An Industrial Framework for Personalized Serendipitous Recommendation in E-commerce

ZONGYI WANG, YANYAN ZOU*, LINFANG HOU, ANYU DAI, NAN QIAO, LUOBAO ZOU, MIAN MA, ZHUOYE DING, and SULONG XU, JD.com, China

Classical recommendation methods typically face the filter bubble problem where users likely receive recommendations of their familiar items, making them bored and dissatisfied. To alleviate such an issue, this applied paper introduces a novel framework for personalized serendipitous recommendation in an e-commerce platform (i.e., JD.com), which allows to present user unexpected and satisfying items deviating from user's prior behaviors, considering both accuracy and novelty. To achieve such a goal, it is crucial yet challenging to recognize when a user is willing to receive serendipitous items and how many novel items are expected. To address above two challenges, a two-stage framework is designed. Firstly, a DNN-based scorer is deployed to quantify the novelty degree of a product category based on user behavior history. Then, we resort to a potential outcome framework to decide the optimal timing to recommend a user serendipitous items and the novelty degree of the recommendation. Online A/B test on the e-commerce recommender platform in JD.com demonstrates that our model achieves significant gains on various metrics, 0.54% relative increase of impressive depth, 0.8% of average user click count, 3.23% and 1.38% of number of novel impressive and clicked items individually.

CCS Concepts: • **Information systems** → **Recommender systems**.

Additional Key Words and Phrases: Recommender System, Personalization, Serendipity

ACM Reference Format:
Zongyi Wang, Yanyan Zou, Linfang Hou, Anyu Dai, Nan Qiao, Luobao Zou, Mian Ma, Zhuoye Ding, and Sulong Xu. 2023. An Industrial Framework for Personalized Serendipitous Recommendation in E-commerce. In *Seventeenth ACM Conference on Recommender Systems (RecSys '23), September 18–22, 2023, Singapore, Singapore.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3604915.3610234

1 INTRODUCTION

Nowadays, e-commerce recommender systems play a vital role in recommending users the best and most appropriate content retrieved from billions of product candidates. To achieve this goal, collaborative filtering (i.e., CF) based [17] and click-through-rate (i.e., CTR) prediction based [19] models have been successfully deployed in industry for decades. However, such classical algorithms might suffer the problems of over-specialization [1], filter bubbles [12, 13] and user boredom [3, 4], due to the feedback loop issue of existing recommender systems.

One way to alleviate user boredom is serendipitous recommendation [11], which aims to improve deviations of recommended items from user expectations and thus captures the concept of user surprise and allows recommender systems to break from the feedback loops. Several literature [2, 3, 8, 9] has shown that introducing serendipity into recommendations leads to significant increase of user satisfaction. Unfortunately, very few of them have been successfully applied to real-world applications due to their complexity and lack of scalability. Moreover, the personalization of serendipitous recommendation is crucial to consider the heterogeneous user preferences to unexpectedness during deployment in industrial applications. For example, some people tend to be "variety-seekers" [10] and are more willing

*Corresponding to zouyanyan6@jd.com

to click on novel items, while others prefer to stay in their own comfort zones with familiar recommendations. Even, the same person might have different perceptions of novelty under different contexts.

Recently, Li et al. [7] designs a personalized unexpectedness module that is incorporated into a session-based ranking algorithm to capture user's perception of unexpectedness for video recommendations. However, we argue that the situation gets more tricky for e-commerce feed recommender systems, where a user could continuously slide down and quickly browse a long sequence of items in a short time. The user's perception of serendipity might change frequently even within a single session. For example, a user might want to check out some items highly related to his/her prior interests at the beginning of the current recommendation sequence, while the preferences to novel items could get stronger after browsing some familiar items. The above observations motivate us to carefully capture the change of user's perceptions of serendipity in a more fine-grained manner and to propose a novel framework of serendipitous recommendation for a real-world recommender system in the large-scale e-commerce platform, JD.com. The proposed framework consists of two modules working as two stages. In the first stage, a category novelty scorer is designed to score how novel a product category is to a user, i.e., the serendipity score of the category. During the second stage, a novelty intention detector is employed to first quantify the novelty degree of a recommendation list. Subsequently, we borrow the potential outcome framework and design an uplift-based model to decide when a user is expecting serendipitous recommendation and to what extent the novelty is desired. Extensive offline experiments on real-world industrial data illustrate that the proposed uplift-based model indeed captures user's willingness of surprising items. We also conduct an online A/B test which demonstrates our framework achieves significant improvements over the default production settings with regard to both accuracy and user satisfaction on novelty.

2 METHOD

2.1 Category Novelty Scorer

It is worthy noting that the goal is to discover novel product categories, rather than items, for a user. Assume we have a user u and a set of item categories $c \in C$. For such a user, we have a sequence of user historical behaviors $E_u = (c(u)_1, c(u)_2, \cdots, c(u)_n)$, sorted by time of the occurrence, where $c(u)_i$ records the i_{th} item category interacted by the user, and n is the number of user actions we considered. Given historical interactions E_u, we aim to measure how novel a new category c is to the user u. We define the novelty of a certain category as the distance between the user's known interests and itself in the feature space. In other words, we want to measure how much a user is interested in a category according to observed user behaviors. Since user's attention can be diverse over a time period, we adopt multi-interest extraction framework to capture user's multiple interests for user behavior sequence E_u. To be specific, similar to the recent work [6], we resort to the dynamic routing method [16] to group actions from user's historical behaviors E_u into several clusters M_u, each of which represents one particular aspect of user interests. The observed preference score of category c for user u is defined as:

$$\text{prefer}_u^c = \max_{m_u^i \in M_u} f(m_u^i, c)$$

where m_u^i represents the i_{th} interest cluster, and f is the cosine similarity function. Thus, prefer_u^c implies to what extent user u is familiar with and also interested in category c. Then the novelty of category c for user u, i.e., the serendipity score, is calculated as:

$$S_u^c = 1 - \text{prefer}_u^c$$

A larger value of S_u^c means the category c is more novel to user u.

2.2 Novelty Intention Detector

It is reasonable to merely recommend novel items to someone who might be interested in novel recommendation sometime. To satisfy personalized preferences to serendipity, we argue there are two crucial challenges: 1) at what time a user is willing to see more novel items; 2) how many novel items are expected.

To address the first challenge, we model the demand of novelty as the user engagements of serendipitous recommendation, which is represented by the number of clicks on the whole recommendation list (i.e., including both novel and familiar items). More clicks stand for better user satisfaction and experience. Moreover, to capture the change of user preference to serendipity within a session, for every N impressive items (e.g., a single request), we make a prediction to check if user's preference to novelty is changed. For the second challenge, in our scenario, for simplicity, we first group the novelty of categories into four grades according to the serendipity score S_u^c:

- Group-1: familiar category where $S_u^c \leq 0.5$; Group-2: slightly novel category where $0.5 < S_u^c \leq 0.7$;
- Group-3: medium novel category where $0.7 < S_u^c \leq 0.9$; Group-4: extremely novel category where $S_u^c > 0.9$.

which fits our scenario according to extensive preliminary experiments. Intuitively, we consider each N continuous items as a recommendation list (in other words, a single request), where item categories belong to one of the above four category groups. Therefore, we denote the novelty degree of a recommendation list as $D = \{d_0, d_1, d_2, d_3\}$, where d_i stands for the ratio of the count of item categories from Group-i and the recommendation list length N and the summation of d_i is always equal to 1. For example, $D = \{0.25, 0.25, 0.25, 0.25\}$ stands for a special case where each novelty group accounts for the same percentage. Now, our goal is to investigate at what time a user would prefer to serendipitous items and which novelty degree of a recommendation list is desired. To solve such a problem, we resort to the uplift modeling which is a predictive modeling technique and directly models the incremental impact of a treatment on individual's behaviors. In our scenario, the incremental impact is the number of increased clicks made by users who are presented with serendipitous recommendation, i.e. a treatment. We denote that recommendation lists with different novelty degrees stand for different treatments. To measure the impact on user clicks of different novelty degrees, we borrow the potential outcome framework [14, 15].

Assume $Y_u(t)$ represents the number of clicks of user u under treatment condition t. For simplicity, we consider a two-arm trial where $t = 0$ refers to the control condition, i.e., the default product settings without serendipitous recommendation, and $t = D$ stands for the treatment condition where the novelty degree of serendipitous recommendation is D. Then, the average causal effect of the treatment is calculated as

$$\tau_u = E[Y_u(D) - Y_u(0)]$$

which is actually the click variations of user u under the treatment condition. For a multi-arm scenario, we would like to find a treatment with novelty degree of D that maximizes τ_u.

To predict the click variation, we implement a neural network based novelty intention detector model with two-layer preceptrons and a softmax layer. A typical uplift-based model is learned to directly predict the value of click variation [18]. Nevertheless, we train the model to predict under which group, control or a certain treatment group, the click variation is the maximum, which is essentially a classification task. We empirically observed that prediction classification labels is more suitable for our scenario due to the data-sparsity issue. Take the two-arm case as an example, the task is converted into a binary classification problem. Samples with $\tau_u > \delta$ are considered as positive examples and others as negative ones. Then the model can be trained with the classical cross-entropy loss function.

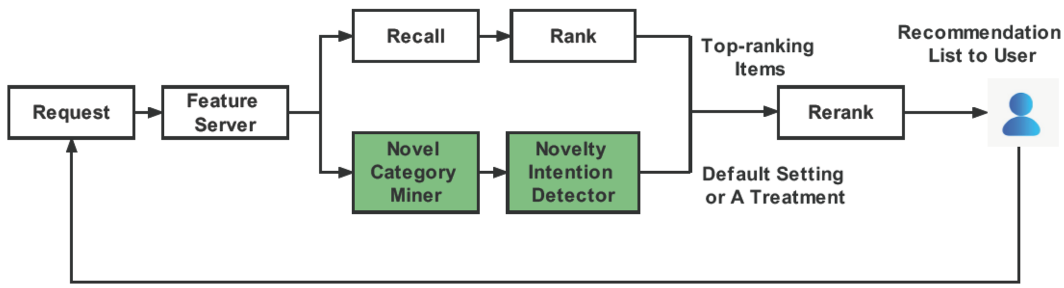

Fig. 1. Deployment of the proposed framework (highlighted in green) on the JD.com recommender system.

Model	AUC	AUUC
Single-Learner	0.544	0.121
Two-Learner	0.571	0.124
Cross-Learner	0.541	0.131
Ours	**0.663**	**0.226**

Table 1. Performance of different model settings.

2.3 Serendipitous Recommendation List

Once we capture user' preference to serendipitous recommendation with a novelty degree D, we reorder the ranking results based on the KL divergence score:

$$KL(p||q) = \sum p(x) \log \frac{p(x)}{q(x)}$$

where $p(x)$ is the distribution of D, and $q(x)$ is the category distribution of the serendipitous recommendation list.

3 DEPLOYMENT

We have successfully deployed the proposed method on the JD.com recommender system, one of the popular large-scale e-commerce platforms in China. Figure 1 demonstrates the overview architecture of the deployed workflow. To save serving latency, the proposed framework works with the existing recommender system, i.e., Recall and Ranking stages, in parallel. Since the user preference to novelty varies over time, the model is regularly updated online.

4 EXPERIMENT

Since there are available public data for our scenario, we conduct experiments on real-world industrial data collected from the "Recommended to You" of JD.com, with around 3 million of samples for training and 870 thousand of records for validation. The features including user profile (e.g., click preference, purchase level), user recent behavior (e.g., recent 7 days), as well as context features (e.g., impressive and click items for the current session, current request page index, and request time).

Offline Experiments. It is common that there is gap between offline and online performance of recommendation strategies. On the other hand, the novelty intention detector model captures user preferences to serendipity which significant matter the performance of serendipitous recommendation. Thus, we verify the accuracy of the proposed novelty intention detector model. For simplicity, we consider two cases: the default production setting and a serendipitous

Model	Impression Depth	UCTR	#Impression Novel Category	#Clicked Novel Category
PURS (baseline)	0.00	0.00	0.00	0.00
+ Category Novelty Scorer	+1.16%*	-0.80%*	+10.57%*	+3.98%*
Ours	+0.54%*	+0.80%*	+3.23%*	+1.38%*

Table 2. Improvements of different system settings through online A/B test. $*$ refers to significant difference with $p < 0.01$.

recommendation with novelty degree of $D = \{0.30, 0.30, 0.20, 0.20\}$, which is essentially a binary classification task. We also make comparisons with several variants of uplift-based models, including Single-Learner, Two-Learner, Cross-Learner (Refer [5] for details), which predicts the variation of user click counts and can be directly converted to predicted labels. As listed in Table 1, we report two offline metrics: AUC for classification prediction and AUUC for the cumulative gains of selected treatments. Our proposed method achieves the best performance on both metrics.

Online A/B Test. To directly evaluate the payoff of our proposed framework, a standard online A/B test is conducted at "Recommended to You" in JD.com from 2023-01-06 to 2023-02-06. As listed in Table 2, we compare our method with the default production settings which is a implementation of PURS [7]. To evaluate the effects of each module, we also consider two settings: 1) only category novelty scorer, which is non-personalized setting where all traffic share the same novelty degree, denoted as "+Category Novelty Scorer". 2) with both Category Novelty Scorer and Novelty Intention Detector modules, which is a setting considering personalized needs of serendipitous recommendation, denoted as "Ours". To be specific, we measure the performance using standard business metrics: Impression Depth, average impression through one session and UCTR (a variant of Click-Through-Rate, the average number of clicks over impressive users). We also measure the novelty of the recommended items using two statistics metrics: average number of Impressive Novel items and average number of Clicked Novel items. Compared to the default production setting, the non-personalized setting (denoted as "+Category Novelty Scorer") largely improves the impressive depth and novelty metrics, while the user satisfaction (i.e., UCTR) drops a lot. Differently, our method achieves significant improvements on both business and novelty metrics ($p < 0.01$), which demonstrate the effectiveness of our personalized framework for serendipitous recommendation.

5 CONCLUSION

This work proposes a novel framework for personalized serendipitous recommendation, consisting of two modules, i.e., category novelty scorer and novelty intention detector. The category novelty scorer aims to rank how novel a category is to a user, after which the novelty intention detector is employed to recognize at what time a user is prone to surprises and what kind of novel items. Online experiments illustrate the superiority of our proposed framework. In the future, we plan to involve multiple optimization goals (e.g., conversion rate) into the uplift modeling.

Presenter Biography. Yanyan Zou is an applied scientist in Recommendation Platform at JD.com since 2020, launching cutting-edge AI models into practical productions. Her research interests primarily lie in the areas of natural language processing and recommendation, with around 20 papers published in top-tier conferences (e.g., ACL, EMNLP, AAAI). Singapore University of Technology and Design in 2020.

REFERENCES

[1] Panagiotis Adamopoulos and Alexander Tuzhilin. 2014. On Over-Specialization and Concentration Bias of Recommendations: Probabilistic Neighborhood Selection in Collaborative Filtering Systems. In *8th ACM Conference on Recommender Systems: RecSys 2014*.

[2] Panagiotis Adamopoulos and Alexander Tuzhilin. 2014. On unexpectedness in recommender systems: Or how to better expect the unexpected. *ACM Transactions on Intelligent Systems and Technology (TIST)* 5, 4 (2014), 1–32.

[3] Komal Kapoor, Vikas Kumar, Loren Terveen, Joseph A Konstan, and Paul Schrater. 2015. " I like to explore sometimes" Adapting to Dynamic User Novelty Preferences. In *Proceedings of the 9th ACM Conference on Recommender Systems*. 19–26.

[4] Komal Kapoor, Karthik Subbian, Jaideep Srivastava, and Paul Schrater. 2015. Just in time recommendations: Modeling the dynamics of boredom in activity streams. In *Proceedings of the eighth ACM international conference on web search and data mining*. 233–242.

[5] Sören R Künzel, Jasjeet S Sekhon, Peter J Bickel, and Bin Yu. 2019. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the national academy of sciences* 116, 10 (2019), 4156–4165.

[6] Chao Li, Zhiyuan Liu, Mengmeng Wu, Yuchi Xu, Huan Zhao, Pipei Huang, Guoliang Kang, Qiwei Chen, Wei Li, and Dik Lun Lee. 2019. Multi-interest network with dynamic routing for recommendation at Tmall. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 2615–2623.

[7] Pan Li, Maofei Que, Zhichao Jiang, Yao Hu, and Alexander Tuzhilin. 2020. PURS: personalized unexpected recommender system for improving user satisfaction. In *Fourteenth ACM Conference on Recommender Systems*. 279–288.

[8] Pan Li and Alexander Tuzhilin. 2019. Latent Modeling of Unexpectedness for Recommendations.

[9] Qiuxia Lu, Tianqi Chen, Weinan Zhang, Diyi Yang, and Yong Yu. 2012. Serendipitous personalized ranking for top-n recommendation. In *2012 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology*, Vol. 1. 258–265.

[10] Leigh McAlister and Edgar Pessemier. 1982. Variety seeking behavior: An interdisciplinary review. *Journal of Consumer research* 9, 3 (1982), 311–322.

[11] Tomoko Murakami, Koichiro Mori, and Ryohei Orihara. 2007. Metrics for evaluating the serendipity of recommendation lists. In *Proceedings of the 2007 conference on New frontiers in artificial intelligence*. 40–46.

[12] Tien T Nguyen, Pik-Mai Hui, F Maxwell Harper, Loren Terveen, and Joseph A Konstan. 2014. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *Proceedings of the 23rd international conference on World wide web*.

[13] Eli Pariser. 2012. The Filter Bubble: How the New Personalized Web Is Changing What We Read and How We Think.

[14] DONALD B RUBIN. 1974. ESTIMATING CAUSAL EFFECTS OF TREATMENTS IN RANDOMIZED AND NONRANDOMIZED STUDIES1. *Journal of Educational Psychology* 66, 5 (1974), 688–701.

[15] Donald B Rubin. 2005. Causal inference using potential outcomes: Design, modeling, decisions. *J. Amer. Statist. Assoc.* 100, 469 (2005), 322–331.

[16] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. 2017. Dynamic routing between capsules. *Advances in neural information processing systems* 30 (2017).

[17] J Ben Schafer, Dan Frankowski, Jon Herlocker, and Shilad Sen. 2007. Collaborative filtering recommender systems. In *The adaptive web*. Springer, 291–324.

[18] Zhenyu Zhao and Totte Harinen. 2019. Uplift modeling for multiple treatments with cost optimization. In *2019 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*. IEEE, 422–431.

[19] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1059–1068.