

NLGR: Utilizing Neighbor Lists for Generative Rerank in Personalized Recommendation Systems

Shuli Wang*
Meituan
Chengdu, China
wangshuli03@meituan.com

Xue Wei
Meituan, Chengdu, China
weixue06@meituan.com

Senjie Kou
Meituan
Chengdu, China
kousenjie@meituan.com

Chi Wang
Meituan
Chengdu, China
wangchi06@meituan.com

Wenshuai Chen
Meituan
Chengdu, China
chenwenshuai@meituan.com

Qi Tang
Meituan
Chengdu, China
tangqi22@meituan.com

Yinhua Zhu
Meituan
Chengdu, China
zhuyinhua@meituan.com

Xiong Xiao
Meituan
Chengdu, China
xiaoxiong02@meituan.com

Xingxing Wang
Meituan
Beijing, China
wangxingxing04@meituan.com

Abstract

Reranking plays a crucial role in modern multi-stage recommender systems by rearranging the initial ranking list. Due to the inherent challenges of combinatorial search spaces, some current research adopts an evaluator-generator paradigm, with a generator generating feasible sequences and an evaluator selecting the best sequence based on the estimated list utility. However, these methods still face two issues. Firstly, due to the goal inconsistency problem between the evaluator and generator, the generator tends to fit the local optimal solution of exposure distribution rather than combinatorial space optimization. Secondly, the strategy of generating target items one by one is difficult to achieve optimality because it ignores the information of subsequent items.

To address these issues, we propose a utilizing Neighbor Lists model for Generative Reranking (NLGR), which aims to improve the performance of the generator in the combinatorial space. NLGR follows the evaluator-generator paradigm and improves the generator's training and generating methods. Specifically, we use neighbor lists in combination space to enhance the training process, making the generator perceive the relative scores and find the optimization direction. Furthermore, we propose a novel sampling-based non-autoregressive generation method, which allows the generator to jump flexibly from the current list to any neighbor list. Extensive experiments on public and industrial datasets validate NLGR's effectiveness and we have successfully deployed NLGR on the Meituan food delivery platform.

* Corresponding author.

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

WWW Companion '25, Sydney, NSW, Australia

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1331-6/25/04
<https://doi.org/10.1145/3701716.3715251>

CCS Concepts

• Information systems → Retrieval models and ranking; Computational advertising.

Keywords

Recommender Systems, Reranking, Generative Model

ACM Reference Format:

Shuli Wang, Xue Wei, Senjie Kou, Chi Wang, Wenshuai Chen, Qi Tang, Yinhua Zhu, Xiong Xiao, and Xingxing Wang. 2025. NLGR: Utilizing Neighbor Lists for Generative Rerank in Personalized Recommendation Systems. In *Companion Proceedings of the ACM Web Conference 2025 (WWW Companion '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3701716.3715251>

1 Introduction

E-commerce platforms, such as Meituan and Taobao, need to provide users with personalized services from millions of items. To improve recommendation efficiency, personalized recommendation systems generally include three stages: matching, ranking, and reranking. The ranking models (e.g., Wide&Deep [10], DeepFM [14], DIN [35]) evaluate the point-wise items respectively based on the Click-Through Rate (CTR), but they ignore the crucial mutual influence among items [7, 27]. Research [1, 5, 6, 25] indicates that optimizing a list-wise utility is a more advantageous strategy, as it capitalizes on the mutual influences between items within the list to enhance overall performance.

The key challenge of reranking is to explore the optimal list in the huge combinatorial space [21, 28]. Initially, some list-wise methods [2, 26, 36] re-evaluate and score items within lists by modeling the context. These list-wise methods can obtain more accurate scores than point-wise methods, then they use a greedy strategy to reorder based on the list-wise score. However, these methods face the evaluation-before-reranking problem [13, 32] and cannot achieve optimization in combinatorial space. To resolve the problem, a straightforward solution is to evaluate every possible permutation, which is global-optimal but is too complex to meet

the strict inference time constraint in industrial systems. Therefore, most existing evaluator-based reranking framework uses a two-stage architecture [12, 13, 20, 28, 32] which consists of a generator and an evaluator. Within the generator-evaluator paradigm, the generator plays a crucial role [28]. Some methods use heuristic methods as generators, such as beam-search [24] and SimHash [9]. The generators of these methods do not utilize the information of the evaluator, resulting in limited effectiveness. Recently, some methods [13, 20, 28] utilize generative models as generators and achieve better results than heuristic methods.

However, existing generative reranking methods face two significant challenges. Firstly, due to the goal inconsistency problem, the generator has difficulty finding the optimal list in the combinatorial space. While the evaluator is trained to fit list-wise scores of items, the generator is tasked with transforming any candidate list into the optimal one. This disparity in objectives between the evaluator and the generator complicates the transfer of guidance, often causing the generator to merely fit exposure distributions in extreme cases. Secondly, the strategy of generating target items sequentially, one by one, hinders the achievement of optimal results. The sequential decoding process focuses solely on preceding items, neglecting information from succeeding items. This limitation leads to suboptimal performance as the model fails to fully leverage the available context.

To address the aforementioned challenges, we propose a novel Generative Reranking method that utilizes Neighbor Lists, named NLGR. Our approach still follows the evaluator-generator paradigm, with the evaluator solely assisting in the generator's offline training. Our improvements are twofold as follows. First, we introduce an enhanced training process that utilizes neighbor lists, enabling the generator to perceive relative scores within the combinatorial space and identify the optimal direction. As depicted in Figure 1, by evaluating multiple neighboring lists and iterating several times, the generator will converge to an optimal state. Second, we propose a sampling-based non-autoregressive generation method. This method first determines the position of the item that needs to be replaced using the Position Decision Unit (PDU), and then retrieves new replacement items from the candidate item set using the Candidate Retrieval Unit (CRU), which allows for a more flexible exploration of the combinatorial space to find the optimal list.

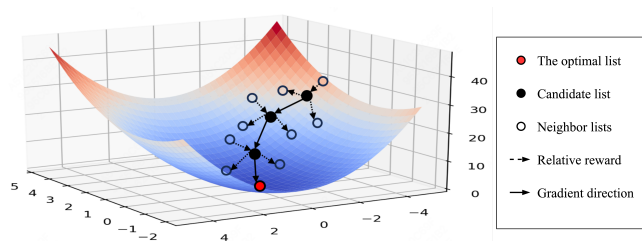


Figure 1: Generator optimization legend: reach the optimal step by step under the guidance of the neighbor lists.

In summary, the contributions of this work are as follows:

- We propose a novel generative reranking method that utilizes neighbor lists to address the goal inconsistency problem between

evaluators and generators. To the best of our knowledge, we are the first to propose and attempt to solve this problem.

- We propose a novel sampling-based non-autoregressive generation method that generates the optimal list more flexibly in the combination space.
- We have verified the superior performance of NLGR through extensive experiments on both offline and online datasets. It is notable that NLGR has been deployed in the Meituan food delivery platform and has achieved significant improvement under various metrics.

2 Related Work

2.1 Reranking Methods

Typical reranking methods can be divided into two categories [29]. The first category is the one-stage reranking methods, which only generates one list as output by capturing the mutual influence among items. Seq2slate [4] utilizes pointer-network and MIRNN [37] utilizes GRU to determine the item order one-by-one. Methods such as PRM [26] and DLCM [1] take the initial ranking list as input, use RNN or self-attention to model the context-wise signal, and output the predicted value of each item. Such methods bring an evaluation-before-reranking problem [32] and lead to suboptimal. Similarly, methods such as EXTR [8] estimate pCTR of each candidate item on each candidate position, which are substantially point-wise models and thus limited in extracting exact context. MIR [33] capturing the set2list interactions by a permutation-equivariant module. Another category is the two-stage reranking methods, which tries to evaluate every possible permutation through a well-designed context-wise model. This is a global-optimal method but is too complex to meet the strict inference time constraint in industrial systems. To reduce the complexity, PRS [12] adopts beam-search to generate a few candidate permutations first, and score each permutation through a permutation-wise ranking model. PIER [29] applies SimHash to select top-K candidates from the full permutation.

2.2 Generative Reranking Solutions

In recent years, the generative reranking model [16, 18, 22] for listwise recommendation has been a topic of discussion. To manage the vast combinatorial output space of lists, the generative approach directly models the distribution of recommended lists and employs deep generative models to generate a list as a whole. For instance, ListCVAE [18] utilizes conditional variational autoencoders (CVAE) to capture the positional biases of items and the interdependencies within the list distribution. But Pivot-CVAE [22] indicates that ListCVAE suffers from a trade-off dilemma between accuracy and diversity, and proposes an "elbow" performance to enhance the accuracy-based evaluation.

GFN [21] uses a flow-matching paradigm that maps the list generation probability with its utility. Essentially it is still studying list distributions rather than directly modeling the permutation space, so it still has the challenge mentioned above. GRN [13] proposes an evaluator-generator framework to replace the greedy strategy, but it can't avoid the evaluation-before-reranking problem [32] because it takes the rank list as input to the generator. DCDR [20] introduces diffusion models into the reranking stage and presents a discrete conditional diffusion reranking framework. NAR4Rec [28]

uses a non-autoregressive generative model to speed up sequence generation. However, these methods still face the two problems mentioned above.

3 Problem Definition

In the Meituan food delivery platform, we adhere to the matching, ranking, and reranking recommendation paradigms to present items to users in a list format. Initially, we define a user set U and an item set I . We utilize the session-level users' historical interacted lists B and candidate set C to represent the user's features $u \in U$, which consistently holds in our reranking scenario $C \in I$, and ultimately select the list L for users.

Reranking introduces a combination space with exponential size, represented as $\mathcal{O}(A_n^m)$, where n represents the size of the candidate set C and m represents the size of the output list L . Our optimization objective is to learn a strategy $\pi : C \rightarrow L$ by maximizing the list score reward $R(u, \pi)$. The list score reward $R(u, \pi)$ takes into account factors such as click-through rate (CTR) and conversion rate (CVR).

Neighbor List. If the distance between two lists is 1, that is, the two lists only differ by 1 item, we define the two lists as neighbor lists. If we swap two items in a list, the new list will have a distance of 2.

4 Proposed Method

In this section, we present a detailed introduction of NLGR. We first introduce the evaluator (in Section 4.1) and generator (in Section 4.2) of NLGR, denoted as NLGR-E and NLGR-G respectively. Then, we demonstrate the offline training process of NLGR in Section 4.3. The evaluator is only assisting in the offline training of the generator, so NLGR is a fast inference method.

4.1 Evaluator Model

We use NLGR-E to evaluate a ranked sequence, such as the exposed list which is displayed to the user, or the candidate lists generated by NLGR-G. The structure of NLGR-E is shown on the left side of Figure 2.

NLGR-E includes two inputs: the exposed list to be evaluated and the user session-level behavior sequence, where each session in the user session-level behavior sequence is the user's historical exposed list. The evaluation process of NLGR-E is as follows:

First, we use an embedding layer to get the embedding of the original input, denoted as $\mathbf{X} \in \mathbb{R}^{m \times F \times D}$ and $\mathbf{M} \in \mathbb{R}^{H \times m \times F \times D}$ respectively, where H represents the number of history sessions, m represents the number of items in each list, F represents the number of feature fields for each item (i.g., ID, category, position index), and D represents the dimension of the embedding. Inspired by DIF [34], to avoid feature interference, we propose the D-Attention unit to decouple the feature context information. We first calculate the attention score on i -th attribute $\mathbf{X}_i \in \mathbb{R}^{m \times D}$:

$$\text{Att}_i = \sigma \left(\frac{(\mathbf{X}_i \mathbf{W}_i^Q)(\mathbf{X}_i \mathbf{W}_i^K)^\top}{\sqrt{D}} \right), \forall i \in [F], \quad (1)$$

where \mathbf{W}_i^Q and $\mathbf{W}_i^K \in \mathbb{R}^{D \times D}$ denote the weight matrices.

Then we aggregate all attention matrices $\text{Att}_i \in \mathbb{R}^{m \times m}$ and get the item-level attention score $\text{Att}_{\text{all}} \in \mathbb{R}^{m \times m}$:

$$\text{Att}_{\text{all}} = \frac{1}{F} \sum_{f=1}^F \text{Att}_f. \quad (2)$$

Subsequently, we aggregate ID feature embedding $\mathbf{X}^{id} \in \mathbb{R}^{m \times D}$ and obtain each exposed list's representation $\mathbf{e}^l \in \mathbb{R}^D$:

$$\mathbf{e}^l = \text{reduce_mean} \left(\text{Att}_{\text{all}}(\mathbf{X}^{id} \mathbf{W}^V) \right). \quad (3)$$

Similarly, by performing the above operations on each session in the user session-level behavior sequence, we can obtain the representation $\mathbf{e}_i^s \in \mathbb{R}^D$ of each session. Then we input \mathbf{e}_i^s into a Self-Attention layer (SA) [19, 30] to get the user representation $\mathbf{e}^u \in \mathbb{R}^D$:

$$\mathbf{e}^u = \text{SA}(\mathbf{e}_1^s \| \mathbf{e}_2^s \| \dots \| \mathbf{e}_H^s), \quad (4)$$

where $\|$ represents concatenate operate.

Finally, the predicted Click-Through Rate (pCTR) of the j -th item can be represented as:

$$\hat{y}_j = \sigma(\text{Tiled_MLP}(\mathbf{X}_j \| \mathbf{e}^l \| \mathbf{e}^u \| PE_j)), \quad (5)$$

where PE_j denotes j -th position embedding. Similarly, the pCVR and other evaluations also follow the above process.

4.2 Generator model

We utilize NLGR-G to generate the optimal list in combinatorial space. The structure of NLGR-G is shown on the middle side of Figure 2.

Similar to NLGR-E, NLGR-G includes two inputs: the candidate list to be reranked and the user session-level behavior sequence, where each session in the user session-level behavior sequence is the user's historical exposed list. Theoretically, the candidate list can be any in the combinatorial space, but generally, we use the ranking list as the initial input. The generation process of NLGR-G is as follows:

First, we obtain the user representation \mathbf{e}^u through Eq. 4. These parameters are shared from NLGR-E to ensure that the remaining parameters can be optimized more focused.

Then, we propose a sampling-based non-autoregressive generation method. It first determines the position of the item that needs to be replaced through the Position Decision Unit (PDU), then retrieves new replacement items from the candidate item set through the Candidate Retrieval Unit (CRU).

4.2.1 Position Decision Unit (PDU). First, we use the embedding layer to get the embedding of the candidate list, denoted as $\mathbf{X} \in \mathbb{R}^{m \times F \times D}$. Then, we flatten j -th item embedding $\mathbf{X}_j \in \mathbb{R}^{F \times D}$ and use a Fully Connected layer (FC) to calculate the selected logit of the j -th position:

$$h_j = \text{FC}_1(\mathbf{X}_j \| \mathbf{e}^l \| \mathbf{e}^u \| PE_j). \quad (6)$$

To solve the non-differentiable problem of the sampling distribution, inspired by [15, 17, 23], we use the Gumbel-softmax trick for sampling:

$$r_j^p = \text{softmax} \left(\frac{\log(h_j) + n}{\tau} \right), \forall j \in [m], \quad (7)$$

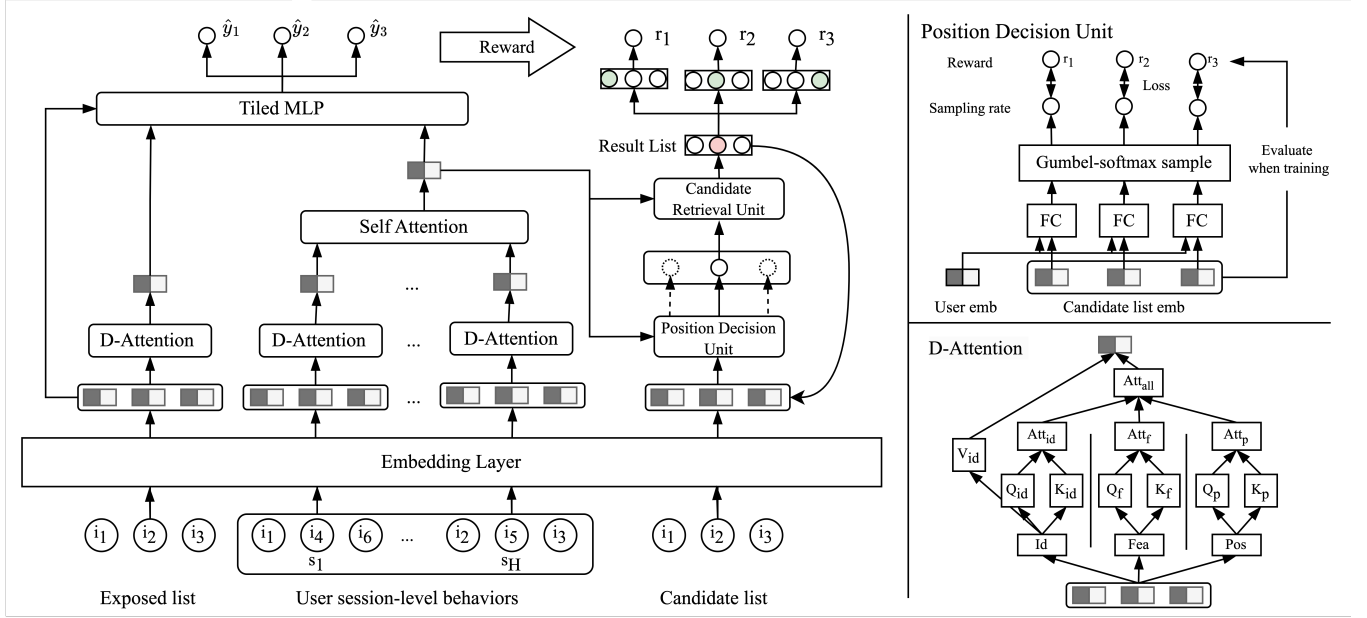


Figure 2: The overall architecture of NLGR.

where $\tau > 0$ is a temperature parameter, $n = -\log(-\log(u))$ represents random noise sampled from the Gumbel distribution, u is a uniform distribution between $[0, 1]$. During backpropagation, the gradient is calculated using Eq. 7. While during forward propagation, the replaced position is $j = \text{argmax}(r_j^p)$.

4.2.2 Candidate Retrieval Unit (CRU). After determining the position j to be replaced, we need to select a suitable one from n candidate items and place it at position j . Since this operation will be repeated multiple times during the generation process, we propose leveraging retrieval-based techniques to quickly achieve this goal for efficiency. First, we mask the position j of the candidate list, denoted as $\mathbf{X}_j^{\text{mask}} \in \mathbb{R}^{m \times F \times D}$, and then extract the list representation $\mathbf{e}_j^{\text{mask}} \in \mathbb{R}^D$:

$$\mathbf{e}_j^{\text{mask}} = \text{SA}(\mathbf{X}_j^{\text{mask}}). \quad (8)$$

We then compute the representation of each candidate item at position j :

$$\mathbf{e}_k^c = \text{FC}_2(\text{flatten}(\mathbf{X}_k^c) || PE_j), \forall k \in [n], \quad (9)$$

where $\mathbf{X}^c \in \mathbb{R}^{n \times F \times D}$ denotes embedding of candidate set C , and \mathbf{X}_k^c denotes k -th candidate item's embedding.

Then we use an FC layer to calculate the selected logit of the $j = k$ -th candidate item:

$$g_k = \text{FC}_3(\mathbf{e}_k^c || \mathbf{e}_j^{\text{mask}} || \mathbf{e}^u || PE_j). \quad (10)$$

Similarly, to overcome non-differentiable problems, we use the Gumbel-softmax trick for sampling:

$$r_k^c = \text{softmax}\left(\frac{\log(g_k) + n}{\tau}\right), \forall k \in [n], \quad (11)$$

During backpropagation, the gradient is calculated using Eq. 10. While during forward propagation, the newly inserted item is $c = \text{argmax}(r_k^c)$.

Stop condition. Note that the generation process may be repeated many times until the newly inserted item equals the replaced item or the values of r_j^p and r_k^c are too low.

4.3 Utilizing Neighbor Lists Training

In this section, we elaborate on the offline training process of NLGR, which includes the training procedures for NLGR-E and NLGR-G. As mentioned before, the evaluator is trained to fit list-wise scores of items, and the generator is tasked with transforming any ranking list into the optimal one. This goal inconsistency between the evaluator and the generator complicates the transfer of guidance. We will introduce our solution in detail below.

4.3.1 Training of NLGR-E. To accurately evaluate the return of the exposure list and estimate the listwise pCTR value of the exposure list, we train NLGR-E using real data collected from online logs. The input is the features of the recommended advertisement sequences exposed in reality online, and the advertising return situation, including exposure, click, conversion, and other performance indicators, is used as the label to supervise the training of NLGR-E, enabling it to accurately evaluate the return of the recommended sequence. The loss of NLGR-E is calculated as follows:

$$\mathcal{L}^E = \sum_{j=1}^m (-y_j \log(\hat{y}_j) - (1 - y_j) \log(1 - \hat{y}_j)), \quad (12)$$

where y_j represents the real label, \hat{y}_j represents the predicted value of NLGR-E, and the evaluation is carried out for the m items in the exposure list in turn.

4.3.2 Training of NLGR-G. To address the problem of goal inconsistency mentioned before, we use neighbor lists to guide NLGR-G within the counterfactual space. For each list generated by NLGR-G, NLGR-E simulates human feedback and provides a reward R to guide NLGR-G training.

Figure 3 shows an example of the NLGR-G training process. First, for each candidate list $L^o = [i_1^o, i_2^o, \dots, i_m^o]$, we sample a replacement item i^* from the candidate set $C = [i_1, i_2, \dots, i_n]$, replace the j -th item L_o 's, and construct a neighbor list $L_j^* = [i_1^o, i_2^o, \dots, i_j^*, \dots, i_m^o]$. Repeating and replacing each position, we can get a set of neighbor lists $L^* = [L_1^*, L_2^*, \dots, L_m^*]$.

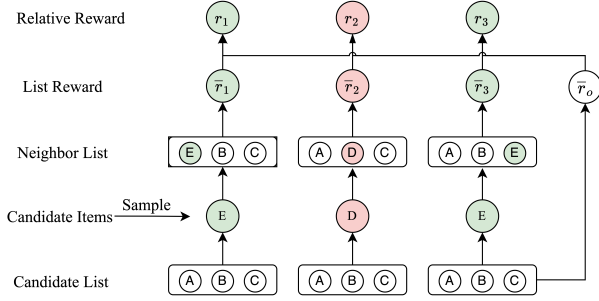


Figure 3: The training process of NLGR-G for a candidate list of length 3.

Then, we use the trained NLGR-E to evaluate the candidate list L^o and the neighbor list L^* , and indicators such as pCTR and pCVR will be estimated. We convert the estimated value into list reward based on business indicators as follows:

$$\bar{r} = \begin{cases} e^{w-1} - 1 & \text{if } w > 1 \\ 0 & \text{if } w = 1 \\ 1 - e^{1-w} & \text{if } w < 1 \end{cases}, \quad (13)$$

$$w = k_1 \cdot L_{ctr} + k_2 \cdot L_{ctr} \cdot L_{cvr}, \quad (14)$$

where L_{ctr} and L_{cvr} represent the list total pCTR and pCVR which are evaluated by NLGR-E, respectively. The parameters k_1 and k_2 are business parameters that depend on the click bid and conversion price in the specific business.

Through Eq. 13, we can get the rewards of the neighbor lists L^* and the original candidate list L^o , denoted as $\bar{r} = [\bar{r}_1, \bar{r}_2, \dots, \bar{r}_m]$ and \bar{r}_o respectively. NLGR-G is tasked to iterate from the candidate list to a more optimal list, so we calculate the relative reward for each position j :

$$r_j = \bar{r}_j - \bar{r}_o, \forall j \in [m]. \quad (15)$$

The authentic evaluation R for the candidate list L^o is obtained by aggregating the relative rewards of all positions. And we define the counterfactual loss of NLGR-G as $-R$:

$$\mathcal{L}_1^G = -R = -\sum_{j=1}^m r_j. \quad (16)$$

Furthermore, to increase the stability of the NLGR-G's generation process, we propose using the position reward r_j to provide additional guidance to the Position Decision Unit (PDU). Specifically, we introduce cross-entropy loss as an auxiliary loss to measure the

distribution difference of position sampling r^p and position reward r_j :

$$\mathcal{L}_2^G = -\sum_{j=1}^m \text{Norm}(r_j) \cdot \log r_j^p, \quad (17)$$

where $\text{Norm}(r_j) = \frac{r_j}{\sum r_j}$.

The final loss of NLGR-G in each batch is:

$$\mathcal{L}^G = \frac{1}{|B|} \sum_B (\mathcal{L}_1^G + \alpha \cdot \mathcal{L}_2^G), \quad (18)$$

where α is a coefficient to balance the two losses.

5 Experiments

To validate the superior performance of NLGR, we conducted extensive offline experiments on the Meituan dataset and verified the superiority of NLGR in online A/B tests. In this section, we first introduce the experimental setup, including the dataset and baseline. Then, in Section 5.2, we present the results and analysis of various reranking methods in both offline and online A/B tests.

5.1 Experimental Setup

5.1.1 Dataset. In order to verify the effectiveness of NLGR, we conduct sufficient experiments on both public dataset and industrial dataset. For public dataset, we choose Taobao Ad dataset¹. For the industrial dataset, we use real-world data collected from Meituan food delivery platform.

- Taobao Ad. It is a public dataset collected from the display advertising system of Taobao. This dataset contains more than 26 million interaction records of 1.14 million users within 8 days. Each sample comprises five features: user ID, timestamp, behavior type, item brand ID, and category ID.
- Meituan. It is an industrial dataset collected from the Meituan food delivery platform during October 2023, which contains 1.3 billion interaction records of 130 million users within 30 days. The dataset includes 239 features, two labels: click and conversion, and collects all items on the same page as one record. We divide the dataset into training and test sets with a proportion of 9:1.

Table 1 gives a brief introduction to the datasets.

Table 1: Statistics of datasets

Dataset	#Users	#Items	#Records
Taobao Ad	1,141,729	99,815	26,557,961
Meituan	130,648,310	14,054,691	1,331,247,488

5.1.2 Baseline. The following six baselines are chosen for comparative experiments and divided into three groups. We select DNN and DeepFM as point-wise baselines (Group I), PRM and MIR as list-wise baselines (Group II), and GRN and DCDR as generative baselines (Group III). A brief introduction of these methods is as follows:

- DNN[11] is a basic deep learning method for CTR prediction, which applies MLP for high-order feature interaction.

¹ <https://tianchi.aliyun.com/dataset/56>

Table 2: Performance comparison. The best result and the second-best result in each column are in bold and underlined

Dataset	Metric	Group I		Group II		Group III		NLGR
		DNN	DeepFM	PRM	MIR	GRN	DCDR	
Taobao Ad	AUC	0.5869	0.5871	0.6052	0.6047	0.6101	<u>0.6217</u>	0.6344 $\pm 0.5\%$
	LogLoss	0.1878	0.1866	0.1842	0.1853	0.1820	<u>0.1792</u>	0.1749 $\pm 0.1\%$
	NDCG@10	0.1527	0.1548	0.1805	0.1769	0.1896	<u>0.2038</u>	0.2323 $\pm 0.2\%$
	NDCG@5	0.1092	0.1110	0.1206	0.1193	0.1273	<u>0.1453</u>	0.1830 $\pm 0.1\%$
Meituan	AUC	0.7034	0.7070	0.8096	0.8031	0.8034	<u>0.8181</u>	0.8349 $\pm 0.4\%$
	LogLoss	0.1162	0.1162	0.1096	0.1102	0.1102	<u>0.1087</u>	0.1039 $\pm 0.1\%$
	NDCG@10	0.2015	0.2019	0.2743	0.2742	0.2744	<u>0.2793</u>	0.2857 $\pm 0.2\%$
	NDCG@5	0.1569	0.1580	0.2378	0.2365	0.2353	<u>0.2400</u>	0.2431 $\pm 0.2\%$

Table 3: Hit ratio comparison. The best result and the second-best result in each column are in bold and underlined

Model	Taobao Ad		Meituan	
	HR@10%	HR@1%	HR@10%	HR@1%
PRM	0.1125	0.0794	0.5702	0.4412
GRN	<u>0.2160</u>	0.0829	0.7386	0.5783
DCDR	0.2159	<u>0.0835</u>	<u>0.7573</u>	<u>0.5802</u>
NLGR	0.4091	0.3220	0.8369	0.7523

- **DeepFM**[14] is a general deep model for recommendation, which combines a factorization machine component and a deep neural network component.
- **PRM**[26] adjusts an initial list by applying the self-attention mechanism to capture the mutual influence between items.
- **MIR**[33] learns permutation-equivariant representations for the inputted items via self-attention mechanism to capture the mutual influence between items.
- **GRN**[13] is a generative reranking model which consists of the evaluator for predicting interaction probabilities and the generator for generating reranking results.
- **DCDR**[20] presents a discrete conditional diffusion reranking framework.

5.1.3 Evaluation Metrics. We adopt several metrics, i.e., **AUC** (Area Under ROC Curve), **Logloss** and **NDCG** (normalized discounted cumulative gain) to evaluate NLGR-E in offline experiments. A larger AUC and NDCG indicate better recommendation performance, while Logloss performs the opposite.

We use **HR** (Hit Ratio) [3] to evaluate NLGR-G in offline experiments. It is worth noting that only one list produced by reranking algorithms can be presented to the user. As a result, the generator cannot be fully and fairly evaluated. A practical workaround is to employ the evaluator to assess the performance of the generator. For each data, we evaluate all candidate lists using NLGR-E. HR@10% is 1 only when the rerank list produced by NLGR-G is ranked within the top 10% as sorted by NLGR-E. The HR metrics are only meaningful with evaluator-based reranking methods. The results of HR can be seen in Table 3.

It is worth noting that AUC and HR can measure the evaluator and generator respectively. AUC measures the model’s ability to

evaluate an ordered list, while HR measures the consistency between the evaluator and generator. A decrease in any indicator will reduce the recommendation effect.

In online experiments, we adopt CTR and GMV (Gross Merchandise Volume) as evaluation metrics.

5.1.4 Implementation Details. We implement all the deep learning baselines and NLGR with TensorFlow 1.15.0 using NVIDIA A100-SXM4-80GB. For all comparison models and our NLGR model, we adopt Adam as the optimizer with the learning rate fixed to 0.001 and initialize the model parameters with normal distribution by setting the mean and standard deviation to 0 and 0.01, respectively. The batch size is 1024, the embedding size is 8, the α is 0.2. The hidden layer sizes of tiled MLP are (1024, 256, 128). For the Taobao Ad dataset, the length of the ranking list and reranking list are both 5, thus the length of full permutation is 120. For Meituan dataset, we select 4 items from the initial ranking list which contains 12 items, thus the length of full permutation is $A_{12}^4 = 11,880$. All experiments are repeated five times and the averaged results are reported.

5.2 Experimental Results

5.2.1 Performance Comparison. Table 2 and Table 4 summarize the results of offline experiments. We have the following observations from the experimental results:

- PRM in Group II outperforms all models in Group I, which verifies the impact of the mutual influence among contextual items. DCDR in Group III outperforms all the other models in Groups I and II, which verifies the effectiveness of generative methods.
- DCDR indeed exhibits robust generative capabilities, thanks to the incorporation of the diffusion model, and achieves the second-best result. Nevertheless, DCDR overlooks the significance of full sight and falls short of leveraging the evaluator’s potential to its fullest, which limits its effect. Our proposed NLGR significantly and consistently outperforms the state-of-the-art approaches in all metrics on both datasets. As presented in Table 2, our proposed NLGR brings 0.8349/0.6344 absolute AUC, 0.2857/0.2323 absolute NDCG@10, 0.2431/0.1830 absolute NDCG@5 on Meituan/Ad dataset, gains significant improvement in industrial recommendation system. NLGR has greater improvements on Meituan dataset because Meituan dataset has more realistic reranking scenarios and richer features.

iii) Our proposed NLGR brings 0.8369/0.4091 absolute HR@10% on Meituan/Ad dataset. This demonstrates our generator achieves extreme improvements via counterfactual evaluation. Compared with GRN and PRM, there is a huge improvement. As a typical greedy reranking algorithm, PRM only has 0.5702/0.1125 absolute HR@10% on Meituan/Ad dataset. This conclusively shows that effective rearrangement cannot be attained by relying solely on a greedy strategy. As a generative reranking model, DCDR lacks full sight, which is why its HR@10% is not optimal.

5.2.2 *Ablation Study.* To verify the impact of different units, we study three ablated variants of NLGR on Meituan dataset.

- NLGR without relative reward r . To verify the effectiveness of the neighbor list training method in addressing the goal inconsistency problem, this variant removes the relative reward defined in Eq. 16 and replaces it with the predicted value returned by the evaluator directly.
- NLGR with autoregressive generation (AG). To verify the sampling-based non-autoregressive generation method in NLGR-G, we replaced it with pointer network [31], which is a sequence generation method widely used in reranking models [4, 13].
- NLGR with \mathcal{L}_2^G . To further verify the effectiveness of the neighbor list training method, we remove \mathcal{L}_2^G in Eq. 18, which means that PDU will lack direct guidance from NLGR-E.

The result is shown in Table 4. From the experimental results, we have the following findings: i) The HR of NLGR w/o r drops the most (8%/17.2%), indicating that neighbor list training is the most important part of NLGR. ii) The HR of NLGR w/ AG dropped significantly (2.2%/4.8%), indicating that the sampling-based non-autoregressive generation method in NLGR-G can significantly improve the generation effect. iii) The HR of NLGR w/o \mathcal{L}_2^G also decreases (1.1%/3.2%), indicating that auxiliary loss can enhance the generation ability of NLGR-G.

Table 4: HR of different methods on Meituan

method	HR@10%	HR@1%
NLGR	0.8369	0.7523
NLGR w/o r	0.7562	0.5809
NLGR w/ AG	0.8142	0.7047
NLGR w/o \mathcal{L}_2^G	0.8255	0.7198

5.3 Hyperparameter Analysis

We analyze the sensitivity of two hyperparameters: α and β , corresponding to the generation process and training process of NLGR. Among them, α is the weight of NLGR-G loss in Eq. 18, and β is the sampling ratio at each position when constructing the neighbor list L^* . By default $\beta = 1$ means that each position is sampled 1 time. The result is shown in Table 5, showing the same trend on the public dataset and industrial dataset and we have the following findings:

- i) Hyperparameter α significantly affects the generator’s HR@10% metric. When $\alpha = 0$, NLGR is equivalent to the method of Group

III, which means the generator has no full sight. As α increases, HR@10% first increases and then decreases.

- ii) We tested several values for β . When $\beta < 1$, we randomly select βm positions in m positions to construct rewards. When $\beta > 1$, we construct β neighbor lists at each position. Increasing β within a certain range can quickly improve the HR@10% performance. As β continues to increase, HR@10% remains stable but increases offline training time. The results show that counterfactual rewards considering all positions are important.

Table 5: HR@10% result of ablation experiment on NLGR

	$\alpha=0$	$\alpha=0.01$	$\alpha=0.2$	$\alpha=0.5$	$\alpha=1.0$
Meituan	0.7562	0.8274	0.8369	0.8301	0.8295
Taobao Ad	0.2192	0.3530	0.4091	0.3973	0.3912
	$\beta=0.1$	$\beta=0.5$	$\beta=1$	$\beta=2$	$\beta=5$
Meituan	0.7763	0.8031	0.8369	0.8369	0.8367
Taobao Ad	0.2814	0.3566	0.4091	0.4091	0.4091

5.4 Online A/B test

We deployed NLGR in Meituan App, where Figure 4 shows the online serving system architecture. It is worth noting that although we involve the evaluator guiding the generator multiple times during offline training, we only need to use the generator when serving online. In this way, we ensure that its model complexity is comparable to the online model without adding additional calculations to the online service.

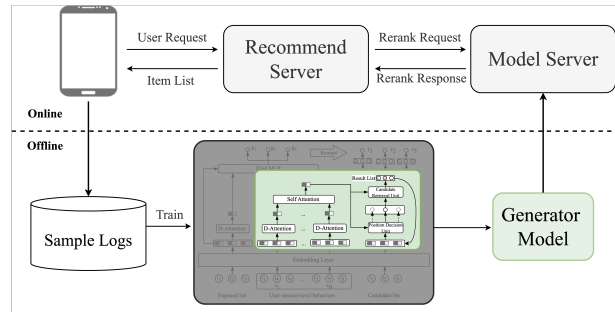


Figure 4: Architecture of the online deployment with NLGR.

We conducted online experiments in an A/B test framework over five weeks from Dec.2023 to Jan.2024. Table 6 shows the online performance of NLGR. Compared to the baseline model (a variant of PRM), NLGR has increased the CTR by 3.25% and the GMV by 3.07%. Moreover, since only NLGR-G is deployed online during online inference, the time-out rate online has no increase, which is acceptable to the recommendation system. Note that NLGR-E is not deployed online which is only utilized for offline guidance of NLGR-G. Now, NLGR has been deployed in the Meituan food delivery platform and serves hundreds of millions of users.

Table 6: Online A/B test result

method	CTR	GMV	Cost (ms)	Time-out
Baseline (PRM)	0.0%	0.0%	0.0	0.0%
NLGR	3.25%	3.07%	1.6	0.0%

6 Conclusion

In this paper, we make the first attempt to solve the goal inconsistency problem in reranking systems. We propose a novel framework called Neighbor List Generative Reranking (NLGR), which uses the relative scores of candidate list and neighboring lists to guide the generator. Furthermore, we propose a sampling-based non-autoregressive generator that can flexibly jump from the current list to any neighbor list. Both offline experiments and online A/B tests show that NLGR significantly outperformed other existing reranking baselines, and we have deployed NLGR on the Meituan food delivery platform.

References

- [1] Qingyao Ai, Keping Bi, Jiafeng Guo, and W Bruce Croft. 2018. Learning a deep listwise context model for ranking refinement. In *The 41st international ACM SIGIR conference on research & development in information retrieval*. 135–144.
- [2] Qingyao Ai, Keping Bi, Jiafeng Guo, and W Bruce Croft. 2018. Learning a deep listwise context model for ranking refinement. In *The 41st international ACM SIGIR conference on research & development in information retrieval*. 135–144.
- [3] Areej Alsini, Du Q Huynh, and Amitava Datta. 2020. Hit ratio: An evaluation metric for hashtag recommendation. *arXiv preprint arXiv:2010.01258* (2020).
- [4] Irwan Bello, Sayali Kulkarni, Sagar Jain, Craig Boutilier, Ed Chi, Elad Eban, Xiyang Luo, Alan Mackey, and Ofer Meshi. 2018. Seq2Slate: Re-ranking and slate optimization with RNNs. *arXiv preprint arXiv:1810.02019* (2018).
- [5] Christopher JC Burges. 2010. From ranknet to lambdarank to lambdamart: An overview. *Learning* 11, 23-581 (2010), 81.
- [6] Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. 2007. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th international conference on Machine learning*. 129–136.
- [7] Jianxin Chang, Chenbin Zhang, Zhiyi Fu, Xiaoxue Zang, Lin Guan, Jing Lu, Yiqun Hui, Dewei Leng, Yanan Niu, Yang Song, et al. 2023. TWIN: TTwo-stage interest network for lifelong user behavior modeling in CTR prediction at kuaishou. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 3785–3794.
- [8] Chi Chen, Hui Chen, Kangzhi Zhao, Junsheng Zhou, Li He, Hongbo Deng, Jian Xu, Bo Zheng, Yong Zhang, and Chunxiao Xing. 2022. EXTR: Click-Through Rate Prediction with Externalities in E-Commerce Sponsored Search. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2732–2740.
- [9] Qiwei Chen, Changhua Pei, Shanshan Lv, Chao Li, Junfeng Ge, and Wenwu Ou. 2021. End-to-end user behavior retrieval in click-through rate prediction model. *arXiv preprint arXiv:2108.04468* (2021).
- [10] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishu Aradhye, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, et al. 2016. Wide & deep learning for recommender systems. In *Proceedings of the 1st workshop on deep learning for recommender systems*. 7–10.
- [11] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*. 191–198.
- [12] Yufei Feng, Yu Gong, Fei Sun, Junfeng Ge, and Wenwu Ou. 2021. Revisit recommender system in the permutation prospective. *arXiv preprint arXiv:2102.12057* (2021).
- [13] Yufei Feng, Binbin Hu, Yu Gong, Fei Sun, Qingwen Liu, and Wenwu Ou. 2021. GRN: Generative Rerank Network for Context-wise Recommendation. *arXiv preprint arXiv:2104.00860* (2021).
- [14] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: a factorization-machine based neural network for CTR prediction. *arXiv preprint arXiv:1703.04247* (2017).
- [15] Iris AM Huijben, Wouter Kool, Max B Paulus, and Ruud JG Van Sloun. 2022. A review of the gumbel-max trick and its extensions for discrete stochasticity in machine learning. *IEEE transactions on pattern analysis and machine intelligence* 45, 2 (2022), 1353–1371.
- [16] Eugene Ie, Vihan Jain, Jing Wang, Sanmit Narvekar, Ritesh Agarwal, Rui Wu, Heng-Tze Cheng, Tushar Chandra, and Craig Boutilier. 2019. SlateQ: A tractable decomposition for reinforcement learning with recommendation sets. (2019).
- [17] Eric Jang, Shixiang Gu, and Ben Poole. 2016. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144* (2016).
- [18] Ray Jiang, Sven Gowal, Timothy A Mann, and Danilo J Rezende. 2018. Beyond greedy ranking: Slate optimization via list-CVAE. *arXiv preprint arXiv:1803.01682* (2018).
- [19] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
- [20] Xiao Lin, Xiaokai Chen, Chenyang Wang, Hantao Shu, Linfeng Song, Biao Li, and Peng Jiang. 2024. Discrete conditional diffusion for reranking in recommendation. In *Companion Proceedings of the ACM on Web Conference 2024*. 161–169.
- [21] Shuchang Liu, Qingpeng Cai, Zhankui He, Bowen Sun, Julian McAuley, Dong Zheng, Peng Jiang, and Kun Gai. 2023. Generative flow network for listwise recommendation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1524–1534.
- [22] Shuchang Liu, Fei Sun, Yingqiang Ge, Changhua Pei, and Yongfeng Zhang. 2021. Variation control and evaluation for generative slate recommendations. In *Proceedings of the Web Conference 2021*. 436–448.
- [23] Xiangyu Liu, Chuan Yu, Zhilin Zhang, Zhenzhe Zheng, Yu Rong, Hongtao Lv, Da Huo, Yiqing Wang, Dagui Chen, Jian Xu, et al. 2021. Neural auction: End-to-end learning of auction mechanisms for e-commerce advertising. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3354–3364.
- [24] Mark F. Medress, Franklin S Cooper, Jim W. Forgie, CC Green, Dennis H. Klatt, Michael H. O'Malley, Edward P Neuburg, Allen Newell, DR Reddy, B Ritea, et al. 1977. Speech understanding systems: Report of a steering committee. *Artificial Intelligence* 9, 3 (1977), 307–316.
- [25] Liang Pang, Jun Xu, Qingyao Ai, Yanyan Lan, Xueqi Cheng, and Jirong Wen. 2020. Setrank: Learning a permutation-invariant ranking model for information retrieval. In *Proceedings of the 43rd international ACM SIGIR conference on research & development in information retrieval*. 499–508.
- [26] Changhua Pei, Yi Zhang, Yongfeng Zhang, Fei Sun, Xiao Lin, Hanxiao Sun, Jian Wu, Peng Jiang, Junfeng Ge, Wenwu Ou, et al. 2019. Personalized re-ranking for recommendation. In *Proceedings of the 13th ACM conference on recommender systems*. 3–11.
- [27] Qi Pi, Guorui Zhou, Yujing Zhang, Zhe Wang, Lejian Ren, Ying Fan, Xiaoqiang Zhu, and Kun Gai. 2020. Search-based user interest modeling with lifelong sequential behavior data for click-through rate prediction. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2685–2692.
- [28] Yuxin Ren, Qiya Yang, Yichun Wu, Wei Xu, Yalong Wang, and Zhiqiang Zhang. 2024. Non-autoregressive generative models for reranking recommendation. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 5625–5634.
- [29] Xiaowen Shi, Fan Yang, Ze Wang, Xiaoxu Wu, Muzhi Guan, Guogang Liao, Wang Yongkang, Xingxing Wang, and Dong Wang. 2023. PIER: Permutation-Level Interest-Based End-to-End Re-ranking Framework in E-commerce. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4823–4831.
- [30] A Vaswani. 2017. Attention is all you need. *Advances in Neural Information Processing Systems* (2017).
- [31] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. 2015. Pointer networks. *Advances in neural information processing systems* 28 (2015).
- [32] Yunjia Xi, Weiwen Liu, Xinyi Dai, Ruiming Tang, Weinan Zhang, Qing Liu, Xiuqiang He, and Yong Yu. 2021. Context-aware reranking with utility maximization for recommendation. *arXiv preprint arXiv:2110.09059* (2021).
- [33] Yunjia Xi, Weiwen Liu, Jieming Zhu, Xilong Zhao, Xinyi Dai, Ruiming Tang, Weinan Zhang, Rui Zhang, and Yong Yu. 2022. Multi-Level Interaction Reranking with User Behavior History. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1336–1346.
- [34] Yueqi Xie, Peilin Zhou, and Sunghun Kim. 2022. Decoupled side information fusion for sequential recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1611–1621.
- [35] 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.
- [36] Tao Zhuang, Wenwu Ou, and Zhirong Wang. 2018. Globally optimized mutual influence aware ranking in e-commerce search. *arXiv preprint arXiv:1805.08524* (2018).
- [37] Tao Zhuang, Wenwu Ou, and Zhirong Wang. 2018. Globally optimized mutual influence aware ranking in e-commerce search. *arXiv preprint arXiv:1805.08524* (2018).