Rule-Guided Knowledge-Graph based Negative Sampling for Outfit Recommendation

Huijing Zhan
zhan_huijing@i2r.a-star.edu.sg
Institute for Infocomm Research

Ling Li
LING014@e.ntu.edu.sg
Nanyang Technological University

Xue Geng
geng_xue@i2r.a-star.edu.sg
Institute for Infocomm Research

Jie Lin
lin-j@i2r.a-star.edu.sg
Institute for Infocomm Research

A. C. Kot
eackot@ntu.edu.sg
Nanyang Technological University

ABSTRACT

Recommender system (RS) has become increasingly prevalent among online service providers. The effectiveness of an accurate RS highly relies on the quality level of the selected negative instances for training. Most of the existing negative sampling strategies either merely leverage the interaction data which suffer from the sparsity challenge or fail to fully utilize the side information. To address these issues, we introduce rule-guided knowledge graph (RuleKG) by integrating the enriched relations of the knowledge graph (KG) and category-aware fashion rules into the outfit recommendation problem. The rules are further incorporated into a newly designed score function which represents the user's preferences toward outfit in a more fine-grained perspective. Given a user-outfit pair, the negative candidates are explored in both outfit- and user-level. Also, the reinforcement learning (RL)-based strategy is developed to automatically choose the next state from the starting point over KG. Experimental results on the new top-bottom outfit dataset demonstrate the superiority of the proposed approach and the generality of the negative sampler model is validated on the music recommendation benchmark.

CCS CONCEPTS

• **Information systems → Recommender systems**; *Recommender systems*; *Personalization*.

KEYWORDS

Recommender system, knowledge graph, negative sampling, reinforcement learning

ACM Reference Format:
Huijing Zhan, Ling Li, Xue Geng, Jie Lin, and A. C. Kot. 2022. Rule-Guided Knowledge-Graph based Negative Sampling for Outfit Recommendation. In *Proceedings of International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'22).* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/1122445.1122456

SIGIR'22, July 11–15, 2022, Madrid, Spain

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

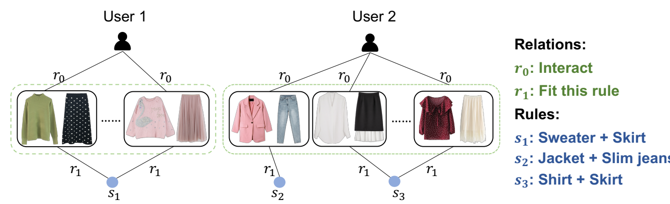

Fig. 1: Examples of users and the outfits they interact with diverse category combinations.

1 INTRODUCTION

In recent years, we have witnessed the widespread application of the recommender system in a variety of real-world scenarios [6, 7, 10, 17, 26, 27, 31]. Even though traditional matrix factorization (MF)-based approaches [9] have achieved excellent performances, it suffers from data sparsity problem. Knowledge-aware framework in the context of knowledge graph aims to resolve these issues by exploiting auxiliary data sources [5]. Moreover, the quality of the negative user-item pairs has a significant impact on the recommendation accuracy. Existing negative sampling strategies [1, 11, 23, 25, 28] suffer from the problem of vanishing gradient [30]. With the huge success of Generative Adversarial Network (GAN) [4] in modeling the dynamic distribution of negative instances, IRGAN [18] and KBGAN [2] propose to view the generator as sampler, which is trained in a minimax game with the discriminator. However, these negative sampling approaches don't fully exploit the advantage of KG.

Recent studies [19, 22, 24, 32] have demonstrated the success of incorporating the knowledge graph into reinforcement learning-based algorithms to formalize the negative sampling processing as a Markov Decision Process (MDP) [15] with properly designed reward functions. It poses great challenges in how to incorporate the RL-based scheme into the knowledge graph framework with rich relations between entities. Also, fine-grained attributes are required to characterize the overall look of the outfit for capturing the user's individual taste. The recent work [22] also employs RL-based strategy for negative sampling. However, the major difference between the proposed approach and theirs is that they explore the negatives based on multi-hop connection over graph and fail to exploit the critical signals of negatives from the user perspective. Fig. 1 illustrates two exemplar users and the outfits they have clicked triggering different category-related fashion matching rules. User 1 prefers the sweater goes with the skirt while User 2 is more interested in matching that with shirt.

To address the above-mentioned issues, we propose a novel Rule-Guided Knowledge-Graph based Negative Sampling for Outfit Recommendation (RuleKG) framework, which is comprised of a reinforcement learning-based sampler to generate high-quality negatives and a rule-guided recommender. More specifically, a set of structured category-aware fashion rules are defined based on fashion domain knowledge, which are further leveraged to expand the knowledge graph and incorporated for user preference modeling via a new rule-guided score function. The attention-based graph convolutional network is utilized to differentiate the importance of neighbors revealing the user's distinctive preference for outfits with different category combinations. And we construct a pseudo exposure data leveraging outfit's textual descriptions and the side information provided by KG, which is utilized to form the action space of a hierarchical RL agent. What is more, the user-outfit pairs are converted into triples for negative sampling in the perspective of both the user and outfit. Experimental results demonstrate the excellence of our proposed method over state-of-the-art methods in the fashion outfit recommendation.

In summary, the contribution of our proposed method can be summarized as below:

- We propose a novel knowledge-aware approach to sample negatives with hierarchical reinforcement learning strategy and introduce category-aware fashion rules.
- We utilize the rich information provided by KG to construct pseudo exposure data and design a novel rule-guided score function for negative sampling.
- The user-outfit pairs are transformed into the KG triples to explore the user- and outfit-aware negatives to improve the diversity of negative instances.

2 METHODOLOGY

2.1 Overview of the Framework

The overall framework of RuleKG is shown in Fig. 2, which consists of a RL-based sampler and a recommender. For the sampler, we develop a hierarchical reinforcement learning strategy to perform the candidate negative instance selection. Given a positive user-outfit (u, o) pair with the "interact" relation r_0^+, the sampler generates the high-quality negatives e_f^u and e_f^o. For the recommender, the user's preference towards outfit matching rules concerning category compositions is encoded.

2.2 Rule Construction

To capture the user's distinctive preference over category combinations of the outfit, we aim to define a set of category-aware fashion rules, which are designed by referring to the outfit descriptions in Chinese provided by the internal dataset (*i.e.,* Alibaba-iFashion [3]) and the structured matching rules in English proposed in the external dataset (*i.e.* FashionVC [14]). There exists a cross-lingual language challenge to unify the rules of these two fashion outfit dataset. The rule defined in [14] is composed of a specific category pair with top and bottom. And we utilize EasyNMT[1], a state-of-the-art neural translation tool, to translate the fashion category name in English to that in mandarin. Category combinations with low co-occurrence frequencies are removed based on their occurrence in Alibaba-iFashion. If the description of a given outfit contains all the keywords in the rule, then it is supposed to trigger the specific rule. Since the outfits can be activated with more than one rule, the most fine-grained rule are determined based on the granularity score. For example, "T-shirt + tight skinny jeans" is selected over "top + bottom" as the final candidate rule for an outfit o. Because the former specifies the category combination in a more fine-grained manner thus achieves higher granularity score. The selected rule, denoted as v_o as shown in Eq. 6, is incorporated into the knowledge graph by establishing a connection between the target outfit node and the specific rule node with the newly appended relation "outfit_has_rule". And the user preference toward the outfit is further characterized by the interaction between user embedding $\mathbf{e}_u$ and rule representation $\mathbf{e}_{v_o}$, capturing the user tastes in a fine-grained and interpretable manner.

[1]https://github.com/UKPLab/EasyNMT

2.3 Sampler

The task of the sampler is formularized as a hierarchical Markov Decision Process (MDP). More specifically, the overall task is decomposed into three levels of sub-tasks: 1) epoch-level 2) batch-level and 3) triple-level. Each task is defined as a 5-tuple MDP $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \pi)$, corresponding to a set of states, actions, state transition probabilities, reward and policy functions, respectively. The final candidate negatives selection aims to combine the uniformly sampled output from 2) and probability-distribution based output from 3) via the rule-guided score function.

Pseudo Exposure Data. The exposure data, indicating whether the exposed item has been interacted by the user, encodes important clues about the user's preference; however, such information is not always available. We propose to construct a pseudo exposure data which consists of three sets of user-outfit pairs: 1) clicked; 2) exposed but not clicked; and 3) non-observed entries, denoted as ξ_1, ξ_2, ξ_3, respectively. Here $\xi = \{\xi_1, \xi_2, \xi_3\}$ denotes the whole entity set. Previous studies merely exploit the data from ξ_1 and ξ_2, ignoring the side information provided by ξ_3.

State. The epoch-level task takes an action to construct the negative sampling neighbor set ξ_e for the triple-level based on the state feature $\mathbf{s}_e$. It is defined as the difference between the consecutive epochs in terms of recall, precision, NDCG, hit ratio@K and the average reward obtained from the previous epoch of data. The state feature $\mathbf{s}_b$ of the batch-level task is calculated as the average reward from the previous epoch of data and that from the current batch, which aims to determine the candidate negative set ξ_b for uniform sampling. As for the triple-level task, which guides the agent to select the high-quality negatives over the neighbor set, the state feature $\mathbf{s}_t$ is represented as a tuple (e_t^u, r_0^-, e_t^o) with the initial state as (u, r_0^+, o). Here r_0^+ and r_0^- indicate the positive and negative interaction between the user u and outfit o, respectively.

Action and Policy. The action of the agent is regulated by the policy function set $\pi = \{\pi_e, \pi_b, \pi_t^o, \pi_t^u\}$. And π_e and π_b, determining which set within S to choose from, are implemented by a fully-connected (FC) layer and a softmax layer with the parameter $\mathbf{W}_e$, $\mathbf{W}_b$, respectively. The negative instances of the triple-level task are generated by replacing either the user or the outfit node, termed

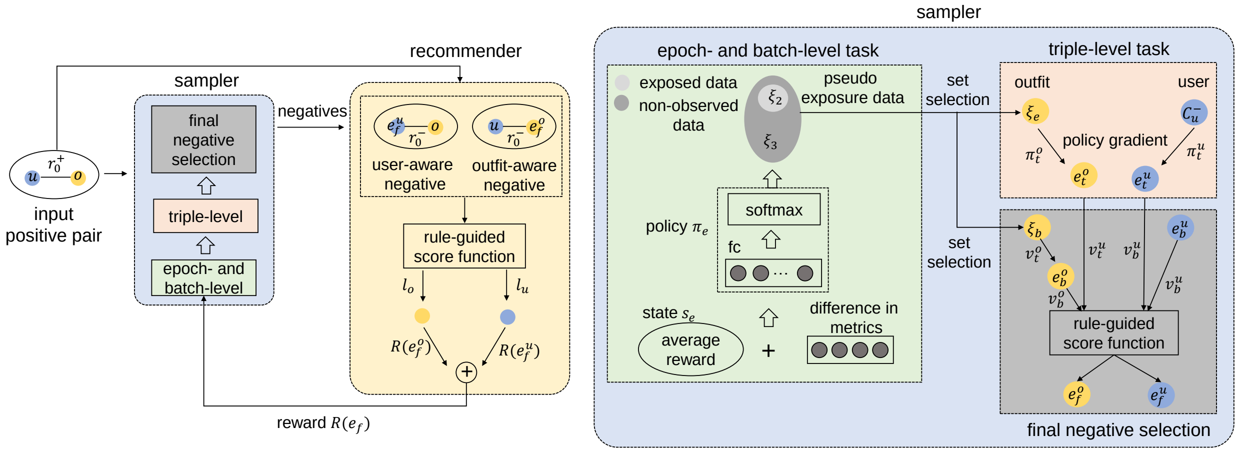

Fig. 2: The framework of the proposed RuleKG for outfit recommendation.

as user- and outfit-aware negatives. The policy π_t^u to select the user-aware negatives e_t^u is represented as the probability distribution among the neighbor set $\mathcal{N}_u$ of target user node u, computed as below:

$$P(u', r_0^-, o|u, r_0^+, o) = \frac{\exp\left(f_G(u', r_0^-, o)\right)}{\sum_{u''\in\mathcal{N}_u}\exp\left(f_G(u'', r_0^-, o)\right)}, \quad (1)$$

where f_G represents the inner product between the embedding of $\mathbf{e}_{u'}$ and $\mathbf{e}_o$, encoding their affinity in the entity space. Due to the space constraint, the description of π_t^o to select the outfit-aware negatives e_t^u is omitted, which performs in the similar manner.

Final Negative Selection. Given a positive triple (u, r_0^+, o), the final negative instances $e_f = \{e_f^o, e_f^u\}$ are chosen based on the rule-guided user preference score of the samples generated from the triple-level and the batch-level task, denoted as (e_t^o, e_t^u) and (e_b^o, e_b^u), respectively. Here we take the outfit-aware negative instances e_f^o for demonstration and that for the user-aware negative sampling of e_f^u follows the similar manner. The final selection of e_f^o is formalized as below:

$$e_f^o = \begin{cases} e_t^o, & y(u, e_t^o, v_t^o) >= y(u, e_b^o, v_b^o) \\ e_b^o, & y(u, e_t^o, v_t^o) < y(u, e_b^o, v_b^o), \end{cases} \quad (2)$$

where v_b^o and v_t^o indicate the mapped rule for e_b^o and e_t^o, respectively. The user preference modeling function incorporating the rules is represented as follows. Here $\mathbf{e}_{v_t^o}$ refers to the rule embedding of outfit e_t^o.

$$y(u, e_t^o, v_t^o) = \langle \mathbf{e}_u, \mathbf{e}_t^o \rangle + w_3 * \langle \mathbf{e}_u, \mathbf{e}_{v_t^o} \rangle, \quad (3)$$

where w_3 is the coefficient balancing the score between the user preference over the outfit and that over the rule.

Reward. Inspired by [22], the reward function $\mathcal{R}(e_t)$, reflecting the quality of the sampled negatives, as shown below:

$$\mathcal{R}(e_f) = w_1 * (\langle \mathbf{e}_u, \mathbf{e}_f \rangle + \langle \mathbf{e}_o, \mathbf{e}_f \rangle) + w_2 * \sigma(\beta \langle \mathbf{e}_u, \mathbf{e}_f \rangle), \quad (4)$$

where β and σ refer to the scaling factor and the sigmoid function, respectively. Here w_1 and w_2 denote the weighting coefficient.

Objective Function. Since the sampling process involves discrete sampling and SGD [13] cannot be directly utilized to compute the gradients. The REINFORCE [16] is introduced to solve this issue, and we maximize the following cumulative discounted reward. The gradients of L_S with respective to the sampler parameter θ_S is calculated as follows:

$$\nabla_{\theta_S} L_S = \nabla_{\theta_S} \sum_{e_f^u \in C_u^-} \mathbb{E}_{\pi_t^u}[\mathcal{R}(e_f^u)] + \nabla_{\theta_S} \sum_{e_f^o \in \xi_e} \mathbb{E}_{\pi_t^o}[\mathcal{R}(e_f^o)] \quad (5)$$

where N refers to the number of triples (u, r_0, o) in the training set Ω.

2.4 Recommender

To learn the recommender model with the implicit feedback, the Bayesian Personalized Ranking (BPR) [12] is utilized to ensure that the user preference score of the positive feedback (u, r_0^+, o) is larger than of the negative triple (u, r_0^-, e_f^o) or (e_f^u, r_0^-, o). The model is trained by jointly optimizing the user- and the outfit-aware BPR ranking loss, represented as ℓ_u and ℓ_o, respectively.

$$\ell_u = \sum_{\Omega} \log\left(1 + \exp(-(y_{(u,o,v_o)} - y_{(e_f^u,o,v_o)}))\right), \quad (6)$$

and the proposed outfit-aware ranking loss ℓ_u is denoted as below:

$$\ell_o = \sum_{\Omega} \log\left(1 + \exp(-(y_{(u,o,v_o)} - y_{(u,e_f^o,v_f^o)}))\right). \quad (7)$$

Here v_o and v_f^o represent the activated rule node for outfit o and e_f^o, respectively. And the final overall objective loss for the recommender is shown in the following:

$$L_R = \lambda_1 \ell_u + \lambda_2 \ell_o + \gamma ||\theta_R||_2^2, \quad (8)$$

where λ_u and λ_o denote trade-off parameters. The last term is the $L2$ regularizer loss, which imposes the constraint on the recommender parameters θ_R to prevent overfitting.

3 EXPERIMENTS

3.1 Experimental Settings

3.1.1 Dataset Description. Two publicly available datasets from different domains are utilized for evaluating the performance of RuleKG: Alibaba-iFashion dataset [3] and Last-FM released by [21]. Since our focus is on the top+bottom outfit recommendation, we

Table 1: Performance evaluation with state-of-the-art baselines on two benchmark datasets.

Method	Alibaba-iFashion		Last-FM	
	recall@20	ndcg@20	recall@20	ndcg@20
BPRMF	0.0809	0.0468	0.0687	0.0584
CKE	0.0819	0.0471	0.0732	0.0630
KGAT	0.1033	0.0589	0.0873	0.0744
KGPolicy	0.1189	0.0717	0.0957	0.0837
RuleKG	**0.1284**	**0.0766**	**0.1010**	**0.0926**

Table 2: Performance with variants of the framework by removing rule and candidate selection modules.

Model Variant	Alibaba-iFashion	
	recall@20	ndcg@20
RuleKG	**0.1284**	**0.0771**
w/o RULE	0.1263	0.0756
w/o FCS	0.0941	0.0551
w/o UA	0.0946	0.0552
w/o OA	0.0779	0.0443

conduct data preprocessing to filter out irrelevant items (*e.g.,* shoes, accessories) based on the relations (*i.e.,* item categories) provided by knowledge graph triples and outfits' textual descriptions. The final dataset consists of 24 962 outfits with only tops and bottoms retained. Among 114,728 users, about 90% of them have less than 20 interactions.

3.1.2 Implementation Details. The proposed RuleKG is implemented in Pytorch. The embedding size d is set as 64 in the collaborative knowledge graph (*i.e.,* relations, user, item and entity nodes) during representation learning. For the neighbor sampling, the size of the adjacency map for each node is set as 128 to involve sufficient neighbors for candidate selection. Adam optimizer [8] is leveraged for model update and the batch-size is fixed to 1024.

3.1.3 Baselines. We compare the proposed RuleKG with the following traditional negative sampling methods and knowledge-graph based approaches.

- **BPRMF** [12]: the negatives are randomly sampled with uniform probability with the classical matrix factorization (MF) model as the recommender.
- **CKE** [29] integrates the latent representation learnt from collaborative filtering algorithm and structural representations from TransR.
- **KGAT** [20] proposes an attention-based GCN and models the high-order connectivity among entities.
- **KGPolicy** [22] proposes to discover the high-quality negatives over the knowledge graph using the structure neighborhood information.

3.2 Performance Comparison and Ablation Study

Table 1 shows the comparison results and we can find that the proposed RuleKG outperformed the best performing state-of-the-art KGPolicy [22] by a large margin. It demonstrates the effectiveness of the proposed rule-guided recommender and the negative sampler in obtaining the high-quality negative instances. And the knowledge embedding method (i.e., CKE) achieves much inferior results than GNN-based emthod, which proves that the capability of graph-based neural networks in information propagation.

To demonstrate the significance of the key component in the proposed framework, we compare the variants of RuleKG by removing the core module: constructed fashion rules (w/o RULE), final candidate selection (w/o FCS), outfit-aware negatives (w/o OA) and user-aware negatives (w/o UA). The results are shown in Table 2. From the results, we can observe that the user-aware negatives improve the performance substantially. The final candidate selection module with sampling on two different sets also enhances the performance.

3.3 Case Study

We utilize the domain knowledge embedded in the knowledge graph and category-aware rich semantics provided by the constructed rules. Fig. 3 qualitatively demonstrates the importance of such auxiliary information. We randomly choose a user, the exemplar outfits compositions shown indicate that the target user has a preference over outfits with sweater + skirts. From the top-3 ranked results recommended by our model, we can find that it accurately captures the user's taste for such category combinations.

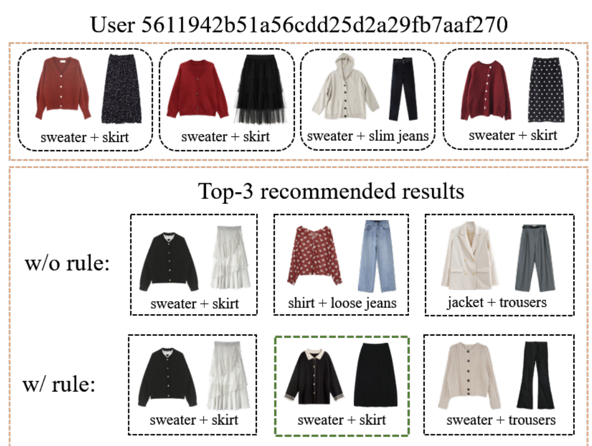

Fig. 3: Visualization of top-3 recommendation results (w/ and w/o rule). Best viewed in color.

4 CONCLUSION AND FUTURE WORK

In this paper, we present a novel rule-guided recommender with the knowledge-aware sampler to perform personalized outfit recommendation. A novel user preference score function is designed, capturing the user's taste toward category combinations in both the implicit and explicit manner. Experiments demonstrate the effectiveness of the proposed framework in modeling the user's preference and fully utilizing the rich facts provided by knowledge graph. In the future, we aim to incorporate user reviews of the outfits to develop explainable knowledge-aware recommendation models and generalize it to more real-world scenarios.

REFERENCES

[1] Antoine Bordes, Nicolas Usunier, Alberto Garcia-Duran, Jason Weston, and Oksana Yakhnenko. 2013. Translating embeddings for modeling multi-relational data. *Advances in neural information processing systems* 26 (2013).

[2] Liwei Cai and William Yang Wang. 2017. Kbgan: Adversarial learning for knowledge graph embeddings. *arXiv preprint arXiv:1711.04071* (2017).

[3] Wen Chen, Pipei Huang, Jiaming Xu, Xin Guo, Cheng Guo, Fei Sun, Chao Li, Andreas Pfadler, Huan Zhao, and Binqiang Zhao. 2019. POG: personalized outfit generation for fashion recommendation at Alibaba iFashion. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 2662–2670.

[4] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).

[5] Qingyu Guo, Fuzhen Zhuang, Chuan Qin, Hengshu Zhu, Xing Xie, Hui Xiong, and Qing He. 2020. A survey on knowledge graph-based recommender systems. *IEEE Transactions on Knowledge and Data Engineering* (2020).

[6] Jin Huang, Wayne Xin Zhao, Hongjian Dou, Ji-Rong Wen, and Edward Y Chang. 2018. Improving sequential recommendation with knowledge-enhanced memory networks. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 505–514.

[7] Zhenya Huang, Qi Liu, Chengxiang Zhai, Yu Yin, Enhong Chen, Weibo Gao, and Guoping Hu. 2019. Exploring multi-objective exercise recommendations in online education systems. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 1261–1270.

[8] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[9] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 42, 8 (2009), 30–37.

[10] Chuan Qin, Hengshu Zhu, Chen Zhu, Tong Xu, Fuzhen Zhuang, Chao Ma, Jingshuai Zhang, and Hui Xiong. 2019. Duerquiz: A personalized question recommender system for intelligent job interview. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2165–2173.

[11] Steffen Rendle and Christoph Freudenthaler. 2014. Improving pairwise learning for item recommendation from implicit feedback. In *Proceedings of the 7th ACM international conference on Web search and data mining*. 273–282.

[12] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618* (2012).

[13] Sebastian Ruder. 2016. An overview of gradient descent optimization algorithms. *arXiv preprint arXiv:1609.04747* (2016).

[14] 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.

[15] Richard S Sutton and Andrew G Barto. 2018. *Reinforcement learning: An introduction*. MIT press.

[16] Richard S Sutton, David McAllester, Satinder Singh, and Yishay Mansour. 1999. Policy gradient methods for reinforcement learning with function approximation. *Advances in neural information processing systems* 12 (1999).

[17] Hongwei Wang, Fuzheng Zhang, Xing Xie, and Minyi Guo. 2018. DKN: Deep knowledge-aware network for news recommendation. In *Proceedings of the 2018 world wide web conference*. 1835–1844.

[18] Jun Wang, Lantao Yu, Weinan Zhang, Yu Gong, Yinghui Xu, Benyou Wang, Peng Zhang, and Dell Zhang. 2017. Irgan: A minimax game for unifying generative and discriminative information retrieval models. In *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*. 515–524.

[19] Pengfei Wang, Yu Fan, Long Xia, Wayne Xin Zhao, ShaoZhang Niu, and Jimmy Huang. 2020. KERL: A knowledge-guided reinforcement learning model for sequential recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*. 209–218.

[20] Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. 2019. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 950–958.

[21] Xiang Wang, Tinglin Huang, Dingxian Wang, Yancheng Yuan, Zhenguang Liu, Xiangnan He, and Tat-Seng Chua. 2021. Learning Intents behind Interactions with Knowledge Graph for Recommendation. In *Proceedings of the Web Conference 2021*. 878–887.

[22] Xiang Wang, Yaokun Xu, Xiangnan He, Yixin Cao, Meng Wang, and Tat-Seng Chua. 2020. Reinforced negative sampling over knowledge graph for recommendation. In *Proceedings of The Web Conference 2020*. 99–109.

[23] Zhen Wang, Jianwen Zhang, Jianlin Feng, and Zheng Chen. 2014. Knowledge graph embedding by translating on hyperplanes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 28.

[24] Yikun Xian, Zuohui Fu, Shan Muthukrishnan, Gerard De Melo, and Yongfeng Zhang. 2019. Reinforcement knowledge graph reasoning for explainable recommendation. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval*. 285–294.

[25] Fajie Yuan, Guibing Guo, Joemon M Jose, Long Chen, Haitao Yu, and Weinan Zhang. 2016. Lambdafm: learning optimal ranking with factorization machines using lambda surrogates. In *Proceedings of the 25th ACM international on conference on information and knowledge management*. 227–236.

[26] Huijing Zhan and Jie Lin. 2021. PAN: Personalized Attention Network For Outfit Recommendation. In *2021 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2663–2667.

[27] Huijing Zhan, Jie Lin, Kenan Emir Ak, Boxin Shi, Ling-Yu Duan, and Alex C Kot. 2021. A3-FKG: Attentive Attribute-Aware Fashion Knowledge Graph for Outfit Preference Prediction. *IEEE Transactions on Multimedia* (2021).

[28] Weinan Zhang, Tianqi Chen, Jun Wang, and Yong Yu. 2013. Optimizing top-n collaborative filtering via dynamic negative item sampling. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. 785–788.

[29] Yongfeng Zhang, Qingyao Ai, Xu Chen, and Pengfei Wang. 2018. Learning over knowledge-base embeddings for recommendation. *arXiv preprint arXiv:1803.06540* (2018).

[30] Yongqi Zhang, Quanming Yao, Yingxia Shao, and Lei Chen. 2019. NSCaching: simple and efficient negative sampling for knowledge graph embedding. In *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. IEEE, 614–625.

[31] Pengpeng Zhao, Anjing Luo, Yanchi Liu, Fuzhen Zhuang, Jiajie Xu, Zhixu Li, Victor S Sheng, and Xiaofang Zhou. 2020. Where to go next: A spatio-temporal gated network for next poi recommendation. *IEEE Transactions on Knowledge and Data Engineering* (2020).

[32] Sijin Zhou, Xinyi Dai, Haokun Chen, Weinan Zhang, Kan Ren, Ruiming Tang, Xiuqiang He, and Yong Yu. 2020. Interactive recommender system via knowledge graph-enhanced reinforcement learning. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*. 179–188.