

Personalized Re-ranking for Improving Diversity in Live Recommender Systems

Yichao Wang, Xiangyu Zhang, Zhirong Liu, Zhenhua Dong, Xinhua Feng, Ruiming Tang, Xiuqiang He

Noah's Ark Lab, Huawei, Shenzhen, China

{wangyichao5,zhangxiangyu6,liuzhirong,dongzhenhua,fengxinhua1,tangruiming,hexiuqiang1}@huawei.com

ABSTRACT

Users of industrial recommender systems are normally suggested a list of items at one time. Ideally, such list-wise recommendation should provide diverse and relevant options to the users. However, in practice, list-wise recommendation is implemented as top- N recommendation. Top- N recommendation selects the first N items from candidates to display. The list is generated by a ranking function, which is learned from labeled data to optimize accuracy. However, top- N recommendation may lead to sub-optimal, as it focuses on accuracy of each individual item independently and overlooks mutual influence between items. Therefore, we propose a personalized re-ranking model for improving diversity of the recommendation list in real recommender systems. The proposed re-ranking model can be easily deployed as a follow-up component after any existing ranking function. The re-ranking model improves the diversity by employing *personalized* Determinantal Point Process (DPP). DPP has been applied in some recommender systems to improve the diversity and increase the user engagement. However, DPP does not take into account the fact that users may have individual propensities to the diversity. To overcome such limitation, our re-ranking model proposes a *personalized* DPP to model the trade-off between accuracy and diversity for each individual user. We implement and deploy the personalized DPP model on a large scale industrial recommender system. Experimental results on both offline and online demonstrate the efficiency of our proposed re-ranking model.

CCS CONCEPTS

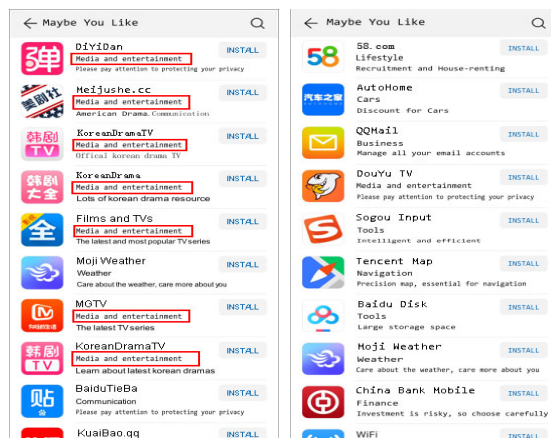
• Information systems → Recommender systems.

KEYWORDS

Diversity, Re-ranking, Determinantal Point Processes, Recommender System

1 INTRODUCTION AND BACKGROUND

Recommender systems are powerful information filters for guiding users to find their interested items from gigantic and rapidly expanding pool of candidates, and they have taken more and more scenarios in our lives [1] [2] [3] [4]. Users in industrial recommender systems are normally recommended a list of items at one time. Ideally, such list-wise recommendation should provide diverse and relevant options to the users. Due to efficiency issue, many industrial recommender systems implement list-wise recommendation as Top- N recommendation, which selects the first N items from an ordered list. The ordered list is generated by a ranking function, which is learned from labeled data to optimize



(a) a top- N recommendation list without considering diversity (b) a top- N recommendation list considering diversity

Figure 1: Comparison of ranking list considering diversity or not.

accuracy and produces a ranking score for each individual item. Such top- N recommendation focuses on relevance of each individual item independently and overlooks mutual influence between items. As observed in [5], recommending a list of items by such a method lead to sub-optimal performance of recommender systems, due to the following two aspects. On the one hand, ranking by relevance is likely to select multiple similar items in the list. However, it is highly possible that at most one of such similar items is needed by a user, while the others are redundant and waste the chance being displayed to the user. We take a real-world example from a mainstream App Store. As shown in Figure 1(a), a top- N recommendation list consists of multiple apps in the category of social community and video, because the ranking function learns that the user likes to chat with people and watch video. However, recommending multiple apps with the same functionality results wasting displayed chances and also degrading user experience. A more reasonable recommendation list should take diversity into account, as presented in Figure 1(b). On the other hand, focusing on relevance of items may lead to information isolation for the users [6], which results in leaving fewer opportunities for exploring new items [7]. To address this problem, diversity [8–11] has been imposed as a complement of accuracy, to model the mutual influence between items and therefore improve the effectiveness of list-wise recommendations.




Figure 2: The distribution of user’s entropy over her download history collected from App Store with user consent

A multitude of approaches have been presented to generate the diverse recommendations. Industrial recommender systems are a very complicated framework so that we have to consider the ease of implementation and the risk of launching online when we deploy new models or components online. Therefore, considering diversity in recommendation, we aim to propose a component that is compatible with all the existing components, instead of replacing some of them. That is to say, we target “diversity” as a re-ranking model, which can be easily deployed as a follow-up component after any existing ranking function. Some existing models exist, which treat “diversity” as a re-ranking model, assuming a ranking of items is available. For example, Maximal Marginal Relevance (MMR) method in [12] selects one item at a time from a ranking list, which considers both of the relevance and the pair-wise similarity. Probabilistic models based on Determinantal Point Process (DPP) in [13, 14] consider the list-wise similarity among items through a kernel matrix, which consists of relevance and pair-wise similarity. Compared with the MMR-based model, DPP-based model can improve the diversity more efficiently without degrading the accuracy [13]. However, we observe that an unrealistic assumption is made in DPP-based model: it is assumed that different users have the same propensity to the degree of diversity. We find some evidence from both literature and real-world data to call for different propensity to diversity with different individuals.

Analyzing user behaviors from the same App Store with user consent, the result is presented in Figure 2. The figure presents the distribution of user’s entropy over her download history. The x -axis represent the entropy value of the app category in her download history, while the y -axis represents the normalized population of this entropy value in the whole population of 10,000 users. It suggests that users’ taste varies significantly. Users with large entropy values have a variety of interests over different categories of apps, while users with small entropy values focus on few categories of apps. It can be implied from this fact that different individuals have different propensity to diversity.

After seeking answer from industrial applications, we check the viewpoint of literature. [15] demonstrates that the personality traits of users significantly correlate with their behaviors in recommendation system. They take into account each users’ personality trait based on a large scale user survey in [16, 17], and present that different users have different propensity to the degree of diversity. Specifically, the users with narrow taste of items way

expect more similar items in the recommendation list, while the users who have a variety of interests may expect more diverse items. Moreover, some researchers have proposed several methods to utilize the users’ behaviors for building the personalized diversified recommendation. In [18], the proposed algorithms focus on the users’ propensity to diversity based on the different attributes of items, and re-rank the recommendation list by MMR methods. A pre-filtering approach proposed in [19] clusters the users into four groups according to each user’s inclination to the diversity, and then apply the user based collaborative filtering algorithm for each group. These methods demonstrate the effectiveness of the personalized diversity with the offline experiment on some public datasets.

However, [18, 19] personalize the diversity propensity on four user clusters, instead of on individuals, where each hyper-parameter in individual user cluster needs to be grid searched. As indicated from Figure 2, it is more reasonable to personalize the diversity propensity on individual users, as users’ propensity to diversity varies significantly. However, it is impossible to extend [18, 19] straightforwardly, personalizing diversity propensity from the granularity of user clusters to a much finer granularity of individual users, as searching for hyper-parameter for each user is impractical. Note that the number of users in an industry recommender system is normally tens or hundreds of millions.

In this paper, we propose a personalized DPP model to improve the diversity of recommendation list, where the personalized granularity is of individual users. The hyper-parameter for each user is factorized to two factors: one is formulated by information entropy of a user’s interaction history, while the other is commonly shared across all the users and tunable.

We summarize the main contributions of our study:

- We propose a personalized re-ranking model for improving diversity of recommendation list, and it can be easily deployed as a follow-up component after ranking function.
- The re-ranking model employs personalized DPP, where the penalization granularity is on individual users, instead of on user clusters in the literature.
- We conduct the experimental evaluations on an offline benchmark to show the superiority of our proposed re-ranking model.
- We deploy our proposed re-ranking model in a live recommender system and demonstrate the significant improvement for both of diversity and accuracy over baselines in online A/B test.

The rest of the paper is organized as follows: in Section 2, we elaborate our re-ranking model in detail. We present our system architecture in live recommender systems in Section 3. Experimental setting and offline/online results are shown and discussed in Section 4. Finally, we give the conclusion in Section 5.

2 RE-RANKING MODEL

2.1 DPP-based Re-ranking

As studied in [13], DPP-based model is more effect and more efficient than other models such as MMR-based model. Therefore, we choose to investigate how to apply DPP-based re-ranking model in our recommender system. In this section, we present DPP-based

re-ranking model, and discuss its limitation, which motivates our personalized DPP-based re-ranking model in the next section.

We summarize some key results about from [13, 14, 20], for readers to better understand our model. A point process \mathcal{P} on a set of items $\mathcal{M} = \{1, 2, \dots, |\mathcal{M}|\}$ is a probability distribution on the powerset of \mathcal{M} . That is, $\forall Y \subseteq \mathcal{M}$, \mathcal{P} assigns a probability $\mathcal{P}(Y)$, such that $\sum_{Y \subseteq \mathcal{M}} \mathcal{P}(Y) = 1$. It is stated in [20] that, finding the set $\max_{Y:|Y|=k, Y \subseteq \mathcal{M}} \mathcal{P}(Y)$ is a way of selecting a relevant and diverse subset of k items from the whole item set \mathcal{M} . Furthermore, \mathcal{P} can be compactly parameterized by a $\mathcal{M} \times \mathcal{M}$ positive semi-definite kernel matrix L , such that $\mathcal{P}(Y) \propto \det(L_Y)$, where $\det(L)$ is the determinants of matrix L and L_Y is a submatrix of L projected to only those rows and columns in Y . Therefore, find the set $\max_{Y:|Y|=k, Y \subseteq \mathcal{M}} \mathcal{P}(Y)$ is equivalent to finding the set $\max_{Y:|Y|=k, Y \subseteq \mathcal{M}} \det(L_Y)$.

The positive semi-definite kernel matrix L is defined as follows:

$$L_{ii} = q_i^2 \quad (1)$$

$$L_{ij} = \alpha q_i q_j S_{ij} \quad (2)$$

where $q_i (i \in [1, |\mathcal{M}|])$ denotes relevance score of item i generated from the ranking function, S denotes a user-defined similarity matrix among the items, α is the hyper-parameter to trade-off relevance and diversity.

As discussed before, we need to select a set of items Y from the whole item set \mathcal{M} , such that

$$\max_{Y:|Y|=k, Y \subseteq \mathcal{M}} \det(L_Y) \quad (3)$$

It is known as a NP-hard problem [20] with complexity $O(C_{|\mathcal{M}|}^{|Y|})$ to find the optimal set. To make DPP-based re-ranking model applicable in industrial recommender systems, we choose to use an efficient and effective approximation algorithm, Fast Greedy MAP Inference [21], to perform re-ranking in an acceptable latency. Such an approximation algorithm solves this combination optimization problem approximately in $O(|Y|^2|\mathcal{M}|)$. Although theoretic lower bound is not provided in [21], online A/B test is conducted to demonstrate its superiority.

To present it formally, we summarize DPP-based re-ranking model in Algorithm 1. Each round, FastGreedyMAP selects the one item greedily (as shown in Line 3 of Algorithm 1), which is to say, the selected item promotes the determinants of the updated submatrix most. Formally, it selects the item $y = \operatorname{argmax}_{i \in \mathcal{M}} \left(\log \det(L_{Y \cup \{i\}}) - \log \det(L_Y) \right)$.

2.2 Personalized DPP

In DPP, α is a tunable hyper-parameter to balance the trade-off between relevance and diversity. DPP assumes every individual have the same propensity to the degree of diversity, as the same α value is applied when constructing the kernel matrix L , which is shared when performing re-ranking for all users. However, as we discussed in Section 1, different individuals have different propensity to diversity, so that personalization is needed in DPP.

A straightforward way to implement personalization in DPP is setting a unique hyper-parameter α_u for user u . Unfortunately, this approach is not practical, since the number of hyper-parameters

Algorithm 1 DPP-based Re-ranking model

Input: candidate items, \mathcal{M} ; kernel matrix L ; number of required items k ;

Output: re-ranking list, Y ;

- 1: **Initialize:** $Y = \Phi$;
 - 2: **while** $|Y| < k$ and $|\mathcal{M}| > 0$ **do**
 - 3: $y = \text{FastGreedyMAP}(\mathcal{M}, L, k)$
 - 4: $Y = Y \cup y$
 - 5: $\mathcal{M} = \mathcal{M} - y$
 - 6: **end while**
 - 7: **return** Y
-

α_u 's is too large to be tuned individually. In this paper, we present an effect and efficient method to achieve personalized DPP (For short, we refer it as *pDPP*). We factorize user-wise hyper-parameter α_u to two factors as

$$\alpha_u = f_u \times \alpha_0 \quad (4)$$

where α_0 is a tunable and shared hyper-parameter to trade-off relevance and diversity across all the users (which is of the same functionality as α in DPP) and f_u is a user-wise factor representing diversity propensity of user u .

Next, we elaborate the intuition of defining f_u . As explained in a real-world example in Section 1, users' diversity propensity can be reflected by their historical behavior. As one of the possible choices, Shannon entropy over the distribution of different genres¹ of interacted items by the user is utilized, as

$$\mathcal{H}(u) = - \sum_{g \in \mathcal{G}} P(g|u) \log(P(g|u)) \quad (5)$$

where $P(g|u)$ denotes the probability of user u being interested in genre g , namely, one of user u 's interacted items being of genre g . As shown in [18], a user u with higher $\mathcal{H}(u)$ has higher propensity of diversity and vice versa. Due to this intuition, we define f_u as the normalized $\mathcal{H}(u)$. Formally, we propose to use a *parameterized min-max normalization*, as follows:

$$f_u = \frac{\mathcal{H}(u) - \mathcal{H}_{\min} + l}{\mathcal{H}_{\max} - \mathcal{H}_{\min} + l} \quad (l \geq 0) \quad (6)$$

where $\mathcal{H}_{\max} = \max_u \mathcal{H}(u)$ represent the maximal entropy value over all the users and \mathcal{H}_{\min} denote the minimal value. The hyper-parameter l controls the personalization degree of f_u (and therefore α_u). As shown in Figure 3, a larger l value indicates less personalized f_u values among all the users, e.g., when $l \rightarrow \infty$, it can be seen that $f_u = 1$ and pDPP downgrades to DPP. In practice, we choose to use two special cases: when $l = 0$, f_u is the standard *min-max* normalized $\mathcal{H}(u)$; and when $l = \mathcal{H}_{\min}$, f_u is the *max* normalized $\mathcal{H}(u)$.

To summarize, pDPP is a personalized version of DPP without introducing extra hyper-parameters for tuning. Though the formulation is simple, the experiment results in Section 4 demonstrate its effectiveness.

¹Our formulation can be extended easily by including other features of items.




Figure 3: Paraterized min-max normalized entropy with different hyper-parameter l values.

3 SYSTEM IMPLEMENTATION

3.1 Framework Modifications

An overview of a recommender system with pDPP re-ranking model is shown in Figure 4. We first present the modules without considering the re-ranking model (which is surrounded in green box) and then illustrate how to adapt these modules with pDPP.

The architecture of a recommender system consists of three modules. (1) *Offline training module* processes user-item interaction data, extracts features (user features, item features and context features), trains model and uploads the model. (2) *Online prediction module* receives users' request and returns a list of items. There are usually two steps in this module, namely *retrieval* and *ranking*. Since there may be over millions of items, it is impossible to score every item within a required latency (often within tens of millisecond). The retrieval step returns a short list of items (often hundreds or thousands) of items that is suitable for the user under such context. After reducing the size of candidates, the ranking step computes relevance scores for individual items using the offline trained model. (3) *Nearline updating module*, which updates user features, item features and even the offline trained models with real-time interaction data.

Our proposed pDPP re-ranking model can be integrated into the above architecture easily. Next, we will elaborate how to adapt the three modules in the framework, to deploy this re-ranking model.

- In offline training module, α initializer computes α_u value for individual user u and uploads such values to online *Indexer*.
- In online prediction module, given the relevance scores of candidate items computed by any ranking function and the personalized α_u value from online *Indexer*, pDPP re-ranking model generates the final recommendation list, considering both relevance and diversity.
- In nearline updating module, personalized α_u values are updated based on the real-time user-item interaction data, and the updated α_u values are sent to online *Indexer*.

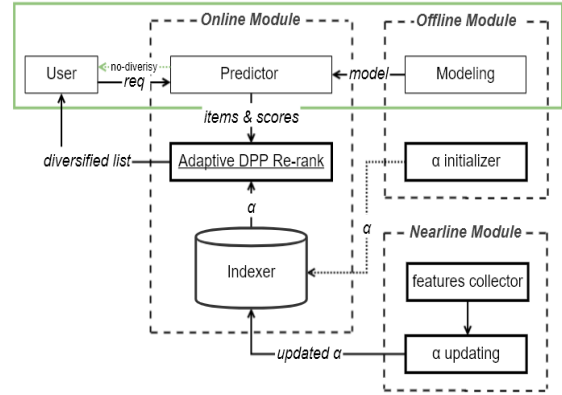


Figure 4: The architecture overview of a recommender system with pDPP re-ranking model. The pDPP re-ranking model can be integrated with the other parts easily, and is compatible with any advanced ranking function.

Developing accurate ranking function is an essential research topic and attracts many researchers from both academia and industry. As can be seen, our pDPP re-ranking model is compatible with any advanced ranking function, without any modification on such ranking function.

3.2 Practical Issues

To help readers better understand and implement our model in their recommender systems, we summarize several practical issues which should be noticed in real-world applications.

- In research work such as [21], the kernel matrix L is pre-computed and stored in memory, as shown in Algorithm 1. However, such a method cannot be performed in a real-world recommender system, due to the following two reasons. Firstly, the relevance score q_i 's, computed by a ranking function, are personalized and real-time updated. Such industrial-style ranking function makes different relevance scores of individual users to the same item, and furthermore, the relevance score of a user-item pair may be updated in a few seconds as the user feature may be changed. Secondly, our pDPP model has a personalized factor f_u when constructing L so that different users have different L . We need a huge amount of time and storage resources to handle such L 's if we need to pre-compute and store them. Due to these two reasons, we compute the personalized kernel matrix L for a user on-the-fly when this user trigger the request to our recommender system.
- In our experiments, we tried two different approaches to construct the similarity matrix S : one utilizes item features and the other uses user-item interaction information. The method with user-item interaction performs slightly worse than with the other. The reason may due to the fact that user-item interactions are usually very sparse which makes the item representations based on such information not very

reliable. No matter which approach are used, we find that the performance is better when we normalize S_{ij} in $[0, 1]$.

- Cold start problem is one of the common challenges in recommender systems. In our system, we set $\alpha_u = \alpha_0$ if u is a new user. Moreover, users with only few interactions are also regarded as new users by our system. We make such a decision because α_0 is a relatively safe value for exploration, while balancing the trade-off between relevance and diversity.

4 EXPERIMENTAL EVALUATION

To demonstrate the superiority of our pDPP-based re-ranking model, we firstly design offline experiment on two datasets to compare the relevance and diversity of recommendation result of our model with that of baselines. Furthermore, we deploy our model on a live recommender system, to validate its effectiveness in an industry application. In this section, we will present the experiment details and analyze the results in terms of offline and online evaluation, respectively.

4.1 Offline Evaluation

4.1.1 Datasets. For offline evaluation, we prepare two datasets. Besides MovieLens, which is a benchmark in recommendation research community, we also collect user-item interaction log from our commercial App Store. To help reproduce our experiment result, we firstly describe how we process such two datasets.

*MovieLens 1M Dataset*² contains 1,000,209 anonymous ratings with approximately 3,900 movies rated by 6,040 users. As a traditional pre-processing by research work as [13], we eliminate the movies rated by less than 10 users and the users rating less than 20 movies. We randomly split the ratings to two parts, where 70% of the ratings are used for training, and 30% are for testing. A samples with rating greater than or equal to 4 is treated as positive, otherwise negative. We perform item-based collaborating filtering as ranking function to predict the relevance score for each item. The similarity matrix S is built based on the genres of movies, which is to say, $S_{ij} = 1$ if movie i and j are of the same genre.

Company Dataset is collected from our commercial App Store with user consent. This dataset contains approximately 100,000 download records from about 80,000 users in 8 consecutive days. The size of the whole item set is 7,000. Samples in the first 7 days are used for training, while samples in the last day are for testing. We consider all download records as the positive samples, and the others (i.e., the apps that are in the item set but not downloaded by a user) are negative ones. The similarity matrix S is generated by the category of apps, i.e., $S_{ij} = 1$ if app i and j are of the same category.

4.1.2 Baselines. Two baselines are compared. The first baseline considers the ranking function for relevance and disregards re-ranking model for diversity, which is referred as BASE. The other baseline is the standard DPP for re-ranking, which is presented as DPP. Although MMR is a popular state-of-the-art method, we omit it here due to its inferiority compared to DPP [13]. Our personalized DPP model for re-ranking is denoted as pDPP. Note that the ranking

function³ utilized in BASE, DPP and pDPP keeps consistent, for fair comparison. The hyper-parameter α value in DPP and α_0 value in pDPP are found by grid search.

4.1.3 Evaluation Metrics. To compare the models comprehensively, we evaluate them from both *relevance* and *diversity* aspects of their recommendation results. *Precision* is utilized to measure the relevance, which is defined as

$$precision = \frac{\sum_u |R_u \cap T_u|}{\sum_u |R_u|} \quad (7)$$

where R_u denotes recommendation list of user u , T_u denotes download apps of user u in test set.

To measure the diversity, we adopt *intra-list distance (ILD)* [9], which is defined as

$$ILD = avg_u avg_{i,j \in R_u, i \neq j} (1 - S_{ij}) \quad (8)$$

$P@k$ and $ILD@k$ is the *precision* and *ILD* of the first k item in the recommendation list.

Moreover, as a third metric, we measure the balance of *precision* and *ILD* denoted as $avg(P@K, ILD@k)$ in [18, 22], wherein the metrics are standardized to make the scales homogeneous.

Table 1: Experiment Results on MovieLens 1M Dataset

Model	$P@5$	$ILD@5$	$avg(P@5, ILD@5)$
BASE	0.0420	0.5587	0.4000
DPP ($\alpha=0.01$)	0.0420	0.5623	0.4314
DPP ($\alpha=0.02$)	0.0419	0.5735	0.5035
DPP ($\alpha=0.03$)	0.0416	0.5928	0.3717
DPP ($\alpha=0.04$)	0.0415	0.6162	0.5000
pDPP ($l=\mathcal{H}_{\min}$)	0.0420	0.6015	0.7717
pDPP ($l=0$)	0.0421	0.5938	0.8046

4.1.4 Experimental Results. Both offline experiments are performed for multiple times to ensure the results are statistically accurate. During each experiment, we randomly shuffle the data for training and test on the MovieLens Dataset and conduct consecutive experiments within different dates for Company Dataset. Experiment results on MovieLens 1M Dataset with $k = 5$ are shown in Table 1. Due to space limit, we omit the results with other k values, but they are analogous.

DPP model aims to balance the trade-off between accuracy and diversity, for which we can focus more on diversity (i.e., *ILD* metric) by enlarging α , but on the other hand, the relevance performance (i.e., *precision* metric) will be degraded. Compared with BASE, DPP ($\alpha=0.01$) achieves the same accuracy but better diversity. We select a reasonable range for α , to avoid degrading the accuracy significantly, i.e., $\alpha = \{0.01, 0.02, 0.03, \dots, 0.1\}$. Among such values, $\alpha = 0.04$ makes DPP performs the best in terms of $ILD@5$ while $\alpha = 0.02$ enables DPP achieves the best $avg(P@5, ILD@5)$.

As expected, pDPP outperforms all the baselines in terms of $P@5$ and $avg(P@5, ILD@5)$, which demonstrates the superiority of modelling different propensity of diversity for individual users. Specially, pDPP ($l=0$) performs best in terms of $P@5$ and $avg(P@5, ILD@5)$

²<http://grouplens.org/datasets/movielens/1m/>

³In MovieLens 1M Dataset, item-based CF is served as the ranking function; while in Company Dataset, a popular deep learning model is served.

and slightly decreases the performance of $ILD@5$ compared to $pDPP_{(l=\mathcal{H}_{\min})}$.

Table 2: Experiment Results on Company Dataset

Model	$P@5$	$ILD@5$	$avg(P@5, ILD@5)$
BASE	0.0782	0.6554	0.5000
$DPP_{(\alpha=0.1)}$	0.0782	0.6559	0.5014
$DPP_{(\alpha=0.2)}$	0.0782	0.6599	0.5108
$DPP_{(\alpha=0.3)}$	0.0769	0.6695	0.4835
$DPP_{(\alpha=0.4)}$	0.0756	0.6775	0.4525
$DPP_{(\alpha=0.5)}$	0.0769	0.6865	0.5239
$DPP_{(\alpha=0.6)}$	0.0769	0.7048	0.5672
$DPP_{(\alpha=0.7)}$	0.0744	0.7160	0.4937
$DPP_{(\alpha=0.8)}$	0.0744	0.7301	0.5272
$DPP_{(\alpha=0.9)}$	0.0705	0.7682	0.4676
$DPP_{(\alpha=1.0)}$	0.0654	0.8662	0.5000
$pDPP_{(l=\mathcal{H}_{\min})}$	0.0782	0.7025	0.6118
$pDPP_{(l=0)}$	0.0782	0.7051	0.6179

The results on Company Dataset are shown in Table 2. Similar to the experiments on MovieLens 1M Dataset, we select an appropriate range as $\alpha = \{0.1, 0.2, 0.3, \dots, 1.0\}$. Adding DPP-based re-ranking model based on BASE, the DPP models improve the diversity while sacrificing the performance of relevance. Compared with the BASE and DPP models, pDPP models (with 0.6 as the α_0) gain the best performance on $P@5$ and $avg(P@5, ILD@5)$. Remarkably, retaining exactly the same accuracy as BASE, pDPP methods improve the diversity significantly. Specifically, between pDPP family, the model with $l = 0$ achieves better diversity than the one with $l = \mathcal{H}_{\min}$ while keeps the same accuracy performance.

4.2 Online Evaluation

As shown its superior balancing the trade-off between accuracy and diversity in offline evaluation, we deploy pDPP in a live recommender system to verify its effectiveness in an industry application.

4.2.1 Experiment setting. For online evaluation, we conduct online A/B test. We compare three different families of models: BASE, DPP and pDPP. We randomly split all the users into hundreds of bins, each of which consists of more than 100,000 users. A bin of users are served by each of the three compared models. In our live recommender system, the hyper-parameter α of DPP and α_0 of pDPP are set to 0.6 as the performance of $avg(P@5, ILD@5)$ is the best when $\alpha = 0.6$ in offline evaluation (as presented in Table 2).

4.2.2 Evaluation Metrics. To compare the performance of these methods, we evaluate them on the basis two metrics of accuracy and one metric of diversity. The first accuracy metric that we measure is *download ratio (DR)*, defined as

$$DR = \frac{\text{total number of downloads}}{\text{total number of impressions}}. \quad (9)$$

Beyond that, we also measure the engagement of users. More specifically, we study *average number of downloads (AD) per user*, as

$$AD = \frac{\text{total number of downloads}}{\text{total number of users}} \quad (10)$$

Besides these two accuracy metrics, we adopt *ILD* to evaluate the diversity, the same as in offline evaluation.

Table 3: Online A/B testing results

Model	DR	AD	ILD
BASE	-	-	-
DPP	+6.13%	+2.83%	+4.06%
pDPP	+6.52%	+3.54%	+4.32%

4.2.3 Online Performance. The results of A/B online test are shown in Table 3. Considering the commercial concerns, we only present the relative improvement of DPP and pDPP over BASE model in terms of DR , AD and ILD .

We can observe that both DPP and pDPP perform significantly better than BASE in terms of all the three evaluation metrics. It suggests that improving diversity is able to boost the recommendation performance. Between DPP and pDPP, we observe that pDPP is superior than DPP, which indicates that personalized propensity to diversity is more suitable than identical propensity setting. We observe that the improvement of the pDPP over DPP is not as significant as that in offline evaluation. Through detailed analysis, we find that about 35% of the users have only one download record in their behavior history so that it is hard to define their propensity to diversity under such circumstance, which may be one reason for the not-so-significant improvement. However, the daily turnover of our App Store is millions of dollars, therefore even such not-so-significant lift in DR and AD brings extra millions of dollars each year.

5 CONCLUSIONS

Recommender system which only focuses on *accuracy* may lead to sub-optimal, as it too much emphasizes the *accuracy* of each individual items and leads to presenting similar items. *Diversity*, which has been studied to present the users with more diversified items, can be viewed as mutual influence among items. Therefore, combining accuracy and diversity in recommender system is a reasonable and convincing way to improve the performance. Furthermore, different users have different propensity to diversity, which requires personalized diversity. In this paper, we propose a personalized re-ranking model for improving the diversity of the recommendation list based on personalized DPP. This re-ranking model can be easily deployed as a follow-up component after any existing ranking function. The offline experiments over two real-world datasets and the online comparison through A/B testing in an industrial recommender system demonstrate the effectiveness of our proposed re-ranking model.

REFERENCES

- [1] James Davidson, Benjamin Liebald, Junning Liu, Palash Nandy, Taylor Van Vleet, Ullas Gargi, Sujoy Gupta, Yu He, Mike Lambert, Blake Livingston, et al. The youtube video recommendation system. In *Proceedings of the fourth ACM conference on Recommender systems*, pages 293–296, 2010.
- [2] Markus Schedl, Peter Knees, Brian McFee, Dmitry Bogdanov, and Marius Kaminakas. Music recommender systems. In *Recommender systems handbook*, pages 453–492. Springer, 2015.
- [3] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, Xiuqiang He, and Zhenhua Dong. Deepfm: An end-to-end wide & deep learning framework for ctr prediction. *arXiv preprint arXiv:1804.04950*, 2018.
- [4] Qing-Cheng Li, Zhen-Hua Dong, and Tuo Li. Research of information recommendation system based on reading behavior. In *2008 International Conference on Machine Learning and Cybernetics*, volume 3, pages 1626–1631. IEEE, 2008.
- [5] Sean M McNee, John Riedl, and Joseph A Konstan. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI'06 extended abstracts on Human factors in computing systems*, pages 1097–1101. ACM, 2006.
- [6] Eli Pariser. *The filter bubble: What the Internet is hiding from you*. Penguin UK, 2011.
- [7] Tien T Nguyen, Pik-Mai Hui, F Maxwell Harper, Loren Terveen, and Joseph A Konstan. Exploring the filter bubble: the effect of using recommender systems on content diversity. In *Proceedings of the 23rd international conference on World wide web*, pages 677–686. ACM, 2014.
- [8] Cai-Nicolas Ziegler, Sean M McNee, Joseph A Konstan, and Georg Lausen. Improving recommendation lists through topic diversification. In *Proceedings of the 14th international conference on World Wide Web*, pages 22–32. ACM, 2005.
- [9] Mi Zhang and Neil Hurley. Avoiding monotony: improving the diversity of recommendation lists. In *Proceedings of the 2008 ACM conference on Recommender systems*, pages 123–130. ACM, 2008.
- [10] Keith Bradley and Barry Smyth. Improving recommendation diversity. In *Proceedings of the Twelfth Irish Conference on Artificial Intelligence and Cognitive Science, Maynooth, Ireland*, pages 85–94. Citeseer, 2001.
- [11] Gediminas Adomavicius and YoungOk Kwon. Improving aggregate recommendation diversity using ranking-based techniques. *IEEE Transactions on Knowledge and Data Engineering*, 24(5):896–911, 2012.
- [12] Jaime G Carbonell and Jade Goldstein. The use of mmr and diversity-based reranking for reordering documents and producing summaries. 1998.
- [13] Laming Chen, Guoxin Zhang, and Hanning Zhou. Improving the diversity of top-n recommendation via determinantal point process. In *Large Scale Recommendation Systems Workshop at the Conference on Recommender Systems (RecSys)*. <http://arxiv.org/abs/1709.05135>, 2017.
- [14] Alex Kulesza and Ben Taskar. k-dpps: Fixed-size determinantal point processes. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pages 1193–1200, 2011.
- [15] Li Chen, Wen Wu, and Liang He. How personality influences users' needs for recommendation diversity? In *CHI'13 Extended Abstracts on Human Factors in Computing Systems*, pages 829–834. ACM, 2013.
- [16] Wen Wu, Li Chen, and Yu Zhao. Personalizing recommendation diversity based on user personality. *User Modeling and User-Adapted Interaction*, 28(3):237–276, 2018.
- [17] Li Chen, Wen Wu, and Liang He. Personality and recommendation diversity. In *Emotions and Personality in Personalized Services*, pages 201–225. Springer, 2016.
- [18] Tommaso Di Noia, Vito Claudio Ostuni, Jessica Rosati, Paolo Toméo, and Eugenio Di Sciascio. An analysis of users' propensity toward diversity in recommendations. In *Proceedings of the 8th ACM Conference on Recommender systems*, pages 285–288. ACM, 2014.
- [19] Farzad Eskandarian, Bamshad Mobasher, and Robin Burke. A clustering approach for personalizing diversity in collaborative recommender systems. In *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization*, pages 280–284. ACM, 2017.
- [20] Mark Wilhelm, Ajith Ramanathan, Alexander Bonomo, Sagar Jain, Ed H Chi, and Jennifer Gillenwater. Practical diversified recommendations on youtube with determinantal point processes. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pages 2165–2173. ACM, 2018.
- [21] Laming Chen, Guoxin Zhang, and Eric Zhou. Fast greedy map inference for determinantal point process to improve recommendation diversity. In *Advances in Neural Information Processing Systems*, pages 5622–5633, 2018.
- [22] Umberto Panniello, Alexander Tuzhilin, and Michele Gorgoglione. Comparing context-aware recommender systems in terms of accuracy and diversity. *User Modeling and User-Adapted Interaction*, 24(1-2):35–65, 2014.