

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

6-2020

Knowledge enhanced neural fashion trend forecasting

Yunshan MA

Yujuan DING

Xun YANG

Lizi LIAO

Singapore Management University, lzliao@smu.edu.sg

Wai Keung WONG

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Artificial Intelligence and Robotics Commons](#), and the [Databases and Information Systems Commons](#)

Citation

MA, Yunshan; DING, Yujuan; YANG, Xun; LIAO, Lizi; WONG, Wai Keung; and CHUA, Tat-Seng. Knowledge enhanced neural fashion trend forecasting. (2020). *ICMR '20: Proceedings of the 2020 International Conference on Multimedia Retrieval*. 82-90.

Available at: https://ink.library.smu.edu.sg/sis_research/7674

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylids@smu.edu.sg.

Author

Yunshan MA, Yujuan DING, Xun YANG, Lizi LIAO, Wai Keung WONG, and Tat-Seng CHUA

Knowledge Enhanced Neural Fashion Trend Forecasting

Yunshan Ma*
 National University of Singapore
 yunshan.ma@u.nus.edu

Yujuan Ding*
 The Hong Kong Polytechnic University
 dingyujuan385@gmail.com

Xun Yang†
 National University of Singapore
 xunyang@nus.edu.sg

Lizi Liao
 National University of Singapore
 liaolizi.llz@gmail.com

Wai Keung Wong
 The Hong Kong Polytechnic University
 calvin.wong@polyu.edu.hk

Tat-Seng Chua
 National University of Singapore
 dcscs@nus.edu.sg

ABSTRACT

Fashion trend forecasting is a crucial task for both academia and industry. Although some efforts have been devoted to tackling this challenging task, they only studied limited fashion elements with highly seasonal or simple patterns, which could hardly reveal the real fashion trends. Towards insightful fashion trend forecasting, this work focuses on investigating fine-grained fashion element trends for specific user groups. We first contribute a large-scale fashion trend dataset (FIT) collected from Instagram with extracted time series fashion element records and user information. Furthermore, to effectively model the time series data of fashion elements with rather complex patterns, we propose a Knowledge Enhanced Recurrent Network model (KERN) which takes advantage of the capability of deep recurrent neural networks in modeling time-series data. Moreover, it leverages internal and external knowledge in fashion domain that affects the time-series patterns of fashion element trends. Such incorporation of domain knowledge further enhances the deep learning model in capturing the patterns of specific fashion elements and predicting the future trends. Extensive experiments demonstrate that the proposed KERN model can effectively capture the complicated patterns of objective fashion elements, therefore making preferable fashion trend forecast.

CCS CONCEPTS

• Information systems → Specialized information retrieval.

KEYWORDS

Fashion Trend Forecasting; Fashion Analysis; Time Series Forecasting

*Equal contribution.

†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 ACM 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.
 ICMR '20, June 8–11, 2020, Dublin, Ireland

© 2020 Association for Computing Machinery.
 ACM ISBN 978-1-4503-7087-5/20/06...\$15.00
<https://doi.org/10.1145/3372278.3390677>

Fashion trends evolving in social media (Example group: New York Female)

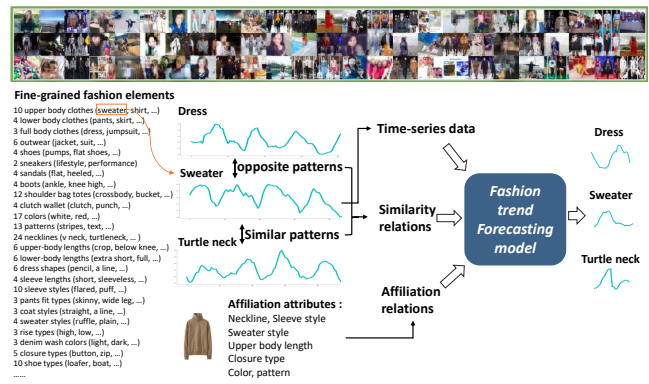


Figure 1: The fashion trend forecasting task aims to predict the future trends of meaningful fashion elements.

ACM Reference Format:

Yunshan Ma, Yujuan Ding, Xun Yang, Lizi Liao, Wai Keung Wong, and Tat-Seng Chua. 2020. Knowledge Enhanced Neural Fashion Trend Forecasting. In *Proceedings of the 2020 International Conference on Multimedia Retrieval (ICMR '20)*, June 8–11, 2020, Dublin, Ireland. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3372278.3390677>

1 INTRODUCTION

Karl Lagerfeld¹ used to say that the essence of fashion is changeability. Fashion trend forecasting, aiming to master such change, is therefore of great significance in fashion industry. It enables fashion companies to develop products and establish marketing strategies more wisely. It also helps fashion consumers make better choices. Traditionally, to predict fashion trends, the staffs of forecasting companies travel across the world to observe the art, music, and other cultural factors that may influence fashion industry. Also, the staffs collect information of consumers' ways of living, thinking, and behaving [18]. However, the existing solutions mainly rely on subjective inferences of these forecasters, which may be less reliable and have large variations.

In the recent decade, technological innovations such as Internet has accelerated the rate of fashion change, which makes fashion trend forecasting even more difficult. On the other hand, the advent of digital age has facilitated the accumulation of huge amounts

¹https://wikipedia.org/wiki/Karl_Lagerfeld

of fashion-related data, which provides an alternative data-driven way of addressing the fashion trend forecasting task [1]. This paper aims to mine useful fashion information from big historical data and predict the possible development of fashion for the future [1, 11, 27]. There are two main research challenges for this task: 1) What kind of data should be used and analyzed in order to make meaningful and relevant fashion trend forecasting? 2) How to effectively model relevant data to make accurate predictions?

For the first challenge, the source data should contain abundant time series fashion information, and should also be of considerable scale to cover a rather long time period in order to reflect the evolution of fashion over time. Compared to e-commerce or fashion show [1, 38], social media is a more appropriate data source because it sensitively and extensively records the fashion development with massive uploaded fashion-related images and comments everyday from multiple sources of end users, fashion bloggers and brands, *etc.* Besides, rich information for both users and fashion items can be extracted from the images, meta data and other source data by the well-developed computer vision or other machine learning techniques. Although there exists datasets based on social media [27], they contain very limited fashion elements and are far from enough for forecasting meaningful and applicable fashion trends. Also, the information of users (such as age, gender or living location) that actually convey most fashion-related data is essential in fashion trend observation. Such user information, however, is neglected in existing datasets. Considering the limitations of existing datasets, in this paper, we build a new dataset with extensive fine-grained fashion elements, including category, attribute and style. It also covers a longer time period with richer user information. More details are introduced in Section 3.

For the second challenge, in order to make accurate data-driven fashion trend forecasting, the underlying patterns in the time series data need to be effectively captured. Though traditional models such as statistical models or matrix factorization have been effectively applied to model simple time series data [27, 28], they fall short of ability to make sound predictions for more complicated fashion trends. Recent advances of deep learning have provided great solutions for many tasks [20]. In particular, the recurrent neural networks (RNN) have demonstrated its superiority in modeling time series data and addressing relevant problems [4, 6, 42]. However, such approaches have not been employed in the area of fashion trend analysis yet. On the other hand, most existing works predict the trend of each fashion element independently. However, according to common sense, fashion elements are not independent but well-correlated with each other in various ways. For example as shown in Figure 1, the trend of *sweater* shows similar pattern with that of *turtle neck*, but nearly opposite with that of *dress*. If we try to predict the trend of *sweater*, we can apply the prediction results on both the *turtle neck* and *dress* to refine the prediction of *sweater* based on their observed correlations. Furthermore, in fashion domain, there naturally exist taxonomic relations between elements, *e.g.*, the affiliation relations between *sweater* and all its affiliation attributes as shown in Figure 1. Such taxonomic relations would result in relations among fashion trend patterns, which we should take advantage of in fashion trend modeling. In short, these types of prior domain knowledge describing the relations among fashion trends are non-trivial to model but helpful.

Driven by the above motivations, this paper presents a novel approach named Knowledge Enhanced Recurrent Network (KERN) for forecasting fashion trends of people in various groups. The proposed approach effectively models the time series data of fashion elements with rather complex patterns by using the Long-Short Term Memory (LSTM) encoder-decoder framework. More importantly, it incorporates two types of knowledge: internal and external knowledge. Specifically, for internal knowledge, it leverages the similarity relations of time series within dataset and introduces a triplet regularization loss based on pattern similarities. For external knowledge, it takes advantage of the affiliation relations of fashion elements within the taxonomy, and incorporates them by updating the embedding of fashion elements via message passing. The proposed KERN model incorporates both the time series information of single fashion element and the connectivity between this element and all related ones. We also exploit the user information for better modeling the different fashion trends for different groups of users by applying the semantic group representation.

In summary, the contributions of this work are three folds: 1) towards meaningful fashion trend forecasting, we contribute a large-scale fashion trend dataset based on Instagram, termed Fashion Instagram Trending (FIT); 2) to make sound trend forecasting, we propose a novel knowledge enhanced LSTM-based model (KERN) to effectively model the time series fashion trend data; and 3) we conduct extensive experiments and analysis of fashion trend forecasting on our FIT dataset and the GeoStyle dataset [27], and show that our KERN model is capable of capturing patterns in time series fashion trends data and effectively forecasting fashion trends.

2 RELATED WORK

2.1 Fashion Trend Analysis

Recently, various tasks in fashion domain has attracted research interests such as fashion recognition [24, 26, 40], retrieval [5, 16, 21, 23], mix-and-match [9, 34, 35, 44?, 45], and visual try-on [10, 14], *etc.* Meanwhile, fashion trend analysis and forecasting has always been a classic research topic in fashion domain due to its significance in guiding the whole fashion industry. Hidayati *et al.* [11] analyzed fashion trends by proposing a framework to automatically detect fashion patterns which frequently occur within a fashion week. Vittayakorn *et al.* [38] extended such task to a larger dataset and studied both runway and real world fashion to produce quantitative analysis for fashion and trends. The main purpose of the above works is to analyze the current fashion styles, but not to make predictions. Likewise, Matzen *et al.* [28] studied clothing trends by statistical analysis, but based on a large-scale social media image dataset. Al-Halah *et al.* [1] proposed a fashion trends prediction model that uses nonnegative matrix factorization (NMF) to discover fashion styles and an exponential smoothing model to forecast the future of a style. However, the fashion styles they discovered are not real fashion style, but the cluster of certain kind of clothes with similar visual appearance. Comparatively, Mall *et al.* [27] explored very specific fashion elements and tried to find detailed fashion trends. They modeled the fashion trend signals of each target element with a basic combination of linear and cyclical components, which, according to the paper, were capable of capturing both coarse-level trends and fine-scale spikes. However,

the limitation is that they only targeted limited fashion elements which showed simple patterns in their trend signals (such as with hat or not) and did not really include fine-grained fashion elements. In summary, existing works are still limited to statistical analysis or predicting trends of specific fashion styles or elements with simple patterns. More meaningful fashion elements involving complex trends are still yet to be explored.

2.2 Time Series Forecasting

Fashion trend forecasting is also closely related to the time series forecasting problem which aims to predict the future based on the historical observations. Statistic models are classic solutions for time series forecasting problems, including the most representative autoregressive (AR) [39], moving averages (MA) [33], improved autoregressive integrated moving average (ARIMA) [3], and others [13, 43]. These models were found to be quite effective for forecasting structural data with high seasonality or simple trend. However, the real-life times-series signals are usually highly volatile and very difficult to model by these traditional methods. Recently, with the success of deep neural networks in a wide range of tasks, RNN, especially its variant LSTM [12], has shown its superiority in modeling sequential data and achieved superior performance in various applications of NLP [29], speech recognition [8], and also time series forecasting [4, 19, 36].

Since fashion trend forecasting is a rather domain-specific task, leveraging abundant fashion knowledge in the forecasting task is a viable approach. Actually, exploiting domain knowledge, or external knowledge, to enhance the performance of deep learning models has achieved promising results lately in many tasks [15, 17, 22, 37, 41]. Specifically, in time series forecasting problems like the stock price prediction, Feng *et al.* found that incorporating domain knowledge of stocks (*e.g.*, companies within the same industry sector) can effectively help stock price forecasting [7]. Despite of many successful applications, domain knowledge has not been well exploited in fashion trend forecasting.

3 PROBLEM FORMULATION AND DATASET

This paper focuses on the fashion trend forecasting problem, which aims to make prediction of future popularity with regard to each fashion element (*e.g.*, *white, dress, off-shoulder, etc.*) for each user group (*e.g.*, *London female of age between 18 and 25*). Given a fashion element $f \in \mathcal{F}$ and a user group $g \in \mathcal{G}$, the temporal popularity of f for g is defined as a time series denoted as $\mathbf{y}_g^f = (y_1, \dots, y_t, \dots)$, where \mathcal{F} is the set of all fashion elements; and \mathcal{G} is the set of all user groups. The value of the time series at each time step t is defined as $y_t = N_t^{g,f} / N_t^g$, where $N_t^{g,f}$ is the number of the fashion elements f at time point t for group g ; N_t^g is the number of all fashion items (*e.g.*, *clothing, bags, shoes, and etc.*) observed at time point t for group g . Given the historical inputs within the time span of $[1, T]$, our aim is to forecast the future values of time $[T+1, T+T']$, where T is the historical sequence length or time span, and T' is termed as the forecast horizon (the number of steps ahead to forecast, $T' > 1$).

Since none of the existing datasets supports this type of study, we contribute a new dataset based on the popular social media




Figure 2: (a) (b) Statistics of the FIT dataset. (c) Two examples of the FIT dataset, where RED curves are from the FIT dataset and BLUE curves are from Google Trends (both examples belong to the group [New York, Female]).

platform Instagram², termed as Fashion Instagram Trends (FIT). Figure 2 (a) and (b) show the statistical comparison between FIT and the Geostyle [27]. It shows that the FIT dataset has more user information, richer fashion elements, and longer time span.

Specifically, we crawl millions of posts uploaded by users from all over the world. To ensure quality of the crawled data, automated and manual filtering are conducted on the collected data, similar to that done in [25, 26]. First, we leverage the pre-trained object detection model to detect person body [31] and face [46]. Images without face or body, or with abnormal-sized face or body are filtered out. Then, we drop posts with people that are not the corresponding account owner. Finally, we keep about 680K images in total. The annotation of the dataset is from two aspects: users and fashion elements, which will be introduced in detail below.

For users, we collect three types of user information (*i.e.*, age, gender and location), and then separate users into different groups based on the information. For each user we first apply the off-the-shelf age and gender detector tools [2, 30, 32] on all of the users' posts (images), and then choose the dominant gender and the average age as the final gender and age. Posts detected as the opposite gender and with age differing from the detected age by over five years are dropped. We categorize the age of each user into four groups, that is: 0 to 18, 18 to 25, 25 to 40, and above 40. Next, we obtain the location based on the longitude and latitude data that comes with the post, and choose the most frequent one as the location of the user. Finally we keep a location set with 14 main cities across the world. The combination of the three types of user attributes forms a group, resulting in 74 groups in the whole dataset.

For fashion elements, we apply a commercial fashion tagging tool³ and extract three types of fashion elements (category, attribute and style) from the images, resulting in a total 197 of different fashion elements for the whole dataset. Each image is labelled with user group, time, and fashion elements after the annotation.

²instagram.com

³visenze.com

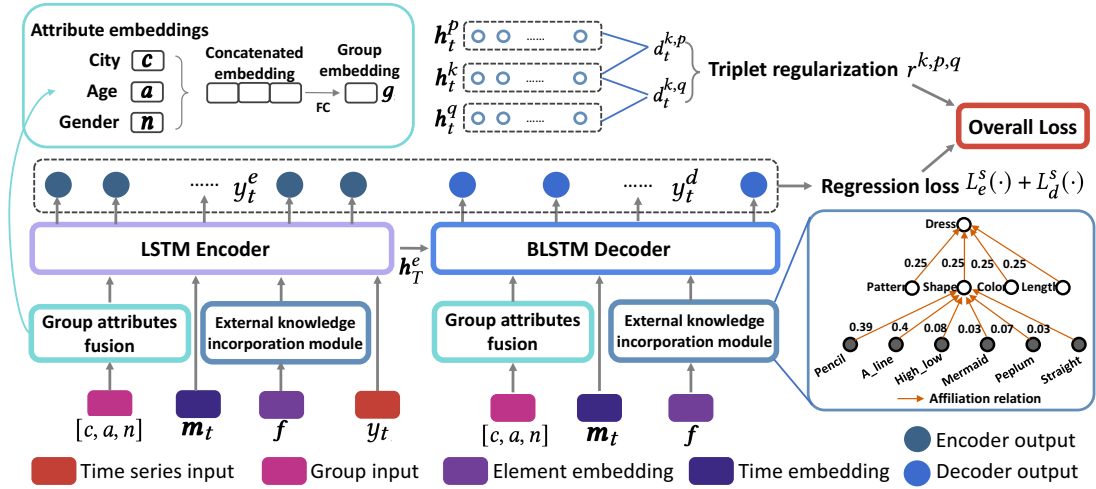


Figure 3: The Knowledge Enhanced Recurrent Network (KERN) framework. We utilize a basic LSTM encoder decoder framework to do multi-horizon trend forecasting and incorporate both internal and external knowledge via a triplet regularization term and a message passing module respectively.

We then calculate the popularity of each fashion element for each user group for every half month, resulting in a time series data. The post time of FIT dataset ranges from July 2014 to June 2019, spanning five years, which means that each time series has 120 data points. We further drop sparse time series with over 50% of time points with no data. Finally, we obtain around 8000 time series in total. Note that as our tags come from an existing tagging tool, which might contain some noise and result in a small bias of real fashion trends. However, we manually check part of the recognition results, and find that the average accuracy is relatively satisfying. More importantly, each time series data is a statistical ensemble of a group of users’ data, thus the noise of each user on the final time series is hugely weakened. Besides, we comprehensively analyze fashion trends in FIT and compare them with that from **Google Trends**⁴, and observe highly similar patterns, which further validate the credibility of our FIT dataset (see examples in Figure 2).

4 APPROACH

This paper aims to develop an end-to-end model to forecast the fashion trends given the historical inputs. First, we adopt the basic LSTM encoder decoder framework, which is able to incorporate both time series inputs and the associated sequence information into a unified model and make multi-horizon forecasting. Second, we add a triplet regularization term to explicitly incorporate the internal knowledge. Moreover, we introduce a message passing module to leverage the external knowledge extracted from the taxonomy. We name our proposed framework as **Knowledge Enhanced Recurrent Network (KERN)**, as shown in Figure 3.

4.1 Basic LSTM Encoder Decoder Framework

Given a time series (y_1, \dots, y_T) indicating the past trend of fashion element f for group g within time period $[1, T]$, we aim to forecast the future values of the trend $(y_{T+1}, \dots, y_{T+T'})$. The group g is defined by the combination of three attributes: the city c , the

⁴trends.google.com

age a and the gender n , where $c \in C$, $a \in \mathcal{A}$ (C, \mathcal{A} denote all cities and all age groups) and $n \in \{male, female\}$. We adopt a LSTM encoder decoder framework, including two main components: sequence feature embedding and LSTM encoder decoder network.

4.1.1 Sequence Feature Embedding. Each time sequence s is characterized by a group $g_s = [c_s, a_s, n_s]$ and a fashion element f_s (we omit the subscript s if there is no confusion thereafter). All the categorical features are converted into dense vector representations. Particularly, to get the group representation, we first convert the group features c , a , and n into their embeddings $\mathbf{c} \in \mathbb{R}^D$, $\mathbf{a} \in \mathbb{R}^D$, and $\mathbf{n} \in \mathbb{R}^D$ separately, where D is the dimensionality of sequence feature embedding. We then adopt a linear layer to aggregate the three embeddings into one unified group representation:

$$\mathbf{g} = \mathbf{W}_g [\mathbf{c}, \mathbf{a}, \mathbf{n}] + \mathbf{b}_g \quad (1)$$

where $\mathbf{W}_g \in \mathbb{R}^{3D \times D}$, $\mathbf{b}_g \in \mathbb{R}^D$, and $\mathbf{g} \in \mathbb{R}^D$. For each fashion element f , we directly convert it into a vector $\mathbf{f} \in \mathbb{R}^D$.

4.1.2 LSTM Encoder Decoder Network. Most of the current methods for fashion trend forecasting [1, 27] model each time series independently, overlooking the correlations among them. However, many fashion elements or groups have high correlations with each other, and the correlations can help to learn the trend patterns. For example, the seasonal trends of *sweater* and *t-shirt* are opposite with each other, if we learn one trend well and the other one will also perform well with high probability. Therefore, in this paper, we utilize a deep learning model: LSTM encoder decoder framework, by designing one model for all the time series instead of one for each, to implicitly capture the correlations among time series.

The LSTM encoder decoder framework consists of two parts: encoder and decoder, as shown in Figure 3. The encoder is a LSTM network, which aims to map the historical inputs to latent representations \mathbf{h}_T^e . Specifically, we concatenate the group representation \mathbf{g} , the fashion element representation \mathbf{f} , the timestep feature \mathbf{m}_t (the position of each point within one year, converted to vector representation thus $\mathbf{m}_t \in \mathbb{R}^D$), and the trend value y_t as the input

of the encoder network at timestep t :

$$\mathbf{v}_t^e = [\mathbf{g}, \mathbf{f}, \mathbf{m}_t, y_t] \quad (2)$$

where $\mathbf{v}_t^e \in \mathbb{R}^{3D+1}$. The output of the encoder LSTM is the hidden representations for the input sequence at timestep t , denoted as:

$$\mathbf{h}_t^e = LSTM^e(\mathbf{v}_t^e; \mathbf{h}_{t-1}^e) \quad (3)$$

where $\mathbf{h}_{t-1}^e, \mathbf{h}_t^e \in \mathbb{R}^H$, and H is the size of the hidden state. \mathbf{h}_{t-1}^e is the encoder hidden state one step ahead of \mathbf{h}_t^e ,

The decoder network is a bi-directional LSTM, of which the initial hidden state is \mathbf{h}_T^e , i.e., the last hidden state of encoder, and at each decoding step it takes the input feature and outputs the trend forecasting value. The input feature of decoder network at timestep t is: $\mathbf{v}_t^d = [\mathbf{g}, \mathbf{f}, \mathbf{m}_t]$, which is different from \mathbf{v}_t^e by removing the trend value y_t and thus $\mathbf{v}_t^d \in \mathbb{R}^{3D}$. The bi-directional LSTM can propagate information from both forward and backward directions. Formally, we denote the hidden state from forward LSTM as $\overrightarrow{\mathbf{h}}_t^d$ and from backward as $\overleftarrow{\mathbf{h}}_t^d$. We can obtain the final hidden state \mathbf{h}_t^d by concatenating them as follows:

$$\begin{cases} \overrightarrow{\mathbf{h}}_t^d = LSTM^d(\overrightarrow{\mathbf{v}}_t^d; \overrightarrow{\mathbf{h}}_{t-1}^d) \\ \overleftarrow{\mathbf{h}}_t^d = LSTM^d(\overleftarrow{\mathbf{v}}_t^d; \overleftarrow{\mathbf{h}}_{t+1}^d) \\ \mathbf{h}_t^d = [\overrightarrow{\mathbf{h}}_t^d, \overleftarrow{\mathbf{h}}_t^d] \end{cases} \quad (4)$$

where $\overrightarrow{\mathbf{h}}_t^d, \overleftarrow{\mathbf{h}}_t^d, \overrightarrow{\mathbf{h}}_{t-1}^d, \overleftarrow{\mathbf{h}}_{t+1}^d \in \mathbb{R}^H$, and $\mathbf{h}_t^d \in \mathbb{R}^{2H}$. The prediction is made based on the hidden state of each step. Due to the teacher mechanism leveraged in the encoder network, the predictions are made in both the encoder and decoder stages during training. But for testing, predictions only happen in the decoder stage. Particularly, we apply linear layers for making predictions for encoder and decoder respectively:

$$\begin{cases} y_t^e = \mathbf{W}_e \mathbf{h}_t^e + b_e \\ y_t^d = \mathbf{W}_d \mathbf{h}_t^d + b_d \end{cases} \quad (5)$$

where $\mathbf{W}_e, \mathbf{W}_d \in \mathbb{R}^{1 \times 2H}$ and $b_e, b_d \in \mathbb{R}$ are the parameters for the linear layer; $y_t^e, y_t^d \in \mathbb{R}$ are the forecasting value at each timestep for encoder and decoder respectively. We use L1 loss to train the LSTM, including the encoder loss $L_e(\cdot)$ and decoder loss $L_d(\cdot)$:

$$L_{LSTM} = L_e(\mathbf{y}_e, \mathbf{y}_e^*, \boldsymbol{\theta}_e) + L_d(\mathbf{y}_d, \mathbf{y}_d^*, \boldsymbol{\theta}_d) \quad (6)$$

where $\boldsymbol{\theta}_e, \boldsymbol{\theta}_d$ are the model parameters for encoder and decoder respectively; $\mathbf{y}_e, \mathbf{y}_e^* \in \mathbb{R}^{(T-1)}$ are the prediction and ground-truth of the encoder sequence; and $\mathbf{y}_d, \mathbf{y}_d^* \in \mathbb{R}^{T'}$ are the prediction and ground-truth of the decoder sequence.

4.2 Knowledge Incorporation

Formally, we leverage two types of knowledge: internal knowledge within the dataset and external knowledge extracted from the fashion element taxonomy.

4.2.1 Internal Knowledge. Fashion trend sequences demonstrate high correlation with each other. Particularly, some sequences have similar patterns while some sequences have opposite patterns. For example, as shown in Figure 1, the time series pattern of *turtle neck* is similar with *sweater* but opposite with *dress*. Such prior

knowledge is valuable to guide the learning of the model. We deem that the hidden representations of time series with similar patterns should be close to each other, and that with opposite patterns should be far away from each other. We generalize such *similar-opposite* relation to general *close-far* relation, that is: given one fashion trend sequence \mathbf{y}^k , we can always find another two series \mathbf{y}^p and \mathbf{y}^q , where \mathbf{y}^p is closer to \mathbf{y}^k than \mathbf{y}^q , and $\mathbf{y}^k, \mathbf{y}^p, \mathbf{y}^q \in \mathbb{R}^{(T+T')}$. We use L1 distance (denoted as $\|\cdot\|$ thereafter) to measure the similarity. Formally these three sequences should satisfy:

$$\|\mathbf{y}_k - \mathbf{y}_p\| < \|\mathbf{y}_k - \mathbf{y}_q\| \quad (7)$$

We construct a triplet regularization term $r^{k,p,q}$ as follows:

$$\begin{cases} d_t^{k,p} = \|\mathbf{h}_t^k - \mathbf{h}_t^p\| \\ d_t^{k,q} = \|\mathbf{h}_t^k - \mathbf{h}_t^q\| \\ r^{k,p,q} = \frac{1}{T+T'-1} \sum_{t=0}^{T+T'-1} \max(0, d_t^{k,p} - d_t^{k,q}) \end{cases} \quad (8)$$

where $\mathbf{h}_t^k, \mathbf{h}_t^p$, and \mathbf{h}_t^q are the LSTM hidden states in both encoder and decoder stages for sequences k, p, q respectively; $d_t^{k,p}, d_t^{k,q} \in \mathbb{R}$ are the hidden state distances between (k, p) and (k, q) respectively. We randomly sample the (k, p, q) triplet from the whole dataset. Thus the final loss of our framework is as follows:

$$L^{k,p,q} = \sum_{s \in \{k,p,q\}} \left(L_e^s(\cdot) + L_d^s(\cdot) \right) + \lambda r^{k,p,q} \quad (9)$$

where λ is a hyper parameter determining the weight of the regularization term.

4.2.2 External Knowledge. In addition to the internal knowledge of *close-far* relations observed from the dataset, we leverage external knowledge from fashion element taxonomy to help in enhancing the model. Generally speaking, all the fashion elements are usually organized into a hierarchical taxonomy with a tree structure. There exist affiliation relations between children nodes and their associated parent nodes, which will further affect their corresponding trend sequences. For example, if we find that the trend of the attribute *peplum* goes up, it is highly possible that the category *dress* also goes up since *peplum* is an attribute of category *dress*. Different from the aforementioned similarity correlation which is pair-wise, the affiliation relation is many-to-one. For instance in Figure 4, the category *dress* has four attributes and the attribute *shape* has six values. To model such complicated relationships, we propose to construct a tree among all the fashion elements and conduct message passing between nodes with affiliation relationships.

A part of the constructed tree is shown in Figure 4. Basically, we have three types of nodes in this tree: *category*, *attribute*, and *attribute value*, and the affiliation relations are between *attribute* and *category*, *attribute value* and *attribute*. As mentioned in Section 4.1.1, each fashion element f is converted to a vector representation \mathbf{f} . Therefore we conduct message passing among those embeddings, i.e., passing messages from children nodes to their parent nodes. The message passing for node i is as follows:

$$\begin{cases} \mathbf{m}_i = \sum_{j \in N_i} w_j \times \mathbf{f}_j \\ \mathbf{f}_i \leftarrow \mathbf{f}_i + \mathbf{m}_i \end{cases} \quad (10)$$

where $m_i \in \mathbb{R}^D$ is the message passed from its children nodes, N_j^S is the set of nodes that have affiliation relations with node i , and w_j is the weight of each relation and can be learned during training. Note that we initialize the $w_j \in \mathbb{R}$ as the portion of node j out of all nodes affiliated to the parent node, thus $\sum_{j \in N_i} w_j = 1$.

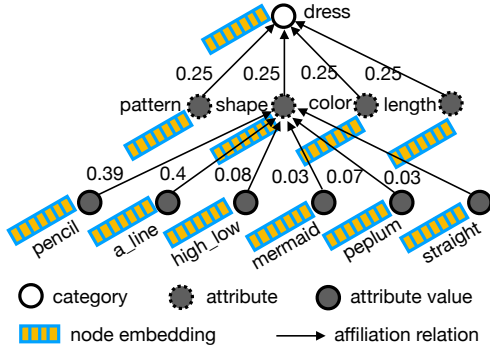


Figure 4: An illustration of external knowledge incorporation module.

5 EXPERIMENTS

To verify the effectiveness of our proposed approach, we conduct extensive experiments on two datasets. In particular, we are interested in the following research questions:

- (1)**RQ1**: Does our KERN model outperform current state-of-the-arts methods in predicting future fashion trend?
- (2)**RQ2**: Whether the introduced relation modules help in improving the performance and how do they help?
- (3)**RQ3**: How does the proposed model perform in trend forecasting in terms of specific fashion elements, and based on that, how can the model produce insightful fashion trend forecasting.

5.1 Experimental Settings

Experimental Setup. We apply two fashion trend forecasting datasets, our proposed FIT dataset and the GeoStyle dataset [27]. For GeoStyle, we take one-year of data (52 data points) as input to predict the output of the following half year (26 data points). We design two settings on FIT, 1) use two years of data (48 data points) as input and predict the output of the following half-year (12 data points); and 2) use two years of data (48 data points) as input and predict the output of the following one-year (24 data points). Since Geostyle dataset only has one attribute (city) and does not have user attributes of age and gender, we simplify the fusion of group attributes and directly use the city embedding as group embedding. Sliding windows strategy was applied on both datasets to generate the training and testing samples. More details of data preparation can be found in Figure 5, which shows the procedure of splitting each time series into shorter samples. We use Mean Absolute Error (MAE) and Mean Absolute Percentage Error (MAPE) as the evaluation metrics [27].

Implementation Details. We set the embedding size of user embeddings (including age embedding, gender embedding and city embedding), fashion element embedding and time embedding to 10 and the hidden size of both encoder and decoder LSTM network




Figure 5: Illustration of generating samples for each time series. Each sample is composed of the input sequence (GREEN) and output sequence (BLUE). Sliding window strategy is applied to generate the training samples. Note that there are two settings on FIT dataset, marked by (1) and (2).

to 50. The hyper parameter λ for regularization is set to 2×10^{-4} . During training, we randomly sample a batch of 400 different time series for each iteration. For each experimental setting, we train one KERN model for all fashion elements. For each testing sample, we calculate the performance of the odd data points as the validation results, and that of even data points as the testing results.

Baseline Methods. We select several state-of-the-arts time series prediction methods to compare with the proposed KERN model: **Mean and Last**: They use the mean value or the value of last point of the input historical data as the forecasting value. **Autoregression (AR)**: It is a linear regressor which uses a linear combination of last few observed values as the forecasting value. **Vector Autoregression (VAR)**: VAR is a generalization of the AR by allowing for more than one evolving variable. **Exponential Smoothing (ES)** [1]: It aggregates all the historical values with an exponential decayed weight, the more recent values have higher impact on the future’s forecast. **Linear and Cyclic** [27]: They are linear or cyclical parametric model which let historical values to fit the specific predefined model. **Geostyle** [27]: It is a parametric model combining a linear component and a cyclical component. It is the state-of-the-art fashion trend forecasting method on Geostyle dataset.




Figure 6: Two examples of trends forecasting made by KERN and KERN-IE models.

Table 1: Performance of KERN and baselines for fashion trend Forecasting (the lower is better)

Dataset	GeoStyle		FIT			
	Half year		Half year		One year	
Method	MAE	MAPE	MAE	MAPE	MAE	MAPE
Mean	0.0292	25.79	0.132	65.31	0.135	63.21
Last	0.0226	21.04	0.125	46.45	0.147	54.04
AR	0.0211	20.69	0.114	54.36	0.119	51.96
VAR	0.0150	17.95	0.157	62.97	0.126	47.35
ES	0.0228	20.59	0.133	55.29	0.150	57.42
Linear	0.0365	24.40	0.112	43.30	0.133	45.89
Cyclic	0.0165	16.64	0.129	49.92	0.143	51.66
GeoStyle	0.0149	16.03	0.136	52.40	0.149	53.14
KERN	0.0134	14.24	0.083	30.02	0.094	33.45
improv(%)	10.07	11.17	25.89	30.67	21.01	27.11

Table 2: Contribution of different knowledge in KERN model (MAE results). ‘-E’ means without external knowledge and ‘-I’ means without internal knowledge and ‘-IE’ means neither is used. Since there is no taxonomy among Geostyle’s fashion elements, KERN-E is left empty.

Dataset	GeoStyle		FIT	
	Half year	Half year	Half year	One year
Prediction				
KERN-IE	0.0137	0.0840	0.0966	
KERN-E	-	0.0835	0.0953	
KERN-I	0.0134	0.0831	0.0942	
KERN	0.0134	0.0836	0.0939	

5.2 Overall Performance on Fashion Trend forecasting(RQ1)

We first evaluate our KERN model in terms of fashion trend forecasting by comparing its performance with several classic and state-of-the-arts baselines. The overall results are shown in Table 1. Based on the results, we have the following observations:

(1) The proposed KERN model consistently yields the best performance on both datasets and under all experimental settings. Specifically, on the GeoStyle, KERN is the only method to achieve the MAE result of lower than **0.014** and MAPE lower than **15**. On the two settings of FIT dataset, MAE and MAPE results of all baselines are over **0.11** and **40**, but our KERN model achieve MAE of under **0.1** and MAPE of around **30**, both show quite superior performance.

(2) On the FIT dataset, the KERN method outperforms all baselines with large margins. As the FIT dataset contains much more fine-grained fashion elements, more user information, and more realistic and complex time series patterns, it is more challenging to model. Therefore, the baseline methods do not perform well. However, our KERN method is able to capture such complex patterns compared with other baselines as it better models nonlinearity in data by the LSTM encoder decoder framework and leverages the abundant domain knowledge. That also explains why our model achieves limited improvement on GeoStyle.

(3) Most methods perform better in half-year prediction than in one-year prediction on the FIT dataset, including our KERN model.

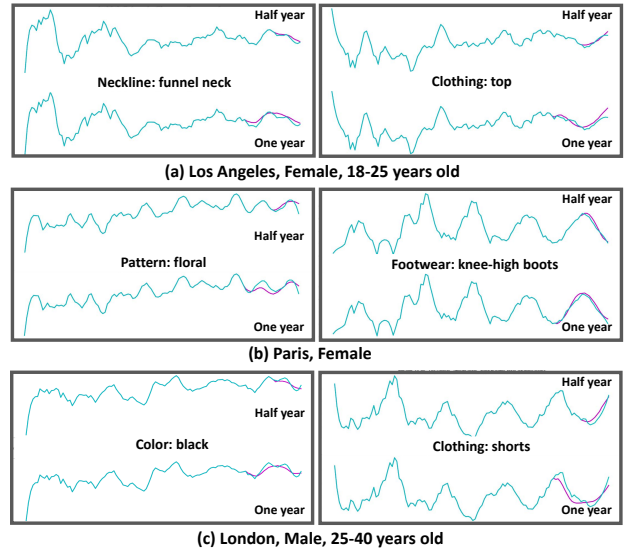


Figure 7: Examples of trend forecasting for different fashion elements and user groups. GREEN curves are ground truth and PURPLE curves are prediction results by KERN.




Figure 8: The prediction results of one fashion element of different user groups. The performance is better when applying user information in data modeling.

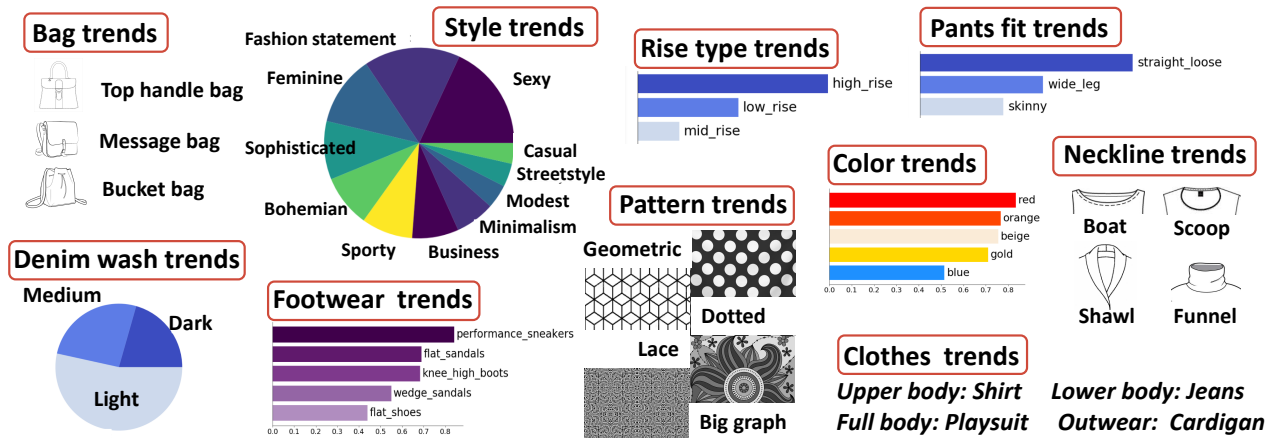


Figure 9: The fashion trend report generated based on KERN forecasting results for 2018.07. This report is targeted on women in Paris, and predicted based on fashion records from 2016.07 to 2018.06.

Such results are reasonable since first, the one-year prediction requires to forecast data with longer time horizon. Second, such setting reduces the quantity of training data (see Figure 5).

5.3 Discussion on Effectiveness of Knowledge Incorporation (RQ2)

Here we conduct experiments and discuss the effectiveness of incorporating knowledge in the fashion trend forecasting model. In particular, two types of knowledge are exploited. The first is the internal knowledge of similarity relations, utilized by introducing triplet regularization term in the loss function. The second is the external knowledge of affiliation relations built from taxonomy, utilized by updating the fashion element embeddings by message passing according to the affiliation tree.

To evaluate the effectiveness of each type of knowledge, we test the model by removing each type one at a time, as well as both simultaneously. Specifically, the model without using internal knowledge (KERN-I) ignores the triplet regularization loss (Eqn. 8) while the model without using external knowledge (KERN-E) is not equipped with the external knowledge incorporation module. The KERN-IE model contains neither. From the performance reported in Table 2, we can see that removing either affiliation relations or similarity relations could degrade the performance of our model. The difference made by incorporating knowledge is more significant for longer forecasting on FIT dataset, which shows that leveraging knowledge could be particularly helpful for complicated and challenging data. Figure 6 shows two examples that compare the trend forecasting results of KERN and the ablated KERN-IE. The prediction results of KERN are clearly better than that of KERN-IE, which further shows that KERN benefits from effectively applying internal and external knowledge.

5.4 Fashion Trend Analysis (RQ3)

To further illustrate the effectiveness of our KERN model for fashion forecasting, we show more visualization results in this part. We first show the one-year prediction and half-year prediction of six representative fashion elements (including clothing category,

footwear category, color, pattern and attributes) for the three user groups as shown in Figure 7. We can see that, in general, KERN can predict the trends of fashion element very well for both the one year or half year prediction, even for those with rather complex patterns. Results in Figure 8 show that the same fashion element between different user groups can be different. It also demonstrates that the KERN method can effectively predict the different trends by leveraging user information because the prediction results using user information is much better than those without.

Based on the forecasting of various types of fashion elements, we can further generate a comprehensive and professional fashion trend report (see Figure 9) that covers a large number of fashion trends ranging from category, color, pattern, style, or even detailed attributes such as denim wash colors.

6 CONCLUSION AND FUTURE WORK

This paper addresses the fashion trend forecasting problem based on social media, aiming to mine the complex patterns in the historical time-series records of fashion elements and accordingly predict the future trends. An effective model, Knowledge Enhanced Recurrent Network (KERN) is proposed to capture the complex patterns in the time-series data and forecast fashion trends.

Although much effort has been made and desirable results have been achieved, there are some aspects that can be further improved in the future. First, more user information should be explored such as occupations or hobbies. Second, multiple sources of knowledge should be considered such as the fashion analysis derived from professional fashion magazines, fashion bloggers and brands.

ACKNOWLEDGEMENT

This research is supported by the National Research Foundation, Singapore under its International Research Centres in Singapore Funding Initiative, and The Hong Kong Polytechnic University (project code: RHQK). Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not reflect the views of National Research Foundation, Singapore. We also appreciate the fashion recognition API service provided by Visenze.

REFERENCES

- [1] Ziad Al-Halah, Rainer Stiefelwagen, and Kristen Grauman. 2017. Fashion forward: Forecasting visual style in fashion. In *Proceedings of the IEEE International Conference on Computer Vision*. 388–397.
- [2] Grigory Antipov, Moez Baccouche, Sid-Ahmed Berrani, and Jean-Luc Dugelay. 2017. Effective training of convolutional neural networks for face-based gender and age prediction. *Pattern Recognition* 72 (2017), 15–26.
- [3] George EP Box and Gwilym M Jenkins. 1968. Some recent advances in forecasting and control. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 17, 2 (1968), 91–109.
- [4] Yagmur Gizem Cinar, Hamid Mirisae, Parantapa Goswami, Eric Gaussier, Ali Ait-Bachir, and Vadim Strijov. 2017. Position-based content attention for time series forecasting with sequence-to-sequence rnns. In *International Conference on Neural Information Processing*. Springer, 533–544.
- [5] Yujuan Ding and Wai Keung Wong. 2019. Fashion Outfit Style Retrieval Based on Hashing Method. In *Artificial Intelligence on Fashion and Textiles*, Wai Keung Wong (Ed.). Springer International Publishing, 187–195.
- [6] Chenyou Fan, Yuze Zhang, Yi Pan, Xiaoyue Li, Chi Zhang, Rong Yuan, Di Wu, Wensheng Wang, Jian Pei, and Heng Huang. 2019. Multi-horizon time series forecasting with temporal attention learning. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2527–2535.
- [7] Fuli Feng, Xiangnan He, Xiang Wang, Cheng Luo, Yiqun Liu, and Tat-Seng Chua. 2019. Temporal relational ranking for stock prediction. *ACM Transactions on Information Systems (TOIS)* 37, 2 (2019), 1–30.
- [8] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech recognition with deep recurrent neural networks. In *2013 IEEE international conference on acoustics, speech and signal processing*. IEEE, 6645–6649.
- [9] Xintong Han, Zuxuan Wu, Yu-Gang Jiang, and Larry S Davis. 2017. Learning fashion compatibility with bidirectional lstms. In *Proceedings of the 25th ACM international conference on Multimedia*. 1078–1086.
- [10] Xintong Han, Zuxuan Wu, Zhe Wu, Ruiqi Yu, and Larry S Davis. 2018. Viton: An image-based virtual try-on network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 7543–7552.
- [11] Shintami C. Hidayati, Kai-Lung Hua, Wen-Huang Cheng, and Shih-Wei Sun. 2014. What Are the Fashion Trends in New York?. In *Proceedings of the 22nd ACM International Conference on Multimedia*. 197–200.
- [12] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [13] Charles C Holt. 2004. Forecasting seasonals and trends by exponentially weighted moving averages. *International journal of forecasting* 20, 1 (2004), 5–10.
- [14] Chia-Wei Hsieh, Chieh-Yun Chen, Chien-Lung Chou, Hong-Han Shuai, Jiaying Liu, and Wen-Huang Cheng. 2019. FashionOn: Semantic-guided Image-based Virtual Try-on with Detailed Human and Clothing Information. In *Proceedings of the 27th ACM International Conference on Multimedia*. 275–283.
- [15] Han Hu, Jiayuan Gu, Zheng Zhang, Jifeng Dai, and Yichen Wei. 2018. Relation networks for object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 3588–3597.
- [16] 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.
- [17] Chenhan Jiang, Hang Xu, Xiaodan Liang, and Liang Lin. 2018. Hybrid knowledge routed modules for large-scale object detection. In *Advances in Neural Information Processing Systems*. 1552–1563.
- [18] Eundeok Kim, Ann Marie Fiore, and Hyejeong Kim. 2013. *Fashion trends: analysis and forecasting*. Berg.
- [19] Martin Långkvist, Lars Karlsson, and Amy Loutfi. 2014. A review of unsupervised feature learning and deep learning for time-series modeling. *Pattern Recognition Letters* 42 (2014), 11–24.
- [20] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature* 521, 7553 (2015), 436.
- [21] Lizi Liao, Xiangnan He, Bo Zhao, Chong-Wah Ngo, and Tat-Seng Chua. 2018. Interpretable multimodal retrieval for fashion products. In *Proceedings of the 26th ACM international conference on Multimedia*. 1571–1579.
- [22] Lizi Liao, Yunshan Ma, Xiangnan He, Richang Hong, and Tat-seng Chua. 2018. Knowledge-aware multimodal dialogue systems. In *Proceedings of the 26th ACM international conference on Multimedia*. 801–809.
- [23] Si Liu, Zheng Song, Guangcan Liu, Changsheng Xu, Hanqing Lu, and Shuicheng Yan. 2012. Street-to-shop: Cross-scenario clothing retrieval via parts alignment and auxiliary set. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 3330–3337.
- [24] 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.
- [25] Yunshan Ma, Lizi Liao, and Tat-Seng Chua. 2019. Automatic Fashion Knowledge Extraction from Social Media. In *Proceedings of the 27th ACM International Conference on Multimedia*. 2223–2224.
- [26] Yunshan Ma, Xun Yang, Lizi Liao, Yixin Cao, and Tat-Seng Chua. 2019. Who, Where, and What to Wear? Extracting Fashion Knowledge from Social Media. In *Proceedings of the 27th ACM International Conference on Multimedia*. 257–265.
- [27] Utkarsh Mall, Kevin Matzen, Bharath Hariharan, Noah Snavely, and Kavita Bala. 2019. GeoStyle: Discovering Fashion Trends and Events. In *Proceedings of the IEEE International Conference on Computer Vision*. 411–420.
- [28] Kevin Matzen, Kavita Bala, and Noah Snavely. 2017. Streetstyle: Exploring worldwide clothing styles from millions of photos. *arXiv preprint arXiv:1706.01869* (2017).
- [29] Tomáš Mikolov, Martin Karafiát, Lukáš Burget, Jan Černocký, and Sanjeev Khudanpur. 2010. Recurrent neural network based language model. In *Eleventh annual conference of the international speech communication association*.
- [30] Gabriel Panis, Andreas Lanitis, Nicholas Tsapatsoulis, and Timothy F Cootes. 2016. Overview of research on facial ageing using the FG-NET ageing database. *let Biometrics* 5, 2 (2016), 37–46.
- [31] Joseph Redmon and Ali Farhadi. 2018. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767* (2018).
- [32] Rasmus Rothe, Radu Timofte, and Luc Van Gool. 2018. Deep expectation of real and apparent age from a single image without facial landmarks. *International Journal of Computer Vision* 126, 2–4 (2018), 144–157.
- [33] Eugen Slutsky. 1937. The summation of random causes as the source of cyclic processes. *Econometrica: Journal of the Econometric Society* (1937), 105–146.
- [34] Xuemeng Song, Fuli Feng, Xianjing Han, Xin Yang, Wei Liu, and Liqiang Nie. 2018. Neural compatibility modeling with attentive knowledge distillation. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 5–14.
- [35] Xuemeng Song, Fuli Feng, Jinhuan Liu, Zekun Li, Liqiang Nie, and Jun Ma. 2017. Neurostylist: Neural compatibility modeling for clothing matching. In *Proceedings of the 25th ACM international conference on Multimedia*. 753–761.
- [36] Ilya Sutskever, James Martens, and Geoffrey E Hinton. 2011. Generating text with recurrent neural networks. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*. 1017–1024.
- [37] Subhashini Venugopalan, Lisa Anne Hendricks, Raymond Mooney, and Kate Saenko. 2016. Improving lstm-based video description with linguistic knowledge mined from text. *arXiv preprint arXiv:1604.01729* (2016).
- [38] Sirion Vittayakorn, Kota Yamaguchi, Alexander C Berg, and Tamara L Berg. 2015. Runway to realway: Visual analysis of fashion. In *2015 IEEE Winter Conference on Applications of Computer Vision*. IEEE, 951–958.
- [39] Gilbert Thomas Walker. 1931. On periodicity in series of related terms. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 131, 818 (1931), 518–532.
- [40] Wenguan Wang, Yuanlu Xu, Jianbing Shen, and Song-Chun Zhu. 2018. Attentive fashion grammar network for fashion landmark detection and clothing category classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4271–4280.
- [41] Xiang Wang, Dingxian Wang, Canran Xu, Xiangnan He, Yixin Cao, and Tat-Seng Chua. 2019. Explainable reasoning over knowledge graphs for recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 5329–5336.
- [42] Ruofeng Wen, Kari Torkkola, Balakrishnan Narayanaswamy, and Dhruv Madeka. 2017. A multi-horizon quantile recurrent forecaster. *arXiv preprint arXiv:1711.11053* (2017).
- [43] Peter R Winters. 1960. Forecasting sales by exponentially weighted moving averages. *Management science* 6, 3 (1960), 324–342.
- [44] Xun Yang, Xiangnan He, Xiang Wang, Yunshan Ma, Fuli Feng, Meng Wang, and Tat-Seng Chua. 2019. Interpretable Fashion Matching with Rich Attributes. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 775–784.
- [45] Xun Yang, Yunshan Ma, Lizi Liao, Meng Wang, and Tat-Seng Chua. 2019. Transfcm: Translation-based neural fashion compatibility modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 403–410.
- [46] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. 2016. Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters* 23, 10 (2016), 1499–1503.