outfit data | 1,013,136 | - | 583,464 |
| Item data | - | - | 4,747,039 |
| User data | 127,169 | 3,569,112 | 4,463,302 |

3 DATASET

Taobao’s fashion experts create thousands of outfits everyday. All these manually created fashion outfits are reviewed before they are exhibited online. About 1.43 million outfits have so far been created and reviewed in this manner. We collect the 80 most frequent leaf categories (e.g. sweater, coat, t-shirt, boots, and ring) from the items in all these outfits. All items in the less frequent categories are removed from the outfits. In the rest of the paper, we only consider the items in these 80 leaf categories. After that, we filter the outfits which contain fewer than 4 items. In total, 1.01 million outfits remained, which are composed of 583 thousand individual items.

Moreover, we collect clicks on the items and outfits from iFashion’s users in recent three months. We select click behaviors from 3.57 million active users who have viewed more than 40 outfits in total. The clicks on outfits are recorded only when more than 10 item clicks happened before from the same user. Then, we build a training sample by pairing an outfit click with the latest 50 item clicks prior to it from the same user. Finally, we obtain 19.2 million training samples which consist of 4.46 million items and 127 thousand outfits. Every item in our dataset is associated with white background image, title, and leaf category.

To the best of our knowledge, our dataset is the largest publicly available dataset of fashion items with rich information compared to existing datasets, such as WoW [15], Fashion-136K [6], FashionVC [24], Maryland Polyvore [3], Polyvore Outfits-D, and Polyvore Outfits [28]. Moreover, we are the first to provide outfit data and associated user behavior data, which can be exploited for future fashion recommendation research. We provide three datasets describing the outfits, the items, and the user behaviors separately. The statistics of the datasets are shown in Table 2.

4 METHODOLOGY

In this section, we introduce our methodology in detail. POG is built within a three-step process: we first embed the items. Second, we build FOM, which learns compatibilities of items within an outfit. Last, once training is completed, we use the resulting pre-trained FOM to initialize POG on a Transformer architecture.

As a first step, we represent all the items using a multi-modal embedding model. Then we introduce FOM and POG in detail. Last, we introduce our Dida platform in this section, which helps to ensure the efficiency and quality requirements in large-scale online application of this work.

4.1 Multi-modal Embedding

For every fashion item f , we compute a nonlinear feature embedding \mathbf{f} . The concept of fashion relies mostly on visual and textual information. Most previous works suggest to leverage image and text to learn multi-modal embeddings [3, 11]. In our work, we use a multi-modal embedding model that takes the following input for

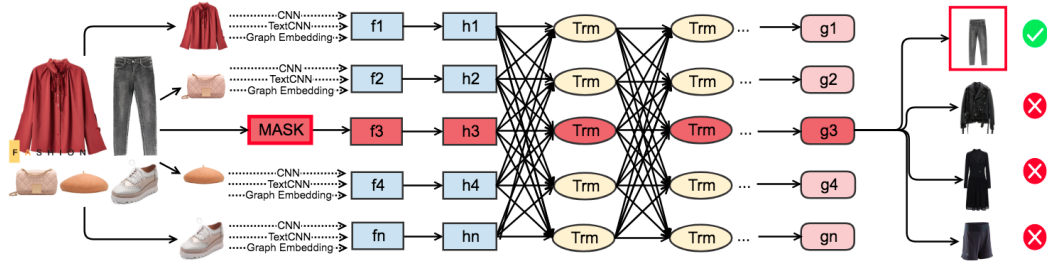


Figure 3: The architecture of FOM. We mask the items in the outfit one at a time. For example, we mask a pair of jeans in the outfit. The model is learned to choose the correct jeans from a candidate pool, to complement other items in the outfit.

every item: (1) dense vector encoding the white background picture of the item from a CNN model, (2) dense vector encoding the title of the item obtained from a TextCNN network, which has been pre-trained to predict an item’s leaf category based on its title, (3) dense vector encoding a collaborative filtering signal for the item using Alibaba’s proprietary Behemoth Graph Embedding platform [31], which generates item embeddings based on the co-occurrence statistics of items in recorded user click sessions in the Taobao Mobile Application.

Our goal is to obtain an embedding space, where similar items are embedded nearby, and different items lie in different regions. We concatenate the embeddings derived from image, text, and CF as input for a final fully connected layer. The output is a d_e -dimension vector f . For each item f , we define positive samples f^+ as those items which belong to the same leaf category as f , and negative samples f^- are hence those items which do not fall into the same leaf category. The entire network is then trained using the triplet loss via:

$$\mathcal{L}_E = \sum_f \max(d(f, f^+) - d(f, f^-) + \alpha, 0) \quad (1)$$

where the distance metric d represents Euclidean distance, and α as the margin. By minimizing \mathcal{L}_E , the distance between f and f^+ in the same category is forced to be smaller than the distance from f^- in a different category by some margin α .

4.2 FOM: Fashion Outfit Model

Outfit is a set of items, where each item should have different weighted interactions to the other items in the outfit. To capture the item interactions in an outfit, we design a masked item prediction task based on a bidirectional Transformer encoder architecture. Masking items one at a time in the outfits, we require the model to fill in the blank with the correct item according to the context. Since every item in the outfit is masked to fuse its left and right context, the compatibility between each item and all the other items within the outfit can be learned from the self-attention mechanism. Hence, the compatibility in the outfit can be learned by combining all these item compatibilities.

Let \mathcal{F} denote the set of all outfits. Given an outfit $F \in \mathcal{F}$, $F = \{f_1, \dots, f_t, \dots, f_n\}$, where f_t is the t -th item. Let f_t be the representation of item f_t derived from multi-modal embedding with dimension d_e . We use a particular embedding [MASK] for the masked

item. Non-masked items are represented by their multi-modal embeddings. We then represent the set of input embeddings as F_{mask} . Given F_{mask} , the task is to predict the masked item rather than reconstructing the entire outfit. More formally, we minimize the following loss function of FOM:

$$\mathcal{L}_F = -\frac{1}{n} \sum_{mask=1}^n \log \Pr(f_{mask} | F_{mask}; \Theta_F) \quad (2)$$

where Θ_F denotes the model parameters, and $\Pr(\cdot)$ is the probability of choosing the correct item conditioned on the non-masked items.

The model architecture is shown in Figure 3. We do not use position embedding like Transformer [29] does, because we take the items in the outfit as a set, not a sequence with position information. Let F_{mask} pass through two fully-connected layers with a Rectified Linear Unit (ReLU) activation in between, to transfer all the input embeddings from the single item space to an outfit space. We call the two fully-connect layers as transition layer, and represent the output as $H^0 \in \mathbb{R}^{n \times d_m}$:

$$H^0 = \text{ReLU}([f_1^T; \dots; f_n^T]^T W_0^F + b_0^F) W_1^F + b_1^F \quad (3)$$

where $W_0^F \in \mathbb{R}^{d_e \times d_m}$, $W_1^F \in \mathbb{R}^{d_m \times d_m}$, $b_0^F \in \mathbb{R}^{n \times d_m}$, and $b_1^F \in \mathbb{R}^{n \times d_m}$ are learnable parameters.

The following Transformer encoder contains multiple layers. Each layer contains a Multi-Head self-attention (MH) sub-layer, and a Position-wise Feed-Forward Network (PFFN) sub-layer, where a residual connection is employed around each of the two sub-layers, followed by Layer Normalization (LN). The definitions of MH and PFFN are identical to the paper [29]. We define the output of the first sub-layer in layer i as H_1^i . Thus each layer H^i can be calculated iteratively:

$$H^i = \text{Transformer}_e(H^{i-1}), \forall i = 1, \dots, l \quad (4)$$

$$\text{Transformer}_e(H^{i-1}) = \text{LN}(H_1^{i-1} + \text{PFFN}(H_1^{i-1})) \quad (5)$$

$$H_1^{i-1} = \text{LN}(H^{i-1} + \text{MH}(H^{i-1}, H^{i-1})) \quad (6)$$

After l layers, we obtain the output $G = H^l$. Let g_{mask} denote the corresponding output of the input [MASK], we then append a softmax layer on top of g_{mask} to calculate the probability of the masked item:

$$\Pr(f_{mask} | F_{mask}; \Theta_F) = \frac{\exp(g_{mask} h_{mask})}{\sum_{h \in \mathcal{H}} \exp(g_{mask} h)} \quad (7)$$

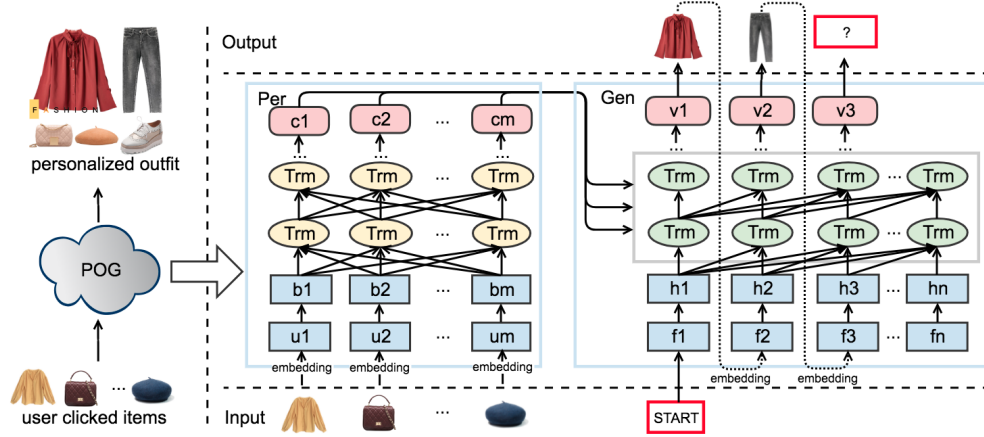


Figure 4: The architecture of POG, which is an encoder-decoder architecture with a Per network and a Gen network. The outfit item is generated step by step according to the user preference signal from the Per network and the compatibility signal from the Gen network.

where h_{mask} is the ground truth transition embedding of the masked item, and \mathcal{H} contains the transition embeddings from all the items in \mathcal{F} . One can choose \mathcal{H} to be the whole set of item transition embeddings of all the items, however, this is not practical due to the large number of high-dimensional embeddings. Therefore, we obtain h by randomly sampling 3 items from the whole transition set, which are not appeared in the outfit F , together with h_{mask} . This allows the model to learn the compatibility information from the outfit by looking at a diverse set of samples. Note that h_{mask} and h do not correspond to the original item embeddings, but rather the outputs of the transition layer.

4.3 POG: Personalized Outfit Generation Model

After modeling outfit compatibility, we now consider a generation model which generates personalized and compatible outfit by introducing user preference signals. Take the advantage of encoder-decoder structure, we aim to translate an user's historical behaviors to a personalized outfit. Let \mathcal{U} denote the set of all users and \mathcal{F} be the set of all outfits. We use a sequence of user behaviors $U = \{u_1, \dots, u_i, \dots, u_m\}$ to characterize an user, where u_i are the clicked items by the user. $F = \{f_1, \dots, f_t, \dots, f_n\}$ is the clicked outfit from the same user, where f_t are the items in the outfit. At each time step, we predict the next outfit item given previous outfit items and user's click sequence on items U . Thus for pair (U, F) , the objective function of POG can be written as:

$$\mathcal{L}_{(U,F)} = -\frac{1}{n} \sum_{t=1}^n \log Pr(f_{t+1} | f_1, \dots, f_t, U; \Theta_{(U,F)}) \quad (8)$$

where $\Theta_{(U,F)}$ denotes the model parameters. $Pr(\cdot)$ is the probability of seeing f_{t+1} conditioned on both previous outfit items and user clicked items.

The model architecture is shown in Figure 4. In POG, the encoder takes user clicked items as the input. Given a special token [START], the decoder then generates an outfit by one item at a time. At each step, the model is auto-regressive consuming the previously generated items as input. The generation stops when a special token [END] appears. In the end, an outfit is generated by composing

the output items. We name the encoder as Per network, and the decoder as Gen network. To be precise, the Per network provides a user preference signal, and the Gen network can generate outfits based on both personalization signal and compatibility signal. In the Per network, after a transition layer, a Transformer encoder architecture is followed with a stack of p identical layers. Thus the user preference can be obtained via:

$$B^0 = \text{ReLU}([u_1^T; \dots; u_m^T]^T W_0^U + b_0^U) W_1^U + b_1^U \quad (9)$$

$$B^i = \text{Transformer}_e(B^{i-1}), i = 1, \dots, p \quad (10)$$

where B^0 is the output of the transition layer. $W_0^U \in \mathbb{R}^{d_e \times d_m}$, $W_1^U \in \mathbb{R}^{d_m \times d_m}$, $b_0^U \in \mathbb{R}^{m \times d_m}$, and $b_1^U \in \mathbb{R}^{m \times d_m}$ are learnable parameters. The output of the Per network is $C = B^p$.

The Gen network is initialized using the aforementioned pre-trained FOM. In the Gen network, after the transition layer, a Transformer decoder architecture is followed. Each Transformer decoder layer contains three sub-layers. The first sub-layer is a Masked Multi-Head self-attention (MMH) mechanism. The masked design ensures that the prediction for a certain item can depend only on the known outputs of previous items. The second sub-layer performs Multi-Head attention (MH) over the output of user behavior C , which aims at introducing user information to the outfit generation process. The third sub-layer is a Position-wise Feed-Forward Network (PFFN). Similarly, we employ residual connections around each of the sub layers, followed by Layer Normalization (LN). Define the output of the first sub-layer in layer i as H_1^i , and the output of the second sub-layer as H_2^i , then each Transformer decoder layer H^i can be calculated iteratively:

$$H^0 = \text{ReLU}([f_1^T; \dots; f_n^T]^T W_0^F + b_0^F) W_1^F + b_1^F \quad (11)$$

$$H^i = \text{Transformer}_d(H^{i-1}), i = 1, \dots, q \quad (12)$$

$$\text{Transformer}_d(H^{i-1}) = \text{LN}(H_2^{i-1} + \text{PFFN}(H_2^{i-1})) \quad (13)$$

$$H_2^{i-1} = \text{LN}(H_1^{i-1} + \text{MH}(H_1^{i-1}, C, C)) \quad (14)$$

$$H_1^{i-1} = \text{LN}(H^{i-1} + \text{MMH}(H^{i-1}, H^{i-1}, H^{i-1})) \quad (15)$$

Algorithm 1 Personalized Outfit Generation.**Input:** Candidate set C ; User clicked item sequence U .**Output:** Personalized fashion outfit F .

- 1: Transform all items in C into outfit space \mathcal{H}_C .
- 2: Input U in the Per network.
- 3: Input [START] mark in the Gen network.
- 4: **repeat**
- 5: Get the output \mathbf{v}_t of the t -th step.
- 6: Output the next item f_{t+1} in \mathcal{H}_C according to Equation 17.
- 7: Input f_{t+1} to the Gen network.
- 8: **until** (The [END] mark appears.)

where H^0 is the output of the transition layer, and q is the number of layers of the Gen network. $W_0^F \in \mathbb{R}^{d_e \times d_m}$, $W_1^F \in \mathbb{R}^{d_m \times d_m}$, $b_0^F \in \mathbb{R}^{n \times d_m}$, and $b_1^F \in \mathbb{R}^{n \times d_m}$ are learnable parameters. The final output of POG is represented as $V = H^q$.

In one training sample with one clicked outfit and several clicked items from the same user, the ground truth prediction is the items in the outfit. We define \mathbf{h}_{t+1} as the ground truth transition vector of the t -th output \mathbf{v}_t . Thus the probability of the next item is calculated as:

$$Pr(f_{t+1}|f_1, \dots, f_t, U; \Theta_{(U,F)}) = \frac{\exp(\mathbf{v}_t \mathbf{h}_{t+1})}{\sum_{\mathbf{h} \in \mathcal{H}} \exp(\mathbf{v}_t \mathbf{h})} \quad (16)$$

Again, \mathcal{H} is composed of 3 randomly selected embeddings from the transition layer, together with the ground truth vector \mathbf{h}_{t+1} .

In the inference process, for each output vector \mathbf{v}_t , we deploy a similarity search to the candidate pool based on the following objective function:

$$\mathbf{h}_{t+1} = \arg \max_{\mathbf{h} \in \mathcal{H}_C} \frac{\exp(\mathbf{v}_t \mathbf{h}_C)}{\sum_{\mathbf{h} \in \mathcal{H}_C} \exp(\mathbf{v}_t \mathbf{h})} \quad (17)$$

where C is the candidate item set. Then the item with the transition embedding \mathbf{h}_{t+1} is chosen to be the next item. The generation stops when the [END] mark appears. A detailed generation process is described in Algorithm 1.

4.4 Dida Platform

In order to ensure an overall smooth user experience, strict quality and efficiency requirements have to be taken into consideration when deploying POG into an online application with millions of users. We thus develop a platform named Dida which is able to generate personalized outfits automatically at very large scales and ensures necessary quality and efficiency standards. Dida platform consists of a number of services including item selection, outfit generation, image composition, and personalized recommendation, as shown in Figure 5. The outfit generation and recommendation workflow, wherein POG is implemented and deployed, can be described as follows:

- We support million scale item pools and assist operators to select qualified images with clean background and chosen categories. Images, texts, and CF-data are extracted for selected items and multi-modal embeddings are computed as described in Section 4.1. In the generation process as described in Algorithm 1, we use Faiss [7] to implement the similarity search. To ensure the quality of the generated outfits, we restrict the similarity search according

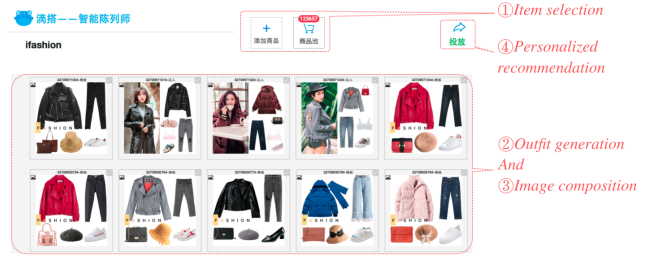


Figure 5: Dida platform. We provide services including item selection, outfit generation, image composition, and personalized recommendation.

to certain domain rules provided by fashion experts. For example, if it specifies that one t-shirt, one pair of jeans, and one sport shoes as a category rule, then only jeans and sport shoes can be searched in Faiss after the t-shirt.

- The generated items from the first step will be composed together with their images in designed templates. After image composition, we recommend outfits to users. The personalized generated outfits from POG are recommended directly to the users. In addition, we also support other recommendation strategies including random recommendation and CF recommendation.

Dida is widely used by more than one million operators at Alibaba. About 6 million personalized outfits are generated everyday with high quality. So far, the outfits have been recommended to more than 5.4 million users.

5 EXPERIMENT

We describe our experiments and analyze the results in this section. We conduct both offline and online experiments to compare the compatibility and recommendation performances.

5.1 Fashion Outfit Compatibility

5.1.1 Implementation Details. For multi-modal embedding process, we use 1536 dimensional CNN features derived from Inception Resnet V2 model [26] as the image features. The text representations are 300 dimensional vectors derived from TextCNN [10] with a vocabulary size of 420,758. Graph embeddings are 160 dimensional CF vectors from [31]. We set $d_e = 128$ as the final dimension of the multi-modal embedding. Triplet margin is fixed as $\alpha = 0.1$. In FOM, the model is composed of a stack of $l = 6$ layers. We use 8 heads in all multi-heads attentions, and $d_m = 64$ as the dimension of hidden layers.

5.1.2 Task Settings & Evaluation Metrics. To evaluate the performances of multi-modal embeddings and models on predicting the outfit compatibility, we adopted two wide-used tasks following the practices in [3, 11, 18, 28].

Fill In the Blank (FITB). FITB is a task predicting an item from multiple choices that is compatible with other items to fill in the blank. For offline evaluation, we split 10% of the data as a test set. We mask one item at a time for each outfit in the test set. Then for each masked item, we randomly select 3 items from other outfits along with the ground truth item to obtain a multiple choice set. It is reasonable to believe that the ground truth item is more compatible

Table 3: FITB and CP results of different modalities.

| Modality | FITB | CP |
|---------------|---------------|---------------|
| Image | 62.85% | 80.78% |
| Text | 68.02% | 85.17% |
| CF | 51.38% | 57.92% |
| Image+text | 68.24% | 85.41% |
| Image+text+CF | 68.71% | 86.09% |

with other items than the randomly selected ones. Finally, the result is evaluated by the accuracy of choosing the correct items.

Compatibility Prediction (CP). The CP task is to predict whether a candidate outfit is compatible or not. For evaluation, we first build a compatible outfit set which is 10% split from the dataset. Then, we produce the same amount of incompatible outfits as the compatible outfits for a balanced test set by randomly selecting fashion items from the compatible outfit set. In the CP task, a candidate outfit is scored as to whether its constitute items are compatible with each other. The performance is evaluated using the AUC (Area Under Curve) of the ROC (Receiver Operating Characteristic) curve.

Specifically, in FOM, we solve the FITB task based on Equation 17, where C is the choice set. For CP task, given an outfit, we simply utilize the negative value of Equation 2 as the compatibility score.

5.1.3 Compare Different Modalities. We compare each modality and their combinations on compatibility evaluations. The performances are compared in Table 3, which shows that: (1) In both FITB and CP tasks, text works best on its own, but image and CF provide complementary information. (2) CF embedding alone does not work very well, partly because it lacks semantic visual and textual information, which is important in fashion compatibility. (3) The 1536 dimensional CNN feature derived from Inception Resnet V2 is compressed to 128 after the fully connected layer. The resulting dimension is relatively small to contain the important visual information of the fashion items, which partly explains that the image only modality performs badly. Similar results can also be observed from [11].

5.1.4 Compare Different Models. To demonstrate the effectiveness of our model FOM, we compared several alternative models. For fair comparison, all the models share the same embeddings and the dimension of hidden layers.

- **F-LSTM**[3] Given a sequence of fashion items in an outfit, a forward LSTM is trained taking each item as an individual time step.

- **Bi-LSTM**[3] A bidirectional LSTM adds a backward LSTM compared to F-LSTM, thus predicting the next item can be performed in the reverse order also.

- **SetNN**[11] SetNN uses a multi-instance pooling model to aggregate information from individual fashion items to produce a holistic embedding for the outfit. We use mean reduction as the pooling method, which performs the best in experiments as [11] shows. The original SetNN predicts popularity of an outfit by labeling the outfit preferences. Since popularity does not always indicate compatibility, for fair comparison, we change its training inputs with compatible outfits, and incompatible outfits by randomly selecting items from the training set.

Table 4: FITB and CP results of different models.

| Model | FITB | | CP | |
|------------|---------------|---------------|---------------|---------------|
| | Unordered | Ordered | Unordered | Ordered |
| F-LSTM[3] | 58.07% | 62.84% | 63.78% | 65.04% |
| Bi-LSTM[3] | 58.21% | 64.91% | 63.82% | 68.61% |
| SetNN[11] | 49.24% | 49.27% | 58.31% | 58.33% |
| FOM (ours) | 68.71% | 68.79% | 86.09% | 86.32% |

- **FOM (ours)** Our Fashion Outfit Model by learning a masked item prediction task described in Section 4.2.

In particular, F-LSTM and Bi-LSTM are sequence based models, which take sequences of items as inputs. SetNN and FOM are set based models, which do not enforce a specific order over the fashion items as inputs. Considering the particular requirement of sequence based models which need ordered inputs, we conduct the experiments with both disordered inputs and ordered inputs. The order follows the guide in [3], which has a fixed order: tops, bottoms, shoes, and accessories, where we use leaf categories to specify.

FITB Results. The middle two columns of Table 4 shows the results of our model compared with alternative models in the FITB task. From this table, we make the following observations: (1) Sequence based models are sensitive to the input order, while set based models are not. Both F-LSTM and Bi-LSTM have better performances with ordered inputs. SetNN and FOM have similar results with different inputs. This demonstrates that sequence based models can only work well in specific order. (2) Bi-LSTM has better results than F-LSTM both in ordered inputs and unordered inputs. The combination of LSTMs in two directions offers higher accuracy than one directional LSTM. Similar result is also observed in [3]. (3) FOM performs the best for both inputs. It exhibits a 18.04% increase compared to the second best result (Bi-LSTM) in unordered inputs, and a 5.98% increase compared to the second best result (Bi-LSTM) in ordered inputs. We attribute this to the self-attention mechanism in Transformer, which facilitates FOM to calculate the weighted interactions of every item to the others, makes it suitable for modeling compatibility in this task.

CP Results. The last two columns in Table 4 shows the results of our model and others for the CP task. Similar to FITB, sequence based models are still sensitive to the input order. FOM obtains the best performance by significant margins. In unordered inputs, FOM has a 34.90% increase compared to the second best (Bi-LSTM), while in ordered inputs, FOM shows 25.81% increase compared to the second best (Bi-LSTM). Although SetNN is directly trained to predict set compatibility, it still performs the worst. Similar results can also be found in [11]. We attribute the performance of FOM to the multi-head self-attention mechanism. It learns the compatibilities between each item and all the other items within the outfit, and then an outfit compatibility can be learned by combining all these item compatibilities.

5.2 Fashion Outfit Generation and Recommendation

In this section, we conduct extensive online experiments on different fashion models and recommendation methods. The performance of outfit compatibility and recommendation can be evaluated via CTR, which is an explicit metric representing the effectiveness of different models and methods. Further, we show some online example cases based on our model.

5.2.1 Implementation Details & Task Settings. In POG, the Per network is composed of a stack of $p = 6$ layers, and the Gen network is composed of a stack of $q = 6$ layers. We use 8 heads in all multi-head attentions, and $d_m = 64$ as the dimension of hidden layers. Our online test data is a manually selected item pool with 1.57 million items. Outfits are generated and recommended via different generation models and recommendation methods on Dida platform. We split the users equally into 8 buckets, each tests a certain generation model with recommendation method.

5.2.2 Compare Different Models and Methods. Since SetNN [11] cannot generate outfits, we focus on the comparisons among generation available models. It is worth mentioned that none of the existing generation models, including F-LSTM and Bi-LSTM, are available to do personalized recommendation. So we Randomly Recommend (RR) these generated outfits to the users. To compare the recommendation performance better, we also deploy classic CF [22] methods to recommend. We first estimate a user's most preferred item via measuring its similarities with the items in his or her interaction history by calculating an item-to-item similarity matrix, according to [22]. Then, the model generated outfit which contains the user's most preferred item is recommended. Online experiments are listed as follows:

- **F-LSTM[3]+RR** Outfits are generated by starting with every item in the test set using the F-LSTM model. The generated outfits are recommended randomly to users.
- **Bi-LSTM[3]+RR** Outfits are generated as paper [3] presents by taking every item in the test set as the query item. The generated outfits are recommended randomly to users.
- **Gen+RR (ours)** Only Gen network is used for generation. We input no user information to POG to evaluate the generation quality only. The generated outfits are recommended randomly.
- **F-LSTM[3]+CF** Outfits are generated by F-LSTM model. The generated outfits are recommended using the CF method.
- **Bi-LSTM[3]+CF** Outfits are generated by Bi-LSTM model. The generated outfits are recommended using the CF method.
- **Gen+CF (ours)** Only Gen network is used for generation. The generated outfits are recommended using the CF method.
- **POG (ours)** Both the Per network and the Gen network are used in POG to generate and recommend personalized outfits.
- **POG+FOM (ours)** The Gen network is initialized using the pre-trained FOM before generating personalized outfits.

We record the CTR on outfits over a period of seven days, and report the results in Figure 6. Three conclusions can be observed from the figure: (1) Although the CTR values vary during the seven days, we still observe an obvious phenomenon that the recommendation performance largely relies on the recommendation methods. CF methods have higher CTR than all the RR methods, and POG

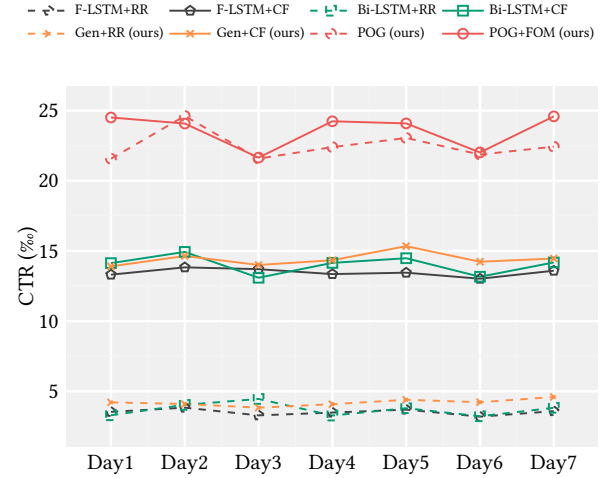


Figure 6: CTR of online experiments. Methods based on POG achieve the best performances. CF based methods are better than RR methods.

methods outperform all the other. POG with pre-trained FOM has the highest CTR in most of the time. It has an improvement of more than 70% in CTR, compared to the state-of-the-art model Bi-LSTM with CF recommendation. We attribute this to the differences between individual item recommendation and outfit recommendation. The most preferred item of one user may not indicate the most preferred outfit. (2) The compatibility performance of these models matters, but not in an obvious way. This is partly because of the rule based search in the generation process described in Section 4.4, which avoids some extreme bad cases. In most of the time, Gen has the highest CTR both in RR and CF methods. (3) The performance of POG with pre-trained FOM has a little improvement over POG alone. We may conclude that the Gen network in POG is already enough to learn the compatibility of outfits.

5.2.3 Case Study. We randomly sample some online cases, and present them in Figure 7. For users with various preferences on the items, (e.g. the first two rows), POG can generate compatible outfits conditioned on the common properties of these items. For users who clicked a lot of similar items, (e.g. the last two rows), POG provides an outfit with the main item, which has similar properties with the clicked items. The outfits may inspire the users with better ideas of how to compose the items more properly, and even influence them to buy the whole outfit.

6 CONCLUSION AND FUTURE WORK

In this paper, we propose a personalized outfit generation model to build the bridge between compatibility and personalization, which are two essential requirements in fashion generation and recommendation. Our model is built in three steps: multi-modal embedding, fashion outfit modeling, and personalized fashion outfit modeling. We use Transformer in our model which helps to find interactions of the items in the outfit, and also helps to connect user behaviors on items and outfits. Our model outperforms other alternative models in outfit compatibility, outfit generation and recommendation by significant margins. Furthermore, we deploy



Figure 7: The online cases of POG. Four rows correspond to four users. The first five columns are user clicked items sampled from the latest 50 clicks, and the last column is the generated outfit by POG.

POG on the Dida platform to assist personalized outfit generation in large-scale online application, which is widely used in Alibaba now. We share our data in this paper for future research.

For the future work, we aim to solve the cold-start problem via building user profiles from similar user groups.

ACKNOWLEDGEMENT

We would like to thank colleagues of our team - Yuchi Xu, Zhiyuan Liu, Wei Li, Jizhe Wang, Mengmeng Wu, Lifeng Wang, and Qiwei Chen for useful discussion and supports. We are grateful to our cooperative team - Wentao Yang, Jun Bao, Yue Hu, Yumeng Liu, Lijie Cao, Yunfeng Zhang, Qichuan Xiao, and Mingkai Wang. We also thank the reviewers for their valuable comments and suggestions that help improve the quality of this manuscript.

REFERENCES

- [1] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [2] M Hadi Kiapour, Xufeng Han, Svetlana Lazebnik, Alexander C Berg, and Tamara L Berg. 2015. Where to buy it: Matching street clothing photos in online shops. In *Proceedings of the IEEE international conference on computer vision*. 3343–3351.
- [3] Xintong Han, Zuxuan Wu, Yu-Gang Jiang, and Larry S Davis. 2017. Learning fashion compatibility with bidirectional lstms. In *Proceedings of the 2017 ACM on Multimedia Conference*. ACM, 1078–1086.
- [4] Yang Hu, Xi Yi, and Larry S Davis. 2015. Collaborative fashion recommendation: A functional tensor factorization approach. In *Proceedings of the 23rd ACM international conference on Multimedia*. ACM, 129–138.
- [5] Junshi Huang, Rogerio S Feris, Qiang Chen, and Shuicheng Yan. 2015. Cross-domain image retrieval with a dual attribute-aware ranking network. In *Proceedings of the IEEE international conference on computer vision*. 1062–1070.
- [6] Vignesh Jagadeesh, Robinson Piramuthu, Anurag Bhardwaj, Wei Di, and Neel Sundaresan. 2014. Large scale visual recommendations from street fashion images. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1925–1934.
- [7] Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2017. Billion-scale similarity search with GPUs. *arXiv preprint arXiv:1702.08734* (2017).
- [8] Wang-Cheng Kang, Chen Fang, Zhaowen Wang, and Julian McAuley. 2017. Visually-aware fashion recommendation and design with generative image models. In *Data Mining (ICDM), 2017 IEEE International Conference on*. IEEE, 207–216.
- [9] Wang-Cheng Kang and Julian McAuley. [n. d.]. Self-Attentive Sequential Recommendation. In *Proceedings of ICDM*. 197–206.
- [10] Yoon Kim. 2014. Convolutional neural networks for sentence classification. *arXiv preprint arXiv:1408.5882* (2014).
- [11] Yuncheng Li, Liangliang Cao, Jiang Zhu, and Jiebo Luo. 2017. Mining fashion outfit composition using an end-to-end deep learning approach on set data. *IEEE Transactions on Multimedia* 19, 8 (2017), 1946–1955.
- [12] Xiaodan Liang, Liang Lin, Wei Yang, Ping Luo, Junshi Huang, and Shuicheng Yan. 2016. Clothes co-parsing via joint image segmentation and labeling with application to clothing retrieval. *IEEE Transactions on Multimedia* 18, 6 (2016), 1175–1186.
- [13] Zhouhan Lin, Minwei Feng, Cicero Nogueira dos Santos, Mo Yu, Bing Xiang, Bowen Zhou, and Yoshua Bengio. 2017. A structured self-attentive sentence embedding. *arXiv preprint arXiv:1703.03130* (2017).
- [14] Greg Linden, Brent Smith, and Jeremy York. 2003. Amazon. com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing* 1 (2003), 76–80.
- [15] Si Liu, Jiashi Feng, Zheng Song, Tianzhu Zhang, Hanqing Lu, Changsheng Xu, and Shuicheng Yan. 2012. Hi, magic closet, tell me what to wear!. In *Proceedings of the 20th ACM international conference on Multimedia*. ACM, 619–628.
- [16] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. 2016. Deep-fashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1096–1104.
- [17] Julian McAuley, Christopher Targett, Qinfeng Shi, and Anton Van Den Hengel. 2015. Image-based recommendations on styles and substitutes. In *Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, 43–52.
- [18] Takuma Nakamura and Ryosuke Goto. 2018. Outfit Generation and Style Extraction via Bidirectional LSTM and Autoencoder. *arXiv preprint arXiv:1807.03133* (2018).
- [19] Ankur P Parikh, Oscar Täckström, Dipanjan Das, and Jakob Uszkoreit. 2016. A decomposable attention model for natural language inference. *arXiv preprint arXiv:1606.01933* (2016).
- [20] Matthew E Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint arXiv:1802.05365* (2018).
- [21] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. [URL https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language_understanding_paper.pdf](https://s3-us-west-2.amazonaws.com/openai-assets/research-covers/language-unsupervised/language_understanding_paper.pdf) (2018).
- [22] Badrul Sarwar, George Karypis, Joseph Konstan, and John Riedl. 2001. Item-based collaborative filtering recommendation algorithms. In *Proceedings of the 10th international conference on World Wide Web*. ACM, 285–295.
- [23] Xueming Song, Fuli Feng, Xianjing Han, Xin Yang, Wei Liu, and Liqiang Nie. 2018. Neural Compatibility Modeling with Attentive Knowledge Distillation. *arXiv preprint arXiv:1805.00313* (2018).
- [24] Xueming Song, Fuli Feng, Jinhuan Liu, Zekun Li, Liqiang Nie, and Jun Ma. 2017. Neurostylist: Neural compatibility modeling for clothing matching. In *Proceedings of the 2017 ACM on Multimedia Conference*. ACM, 753–761.
- [25] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential Recommendation with Bidirectional Encoder Representations from Transformer. *CoRR abs/1904.06690* (2019). arXiv:1904.06690
- [26] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. 2017. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI*, Vol. 4. 12.
- [27] Ivona Tautkute, Tomasz Trzcinski, Aleksander Skorupa, Lukasz Brocki, and Krzysztof Marasek. 2018. DeepStyle: Multimodal Search Engine for Fashion and Interior Design. *arXiv preprint arXiv:1801.03002* (2018).
- [28] Mariya I Vasileva, Bryan A Plummer, Krishna Dusat, Shreya Rajpal, Ranjitha Kumar, and David Forsyth. 2018. Learning Type-Aware Embeddings for Fashion Compatibility. *arXiv preprint arXiv:1803.09196* (2018).
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*. 5998–6008.
- [30] Andreas Veit, Balazs Kovacs, Sean Bell, Julian McAuley, Kavita Bala, and Serge Belongie. 2015. Learning visual clothing style with heterogeneous dyadic co-occurrences. In *Proceedings of the IEEE International Conference on Computer Vision*. 4642–4650.
- [31] Jizhe Wang, Pipei Huang, Huan Zhao, Zhibo Zhang, Binqiang Zhao, and Dik Lun Lee. 2018. Billion-scale commodity embedding for e-commerce recommendation in alibaba. In *SIGKDD*. 839–848.
- [32] Kota Yamaguchi, M Hadi Kiapour, and Tamara L Berg. 2013. Paper doll parsing: Retrieving similar styles to parse clothing items. In *Proceedings of the IEEE international conference on computer vision*. 3519–3526.
- [33] Kota Yamaguchi, M Hadi Kiapour, Luis E Ortiz, and Tamara L Berg. 2012. Parsing clothing in fashion photographs. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*. IEEE, 3570–3577.