

Ada-Ranker: A Data Distribution Adaptive Ranking Paradigm for Sequential Recommendation

Xinyan Fan
xinyan.fan@ruc.edu.cn
Gaoling School of Artificial
Intelligence, Renmin University of
China
Beijing Key Laboratory of Big Data
Management and Analysis Methods

Jianxun Lian
jianxun.lian@microsoft.com
Microsoft Research Asia

Wayne Xin Zhao[✉]
batmanfly@gmail.com
Gaoling School of Artificial
Intelligence, Renmin University of
China
Beijing Key Laboratory of Big Data
Management and Analysis Methods
Beijing Academy of Artificial
Intelligence

Zheng Liu
zhengliu@microsoft.com
Microsoft Research Asia

Chaozhuo Li
cli@microsoft.com
Microsoft Research Asia

Xing Xie
xingx@microsoft.com
Microsoft Research Asia

ABSTRACT

A large-scale recommender system usually consists of *recall* and *ranking* modules. The goal of ranking modules (aka *rankers*) is to elaborately discriminate users' preference on item candidates proposed by recall modules. With the success of deep learning techniques in various domains, we have witnessed the mainstream rankers evolve from traditional models to deep neural models. However, the way that we design and use rankers remains unchanged: offline training the model, freezing the parameters, and deploying it for online serving. Actually, the candidate items are determined by specific user requests, in which underlying distributions (e.g., the proportion of items for different categories, the proportion of popular or new items) are highly different from one another in a production environment. The classical parameter-frozen inference manner cannot adapt to dynamic serving circumstances, making rankers' performance compromised.

In this paper, we propose a new training and inference paradigm, termed as *Ada-Ranker*, to address the challenges of dynamic online serving. Instead of using parameter-frozen models for universal serving, *Ada-Ranker* can adaptively modulate parameters of a ranker according to the data distribution of the current group of item candidates. We first extract distribution patterns from the item candidates. Then, we modulate the ranker by the patterns to make the ranker adapt to the current data distribution. Finally, we use the revised ranker to score the candidate list. In this way, we empower the ranker with the capacity of adapting from a global model to a local model which better handles the current task. As a first study,

we examine our *Ada-Ranker* paradigm in the sequential recommendation scenario. Experiments on three datasets demonstrate that *Ada-Ranker* can effectively enhance various base sequential models and also outperform a comprehensive set of competitive baselines.

CCS CONCEPTS

• **Information systems** → **Learning to rank; Recommender systems**; • **Computing methodologies** → *Neural networks*.

KEYWORDS

Model Adaptation, Sequential Recommendation, Dynamic Ranking

ACM Reference Format:

Xinyan Fan, Jianxun Lian, Wayne Xin Zhao, Zheng Liu, Chaozhuo Li and Xing Xie. 2022. Ada-Ranker: A Data Distribution Adaptive Ranking Paradigm for Sequential Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*, July 11–15, 2022, Madrid, Spain. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3477495.3531931>

1 INTRODUCTION

Recommender systems play an important role in information filtering for online services such as e-commerce, news, movie, and gaming. To support efficient recommendations from a massive set of items, an industrial recommender system usually follows the *recall-then-rank* two-stage paradigm. Given a user request, the recall component proposes a small set of relevant candidates with lightweight methods, then the ranking component further elaborately scores the candidates with more advanced models and returns top-*k* results. This paper discusses models in the ranking component (which we call *rankers* hereafter).

Since online user behaviors are highly dynamic (driven by evolving user preference and short-term interest), enhancing rankers with sequential user modeling becomes a hot research topic and shows significant business value in industry [25, 50]. In recent literature, a number of sequential recommendation models have been proposed based on neural network architectures, including Gated Recurrent Unit (GRU) [8], Convolutional Neural Network (CNN) [21] and Transformers [36]. These approaches can effectively handle

[✉] Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGIR '22, July 11–15, 2022, Madrid, Spain.

© 2022 Association for Computing Machinery.
ACM ISBN 978-1-4503-8732-3/22/07...\$15.00
<https://doi.org/10.1145/3477495.3531931>

sequential interaction data and train the recommendation models in an end-to-end manner. Typically, these neural rankers still follow the “*regular offline training* \Rightarrow *static online inference*” paradigm, in which once a model is trained, the parameters are frozen and deployed into online environment for instant service.

However, the candidate lists to be ranked are determined by specific user requests, and the underlying distributions can be rather different in diverse request scenarios. For example, in news recommendations, for users who are big fans of NBA, their recall candidates will contain more news articles related to basketball than other users; on an e-commerce platform, if a user clicks on the *Women’s Fashion* category, then the retrieved candidates will be highly related to this category. Besides, other factors (e.g., temporal effects [39] and counterfactual settings [10]) might also increase the discrepancy among the data distributions of candidate items for different user requests. Therefore, the parameter-frozen inference paradigm for rankers can only produce sub-optimal performances since data distribution discrepancy will cause significant performance decrease in the rankers. Indeed, such an issue generally exists in supervised learning methods [31, 34]. It is desirable to develop a more capable ranking paradigm that can flexibly adapt to different circumstances.

As existing solutions, several methods curate a few online circumstances (e.g., for different weekdays [39] or different domains [33]), and then let the model dynamically select a suitable one for response. However, these approaches can only cover a limited set of recommendation scenarios by assuming they are pre-given, and we need to maintain multiple models in the service system, which is not parameter-efficient. Besides, it takes a significant cost to switch among different heavy model variants, which also increases the risk of system failures. Therefore, it is not realistic to assume that the circumstances are discrete and enumerable. It is still challenging to design a capable ranker that dynamically adapts to a specific circumstance in a flexible and parameter-efficient manner.

In this paper, we propose **Ada-Ranker** – a new *Adaptive* paradigm for making context-aware inference in *Rankers*. We address the data distribution shift issue by a *model adaptation* approach. Specifically, we regard handling each set of item candidates from recall modules to a specific user request as an individual *task*. During the inference stage for a task, instead of fixing the parameters of a ranker, Ada-Ranker adapts the ranker to the current task in three steps, namely *distribution learning*, *input modulation* and *parameter modulation*. For distribution learning, we learn the underlying data patterns of the current task by Neural Processes [13, 14], which will help the ranker focus on extracting useful users behavior patterns to better discriminate candidate items in current task. For input modulation, by taking the extracted data distribution patterns as adaptation conditions, we learn a feature-wise linear modulation neural network to adjust the input representations, so that input representations are re-located to latent positions where rankers can discriminate the current task better. For parameter modulation, we adopt a model patching approach by generating parameter patches based on a parameter pool of base parameter vectors or matrices. Our adaptation process consists of the above input modulation and parameter modulation procedures, which jointly adapt model’s parameters according to specific tasks in a model-agnostic way. To verify the effectiveness of Ada-Ranker, we design comprehensive

settings for test environments and apply Ada-Ranker to various types of base sequential recommendation models on three real-world datasets. Experiment results demonstrate that Ada-Ranker can effectively adapt to dynamic and diverse test tasks, and significantly boost the base models.

Our contributions are summarized as follows. (1) We highlight the data distribution discrepancy issue and the importance of parameter adaptation during the inference process, which is commonly encountered in industrial recommender systems but overlooked in literature. (2) We propose Ada-Ranker, a new inference paradigm for rankers, which can adaptively adjust the model parameters according to the data patterns in a given set of item candidates. Ada-Ranker satisfies the three properties of being lightweight, model-agnostic and flexible for plug-and-play usage. (3) We conduct extensive experiments on three real-world datasets in the scenario of sequential recommendations. Results demonstrate that Ada-Ranker can significantly boost the performance for various base models, as well as outperform a set of competitive baselines.

2 PROBLEM FORMULATION

Given a request from a user u , the recall module uses multiple approaches (such as popularity-based, item-to-item and approximate nearest neighbor search) to retrieve a small set (usually a few hundreds or thousands) of item candidates: $C = \{v_i\}_{i=1}^m$, which might be relevant to u . The goal of the ranker is to score each candidate item v in C and return the (top- k) ordered list as recommendations.

A Standard Sequential Ranker Architecture. For personalized recommendations, user u is associated with a profiling representation, denoted by \mathbf{x}_u , which is derived based on her historical activities: $\mathbf{x}_u = \{v_1, v_2, \dots, v_n\}$, where n is the length of behaviors and v_i is in chronological order. The ranker adopts a model f to predict the preference score of user u over the target candidate v : $\hat{y}_{uv} = f(\mathbf{x}_u, v)$, where there are usually three major layers in f : an embedding lookup layer $Q(\cdot)$, a sequential encoding layer $g^{SE}(\cdot)$, and a predictive layer $g^{PRED}(\cdot)$. We present a typical architecture of ranker models f in Figure 1(a). The user behavior sequence $\{v_1, v_2, \dots, v_n\}$ will first go through the embedding-lookup layer $Q(\cdot)$ to form the corresponding item embeddings:

$$Q_u = Q(\mathbf{x}_u) = \{\mathbf{q}_{v_1}, \mathbf{q}_{v_2}, \dots, \mathbf{q}_{v_n}\}. \quad (1)$$

Next, the sequential user encoder $g^{SE}(\cdot)$ will encode the item sequence and produce an embedding vector as user representation:

$$\mathbf{p}_u = g^{SE}(Q_u). \quad (2)$$

Then, \mathbf{p}_u will be concatenated with a target item vector \mathbf{q}_v (if there are some attribute features such as item category, user profiles, they can also be appended here) as the input to the predictive layer $g^{PRED}(\cdot)$, which is usually implemented as a simple two-layer Multi layer Perceptron (MLP) (see Figure 3 (a)):

$$\hat{y}_{uv} = g^{PRED}(\mathbf{p}_u, \mathbf{q}_v) = MLP(\mathbf{p}_u, \mathbf{q}_v). \quad (3)$$

In summary, a standard ranker model f can be instantiated as

$$\hat{y}_{uv} = f(\mathbf{x}_u, v) = g^{PRED}(g^{SE}(Q(\mathbf{x}_u)), \mathbf{q}_v). \quad (4)$$

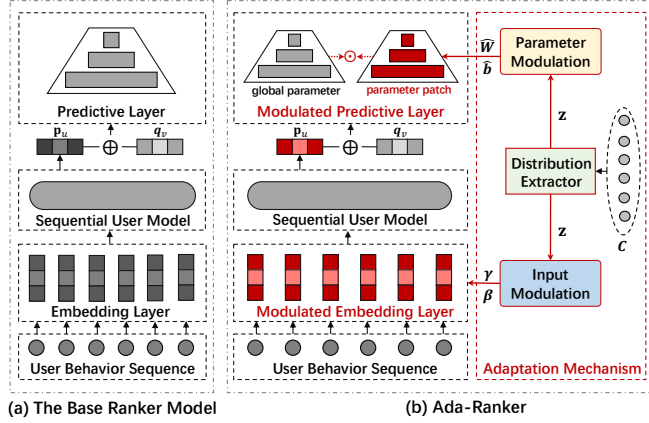


Figure 1: An overview of the traditional sequential model (a) and Ada-Ranker paradigm (b). We use colored elements to indicate the new components in Ada-Ranker.

Potential Issues. However, as introduced in Section 1, the retrieved item candidates C may have different data distributions from diverse recall requests. Existing methods adopt a global ranker f to serve all requests, and for a given user, it will produce the same score for item v , regardless of which candidate set C it draws from. We argue that a better paradigm is to leverage C as a ranking context, and let the model adjust itself according to the specific context to make more fine-grained and accurate prediction scores for the current ranking task.

3 METHODOLOGY

3.1 An Adaptive Ranking Paradigm

The essence of adaptive ranking paradigm is to incorporate a specially-designed adaptation mechanism (aka, adaptor), which encodes the data distribution patterns of C and revises the global model f to a local model f' accordingly, so that f' has a better discriminative capacity of ranking items in C . To implement this new ranking paradigm, it is important to design the adaptation mechanism with special considerations for industrial applications:

- (C1) *Lightweight.* The model should be efficient in both computation and memory usage.
- (C2) *Model-agnostic.* The model can be generalized to various sequential recommendation models, such as GRU-based [15], CNN-based [45] or GNN-based methods [41].
- (C3) *Plug-and-play.* The adaptation mechanism can be applied as a model patch: the global model remains unchanged, so when the adaptor is disabled, the system performs exactly as how it works as usual; when the adaptor is enabled, the global model together with the model patch serve as the adaptation model and produce more accurate results. This feature ensures the flexibility to turn on/off the adaptor for different purposes.

To this end, we introduce a new ranking paradigm *Ada-Ranker*. We first learn the distribution patterns z from C , which will be used to modulate the global model f and derive a local model f' ; then, we use f' to score the candidates in C : $\hat{y}_{uv} = f'(x_u, v)$. Since there are two key components in the global model f , i.e., $g^{SE}(\cdot)$ (Eq. 2) and

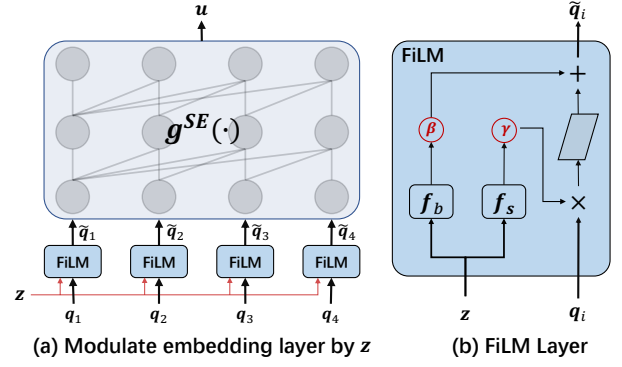


Figure 2: An illustration of the input modulation.

$g^{PRED}(\cdot)$ (Eq. 3), we propose two adaptation components, called *input modulation* and *parameter modulation*, which can incorporate the data distribution z to modulate $g^{SE}(\cdot)$ and $g^{PRED}(\cdot)$ respectively. An overview of Ada-Ranker is illustrated in Figure 1(b) and in Algorithm 1. In the next section, we will introduce the details of each proposed component.

3.2 Data Distribution Learning from the Ranking Candidates

To implement the ranking paradigm, the first issue is how to effectively learn the data distribution patterns in ranking candidates C , which extracts specific data characteristics for adjusting the ranker.

Neural Processes Encoder. We assume that item candidates in C are drawn from a particular instantiation of stochastic process \mathcal{F} , which corresponds to a specific ranking request. In order to characterize the dynamic and diverse data distributions, we borrow the idea of *Neural Processes* (NP) from [13, 14] to approximate a stochastic process with learnable neural networks. Recently, Neural Processes has been utilized for alleviating the user cold-start problem in recommender systems [27]. Different from them, we adopt Neural Processes to model \mathcal{F} which represents the data distribution associated with a ranking request. The fundamental advantages of NP lie in (1) providing an effective way to model data distributions conditioned on representations of observed data; and (2) being parameterized by Multi-layer Perceptron (MLPs) which is more efficient than Gaussian Process (see more details in [13, 14, 27]). Here we introduce the variational approximation implementation of NP encoder [27].

Specifically, we first generate a latent embedding vector \mathbf{r}_j for each item j in C with a two-layer MLP:

$$\mathbf{r}_j = MLP^{(NP)}(q_j). \quad (5)$$

Then, the NP encoder will aggregate these latent vectors to generate a permutation-invariant representation \mathbf{r} via mean pooling:

$$\mathbf{r} = (\mathbf{r}_1 + \mathbf{r}_2 + \dots + \mathbf{r}_m)/m. \quad (6)$$

Reparameterization. The above representation \mathbf{r} will be used to generate the mean vector and the variance vector:

$$\mathbf{s} = ReLU(\mathbf{W}_s \mathbf{r}), \quad (7)$$

$$\boldsymbol{\mu} = \mathbf{W}_\mu \mathbf{s}, \log \boldsymbol{\sigma} = \mathbf{W}_\sigma \mathbf{s}. \quad (8)$$

Finally, the data distribution is modeled by a random variable \mathbf{z} : $\mathbf{z} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}^2)$, which is implemented with the reparameterization trick [19], so that via making randomness in the input node of model, all the computational nodes in the model are differentiable and gradients can be smoothly backpropagated:

$$\mathbf{z} = \boldsymbol{\mu} + \boldsymbol{\epsilon} \odot \boldsymbol{\sigma}, \quad \boldsymbol{\epsilon} \sim \mathcal{N}(0, \mathbf{I}), \quad (9)$$

where \odot means the element-wise product operation.

After extracting the data distribution from the candidates, we next discuss how to adapt a ranker to the given data distribution in two aspects, namely *input modulation* and *parameter modulation*.

3.3 Input Modulation

To ensure the model-agnostic property, we modulate the latent representations of input sequence, and this operation can be shared by different sequential recommendation models.

Modeling Data Distributions as Adaptation Conditions. To revise the item representations according to the underlying data distributions, our idea is to consider the extracted \mathbf{z} in Eq. 9 as a condition to adjust the corresponding input representations. Here, we adopt the FiLM method [30] that is a general-purpose method to model the influence of conditioning information on neural networks. Conditioned on the distribution patterns \mathbf{z} , we perform linear modulation on the latent representations of item sequence, such that input embeddings are adjusted to new representations that are more distinguishable among candidate set \mathcal{C} . The details of this adaptation process for input modulation are shown in Figure 2.

Specifically, We learn to generate two modulation coefficients γ and β through the conditional representation \mathbf{z} :

$$\gamma = f_s(\mathbf{z}), \quad \beta = f_b(\mathbf{z}), \quad (10)$$

where f_s and f_b are two neural networks with different parameters, formulated as: $f(\mathbf{z}) = \mathbf{w}_2 \text{ReLU}(\mathbf{W}_1 \mathbf{z} + \mathbf{b}_1) + b_2$. The latent representations of item sequence are then adjusted by:

$$\tilde{\mathbf{q}}_t = \gamma \mathbf{q}_t + \beta. \quad (11)$$

Here, f_s and f_b are shared for all items in behavior sequence.

A Case with GRU-based Recommender. For example, if the sequential model $g^{SE}(\cdot)$ is GRU, then the modulated user vector is: $\tilde{\mathbf{u}} = \text{GRU}(\{\tilde{\mathbf{q}}_t\}_{t=1}^n)$. Ada-Ranker does not change any parameters in GRU (due to the model-agnostic property). Instead, it makes adaptations on the data input, which creates room for the encoded user vector to be more dedicated to a provided distribution \mathbf{z} .

3.4 Parameter Modulation

For the modulation of scoring function $g^{PREL}(\cdot)$, we adopt the idea of model patch [17, 33] to adjust the parameters, which is a light-weight approach to implement parameter adaptation.

Adaptation by Model Patch. The basic idea of model patch is to learn parameter variations for adapting to new input or task context [44]. As shown in Figure 3 (a)-(b), we can incorporate another two MLPs to generate *parameter patches* for the k -th hidden layer of the predictive layer according to the conditional representation

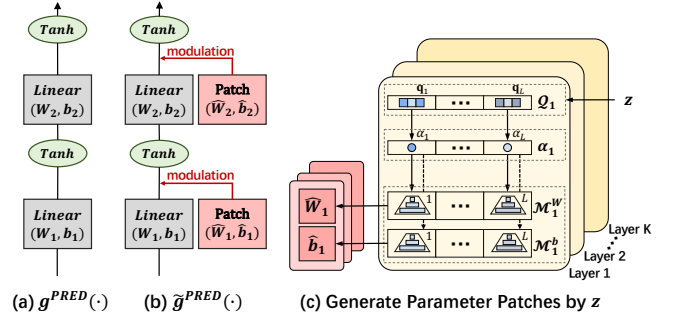


Figure 3: An illustration of the parameter modulation.

\mathbf{z} which represents the underlying data distributions:

$$\widehat{\mathbf{W}}_k = \text{MLP}^{(w_k)}(\mathbf{z}), \quad \widehat{\mathbf{b}}_k = \text{MLP}^{(b_k)}(\mathbf{z}), \quad (12)$$

Then the ranker's MLP parameters (Eq. 3) are modulated by:

$$\widetilde{\mathbf{W}}_k = \mathbf{W}_k \odot \widehat{\mathbf{W}}_k, \quad \widetilde{\mathbf{b}}_k = \mathbf{b}_k + \widehat{\mathbf{b}}_k, \quad (13)$$

where \odot denotes the element-wise multiplication and k means the k -th layer of the ranker's MLP. Let h denote the input dimension and d denote the output dimension, we have $\mathbf{W}_k, \widehat{\mathbf{W}}_k, \widetilde{\mathbf{W}}_k \in \mathbb{R}^{h \times d}$ and $\mathbf{b}_k, \widehat{\mathbf{b}}_k, \widetilde{\mathbf{b}}_k \in \mathbb{R}^d$. We will replace the original MLPs' parameters \mathbf{W}_k by $\widetilde{\mathbf{W}}_k$ and \mathbf{b}_k by $\widetilde{\mathbf{b}}_k$. In this way, a global scoring function $g^{PREL}(\cdot)$ is adapted into a local scoring function $\tilde{g}^{PREL}(\cdot)$.

Learning with Parameter Pool. A major problem of Eq. (12) is that the modulation parameters are generated in an unconstrained way, and it has been shown that it is difficult to directly optimize MLPs to approximate arbitrary real-valued vectors, especially when the training sample is not sufficient [22]. Thus, we propose to construct a parameter pool consisting of multiple base parameters (either vector or matrix), which is a parameter memory network [32]. Furthermore, the parameters are derived based on a linear combination of multiple base parameters. The proposed adaptation mechanism for parameter modulation is depicted in Figure 3 (c). Specifically, we have a parameter pool \mathcal{M} with L base parameters: $\mathcal{M} = \{\mathbf{M}_1, \mathbf{M}_2, \dots, \mathbf{M}_L\}$, where each \mathbf{M}_i has the same shape with the MLP parameter \mathbf{W} of the ranker. Then, the parameter patch $\widehat{\mathbf{W}}$ is composed by a linear combination of \mathcal{M} as follows:

$$\widehat{\mathbf{W}} = \sum_{i=1}^L \alpha_i \mathbf{M}_i, \quad (14)$$

The coefficients α_i are determined by a set of reading heads $\mathcal{Q} = \{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_L\}$:

$$a_i = \mathbf{z} \cdot \mathbf{q}_i, \quad (15)$$

$$\alpha_i = \frac{\exp(a_i)}{\sum_{j=1}^L \exp(a_j)}, \quad \forall i = 1, 2, \dots, L, \quad (16)$$

where \mathbf{q}_* are learnable vectors which have the same shape with \mathbf{z} . We use individual parameter pools $\mathcal{M}_k, \mathcal{Q}_k$ for each layer of the MLP parameters in $g^{PREL}(\cdot)$ (e.g., \mathbf{W}_k in the k -th layer). The case of the bias parameters \mathbf{b}_* is the same, which is omitted here.

3.5 Optimization and Discussion

After the input modulation and parameter modulation, the ranker transforms from a global model $f(u, v) = g^{PREL}(g^{SE}(Q(\mathbf{x}_u)), \mathbf{q}_v)$ to a local model $f'(u, v) = \tilde{g}^{PREL}(\tilde{g}^{SE}(Q(\mathbf{x}_u)), \mathbf{q}_v)$. We next discuss how to optimize the proposed ranking approach.

Parameter Learning. The optimization method remains the same as the original ranker, by minimizing the binary cross entropy loss:

$$\mathcal{L} = \frac{1}{|\Gamma|} \sum_{(u,v) \in \Gamma} -y_{u,v} \log \tilde{y}_{uv} - (1 - y_{u,v}) \log(1 - \tilde{y}_{uv}), \quad (17)$$

where $\tilde{y}_{uv} = f'(u, v)$ and Γ denotes the set of training instances. There are several feasible strategies to train the Ada-Ranker. For clarity, let Θ denote the set of parameters of the base model f , and Φ denote the set of parameters in adaptation modules (including the distribution extractor, input modulation and parameter modulation). We have the following three alternative training strategies:

- (1) $\emptyset \Rightarrow \Theta + \Phi$. A straightforward approach is to train the whole Ada-Ranker model in an end-to-end manner from scratch (\emptyset).
- (2) $\Theta \Rightarrow \Theta + \Phi$. In the first stage, we train the base model until convergence, so only Θ get updated in this stage; in the second stage, we load the pre-trained base model and train the whole Ada-Ranker, so parameters in both Θ and Φ will be updated.
- (3) $\Theta \Rightarrow \Phi$. In the first stage, we train the base model until it converges, so only Θ get updated in this stage; in the second stage, we load the pre-trained base model f and freeze its parameters, and train the Ada-Ranker by only updating Φ .

To achieve the (C3) property mentioned in Section 3.1, we adopt the last training strategy. Pre-training the base model’s parameters can offer a good parameter initialization for the training of Ada-Ranker, so the first strategy, $\emptyset \Rightarrow \Theta + \Phi$, is usually worse than the other two strategies. The major merit of the third strategy is that, after Ada-Ranker’s training, the base model’s parameters remain unchanged, such that an industrial system can flexibly turn on or turn off the adaptation module to meet diverse requirements of different scenarios. The overall training process of Ada-Ranker is summarized in Algorithm 1 (see experimental comparisons of the training strategies in Section 4.4.2). For inference stage, we only need to run Line 8 to Line 17 in Algorithm 1 with a trained Ada-Ranker model.

Complexity Discussion. The number of additional parameters Φ is $O(LK(hd + d) + dd)$, while the original global model has parameter size $|\Theta| = O(Md + K(hd + d) + |\theta|)$, where L, K, h, d, M, θ denote the number of parameter pool slots, the number of hidden layers of the predictive layer, hidden layer size for neural networks, item embedding size, the number of items and the parameters of $g^{SE}(\cdot)$ respectively. For a recommender system, M (usually thousands or millions) is much larger than L (e.g., 10 in this paper), K (e.g., 2 in this paper) and d (e.g., 64 in this paper), compared with other methods that need to train and maintain multiple versions of Θ for serving different scenarios, our Ada-Ranker is much more parameter economic. As for the computational cost, Neural Processes perform the calculation on candidate items individually (see Eq.(5)), without the requirement of interactions among items in the candidate list. Thus, with distribution learning techniques, the computational cost is linear with the number of items in the candidate

Algorithm 1 The training procedure of Ada-Ranker.

Input: Training instances: a set of {behavior sequence $\mathbf{x}_u = \{v_1, v_2, \dots, v_n\}$, labeled items $C = \{v_i, y_{uv_i}\}_{i=1}^m$ }.
Output: An Ada-Ranker model $f'(\mathbf{x}_u, v)$.

- 1: // The first training stage
- 2: Train base model $f(\mathbf{x}_u, v)$ and obtain pre-trained global parameters Θ .
- 3: // The second training stage
- 4: Initialize base model parameters in Ada-Ranker $f'(\mathbf{x}_u, v)$ by Θ , initialize adaptation parameters Φ randomly.
- 5: **while** not convergence **do**
- 6: Fetch a training instance $\{\mathbf{x}_u, C\}$
- 7: // Distribution Extraction
- 8: Extract the distribution pattern \mathbf{z} by Eq.(5 - 9).
- 9: // Input Modulation
- 10: Generate γ, β with \mathbf{z} by Eq.(10).
- 11: Use γ, β to modulate each \mathbf{q}_t by Eq.(11) to get $\tilde{\mathbf{q}}_t$.
- 12: Encode $\{\tilde{\mathbf{q}}_t\}_{t=1}^n$ and generate $\tilde{\mathbf{u}}$ by $g^{SE}(\cdot)$.
- 13: // Parameter Modulation
- 14: Calculate the coefficients α by \mathcal{Q}_k and \mathbf{z} in Eq.(15)&(16).
- 15: Generate parameters $\widehat{\mathbf{W}}_k, \widehat{\mathbf{b}}_k$ in Eq.(14).
- 16: Modulate g^{PREL} parameters $\mathbf{W}_k, \mathbf{b}_k$ in Eq.(13).
- 17: Use the revised \tilde{g}^{PREL} to score the candidates in C .
- 18: Calculate loss \mathcal{L} in Eq.(17).
- 19: Stop gradients for Θ and update Φ by Adam optimizer.

Table 1: Statistics of datasets after preprocessing.

Dataset	# Users	# Items	# Actions	# Avg.len	Sparsity	#Categories
ML10M	69,878	10,676	10,000,047	143	98.66%	18
Taobao	487,813	787,094	24,938,811	34	99.99%	20
Xbox	1,000,000	6,071	25,315,983	15	99.58%	12

list. Besides, for the input modulation and parameter modulation, the forward calculation does not depend on a specific candidate item, so the whole calculation only needs to be conducted once. Thus, the overall computational cost can be effectively controlled at a reasonable level.

4 EXPERIMENT

4.1 Experimental Settings

4.1.1 Datasets. We use three datasets for experiments, covering movie, e-commerce and gaming recommendation scenarios.

- **ML10M Dataset**¹ is a widely used benchmark dataset. It contains 10,000,054 ratings and 95,580 tags applied to 10,681 movies by 71,567 users from an online movie recommender service. There are in total 18 categories (genres in origin file) of items (with overlapping items): *Action, Adventure, Animation, Children, Comedy, Crime, Documentary, Drama, Fantasy, Film-Noir, Horror, Musical, Mystery, Romance, Sci-Fi, Thriller, War* and *Western*.
- **Taobao Dataset**² is an e-commerce dataset released by Taobao. The user-item interactions logs include the type of interaction (e.g., purchase or click), timestamps and items’ categories (in anonymous). The item category distribution is severely uneven, so we keep the top 19 categories which contain most items, and regard the remaining items belong to a virtual category *others*.

¹<https://grouplens.org/datasets/movielens/10m/>

²<https://tianchi.aliyun.com/dataset/dataDetail?dataId=649&userId=1>

Table 2: Performance comparison of seven base models and Ada-Ranker on three datasets.

Datasets	Models	MF		GRU4Rec		SASRec		NARM		NextIttNet		SHAN		SRGNN	
		GAUC	NDCG	GAUC	NDCG	GAUC	NDCG	GAUC	NDCG	GAUC	NDCG	GAUC	NDCG	GAUC	NDCG
ML10M	Base	0.8683	0.6942	0.9187	0.7781	0.9106	0.7616	0.9156	0.7718	0.8620	0.6855	0.8667	0.6973	0.8993	0.7441
	Ada-Ranker	0.8807	0.7121	0.9279	0.7932	0.9214	0.7783	0.9261	0.7879	0.8778	0.7069	0.8816	0.7164	0.9161	0.7676
Taobao	Base	0.8363	0.6605	0.8734	0.7208	0.8707	0.7211	0.8707	0.7177	0.8482	0.6872	0.8609	0.7094	0.8637	0.7105
	Ada-Ranker	0.8512	0.6907	0.8888	0.7490	0.8864	0.7481	0.8881	0.7439	0.8625	0.7079	0.8653	0.7106	0.8791	0.7322
Xbox	Base	0.9036	0.7676	0.9388	0.8302	0.9323	0.8172	0.9368	0.8258	0.9343	0.8222	0.9217	0.8022	0.9336	0.8212
	Ada-Ranker	0.9123	0.7897	0.9451	0.8438	0.9382	0.8291	0.9412	0.8354	0.9401	0.8336	0.9282	0.8131	0.9426	0.8384

- **Xbox Dataset** is a private dataset provided by Xbox, an online gaming platform of Microsoft. It contains one year’s user-game playing logs spanning from September 2020 to September 2021. We sample 1 million users from the original dataset. There are 12 major categories of items in Xbox dataset (with overlapping items), such as *GamePass*³, *Action Adventure* and *Card Board*.

For all datasets, we remove users with fewer than ten interaction records, group the interaction records by users and sort them in chronological order. The basic statistics are reported in Table 1. Following previous works [18, 51], we apply the *leave-one-out* strategy for evaluation. For each user interaction sequence, the last item and the item before it are considered as test and validation data, respectively, and the remaining items is for training. A group of candidate items contains 1 positive instance and 19 negative samples.

Since the focus of this paper is to empower rankers with adaptation ability for dynamic ranking context, to verify this property, the key part of dataset preparation lies in the negative sampling strategy, which reflects that the ranker needs to discriminate positive instance from different groups of item candidates. For example, if instances are sampled by item popularity, we can simulate the situation that the user is visiting the *Most Popular* tab; if instances are all sampled from the *Shooter* category, it indicates that a user has clicked into the *Shooter* channel and searches for games in this scope. Thus, to make the test environment more comprehensive and the distribution of candidates more diverse, we sample negative items for each positive instance by the *distribution-mixer sampling*.

Distribution-mixer sampling: We consider two basic factors that influence the distribution of candidates, namely (1) item categories and (2) item popularity. Our negative sampling strategies are designed as: (Step-1) draw a random number d from $\{1, 2, 3\}$, indicating that d item categories will be involved in the 19 negative instances. If $d > 1$, then randomly sample $d - 1$ item categories which are different with the positive item’s category; (Step-2) conduct a Bernoulli trial with $p = 0.5$, if the outcome is *success*, set the item sampling strategy to *popularity-bias*, otherwise set the item sampling strategy to *uniform* for this round; (Step-3) from each d sampled categories, sample items with strategy returned in Step-2, so that the total number of items is 19. Step-1 to Step-3 is repeated individually for each positive instance in the train/valid/test set.

4.1.2 Evaluation Metrics. Following the previous works [48, 49], we adopt two widely used metrics for evaluation the ranking results:

GAUC (Area Under the ROC curve, group by each user), and NDCG (Normalized Discounted Cumulative Gain).

4.1.3 Implementation Details. All methods are implemented with Python 3.8 and PyTorch 1.8.1. We use four Linux servers with the same hardware configuration: CPU is *Intel(R) Xeon(R) E5-2690 v4 @ 2.60GHz* and GPU is *Tesla P100*. We use the Adam optimizer with a learning rate of 0.001, where the batch size is set as 4096. The dropout rate is set to 0.4. The dimension of the embedding is 64, and the hidden states’ size in the prediction layer is $128 * 64$ and $64 * 1$. The maximum sequence length is 200, 100 and 50 for ML10M, Taobao and Xbox datasets, respectively. For Ada-Ranker, the slots of memory network L is set to 10. The source code of Ada-Ranker is released at <https://github.com/RUCAIBox/Ada-Ranker>.

4.2 Improving Various Base Models

4.2.1 Base Models. Ada-Ranker is a model-agnostic ranking paradigm which can enhance various types of sequential recommendation models. To verify this, we test Ada-Ranker’s performance based on the following representative sequential models:

- **MF** [20] factorizes each item and user into an embedding vector based on the user-item interaction logs. It is the only one base model that is non-sequential.
- **GRU4Rec** [15] uses GRU to model a user’s behavior sequence. It is one of the most popular model for sequential recommendations.
- **SASRec** [18] uses the unidirectional multi-head self-attention model to deal with a user’s behavior sequence.
- **NARM** [24] is a session-based recommendation approach with RNN. It uses an attention mechanism to determine the relatedness of the past purchases in the session for the next purchase.
- **NextIttNet** [45] applies dilated CNN for long sequence modeling.
- **SHAN** [42] is a 2-layer hierarchical attention network to model user’s dynamic long-term preference and sequential behaviors.
- **SRGNN** [41] models session sequences as graph structured data and uses a GNN to capture complex transitions of items.

4.2.2 Results. We train all base sequential models and their corresponding Ada-Rankers on three datasets. The results are reported in Table 2. We observe that Ada-Ranker outperforms all base sequential models substantially and consistently, on all datasets, and in terms of overall metrics. Specifically, the average improvement of Ada-Ranker (GRU4Rec) over GRU4Rec on the three datasets is 2.50% in terms of NDCG and the improvement of Ada-Ranker (SAS-Rec) over SASRec is 2.46%. Since Ada-Ranker and its base model

³GamePass is membership subscription, with which users can play a library of games on Xbox by paying a certain amount of money each month.

Table 3: Performance comparison with eight different baselines based on two base models (GRU4Rec and SASRec). The best performance and the second best performance methods are denoted in bold and underlined, respectively. The p -value of significance test is performed in the NDCG metrics of Ada-Ranker with the corresponding best baseline method (underlined).

Methods	GRU4Rec as base model						SASRec as base model					
	ML10M		Taobao		Xbox		ML10M		Taobao		Xbox	
	GAUC	NDCG	GAUC	NDCG	GAUC	NDCG	GAUC	NDCG	GAUC	NDCG	GAUC	NDCG
Base	0.9187	0.7781	0.8734	0.7208	0.9388	0.8302	0.9106	0.7616	0.8707	0.7211	0.9323	0.8172
DNS	0.9005	0.7512	0.8646	0.7200	0.9254	0.8182	0.8942	0.7411	0.8632	0.7183	0.9219	0.8125
GSF _{concat}	0.9150	0.7744	<u>0.8903</u>	<u>0.7433</u>	0.9335	0.8261	0.9072	0.7623	0.8858	<u>0.7400</u>	0.9308	0.8201
GSF _{avg}	0.9205	0.7760	0.8860	0.7370	0.9390	0.8325	0.9116	0.7622	0.8798	0.7314	0.9353	0.8270
GSF _{trm}	0.9232	0.7815	0.8865	0.7366	0.9429	<u>0.8414</u>	0.9161	0.7700	0.8836	0.7351	0.9366	0.8278
PD	<u>0.9254</u>	0.7818	0.8913	0.7387	0.9373	0.8216	<u>0.9196</u>	0.7705	0.8878	0.7323	0.9325	0.8137
DecRS	<u>0.9247</u>	<u>0.7881</u>	0.8785	0.7306	0.9427	0.8367	<u>0.9195</u>	0.7699	0.8776	0.7299	0.9345	0.8227
DLCM	0.9240	0.7851	0.8835	0.7366	0.9426	0.8374	0.9177	0.7713	0.8793	0.7347	0.9376	0.8272
PRM	0.9244	0.7863	0.8829	0.7397	<u>0.9442</u>	0.8396	0.9181	<u>0.7731</u>	0.8788	0.7345	<u>0.9379</u>	<u>0.8288</u>
AdaRanker	0.9279	0.7932	0.8888	0.7490	0.9451	0.8438	0.9214	0.7783	<u>0.8864</u>	0.7481	0.9382	0.8291
p -value	7.02e-07		7.64e-09		0.00424		0.000215		5.39e-12		0.1586	

have the same embedding layer, user model architectures and prediction layers, it shows that our proposed Ada-Ranker is effective to adapt a base mode according to different distributions of candidates for the current task. Note that we adopt the $\Theta \Rightarrow \Phi$ strategy (see Section 3.5) to train Ada-Ranker, which first trains the base model and then only fine-tune parameters of adaptation networks. The impressive performance gain demonstrates the effectiveness of our adaptation mechanism, even though we just fine-tune the small set of parameters of adaptation networks. In summary, Ada-Ranker processes the virtue of the plug-and-play property: enhancing any given base model with adaptation modules while maintaining the original base model unchanged.

4.3 Comparison with Baselines

4.3.1 Baselines. To the best of our knowledge, there are no studies on dynamic model adaptation for rankers in literature. However, we can extend some existing methods to make them able to handle dynamic candidate sets, so that they can serve as baselines. Generally, we collect baselines from four perspectives and adopt GRU4Rec and SASRec as base models to test how different types of methods perform.

(1) From the view of dynamic negative sampling:

- **DNS** [46] dynamically chooses negative samples from the ranked list produced by the current prediction model and iteratively update it. Considering that DNS dynamically constructs candidate list in the training stage, it is useful to adapt the base model to various types of data distribution.

(2) From the view of group-wise ranking:

- **GSF** [2] uses a groupwise scoring function, in which the relevance score of an item is determined jointly by all items in the candidate list. We implement three types of scoring functions based on the idea of GSF. (1) GSF_{concat}: concatenates all candidates' embeddings and produces scores for all candidates, which is the original implementation of GSF [2]. (2) GSF_{avg}: averages all candidates' embeddings into one embedding vector (which can be regarded as a context vector), and appends it to the input of scoring function. (3)

GSF_{trm}: instead of simply using averaging, it uses a 2-layer Transformer to encoding the candidate list and get all items' contextualized representation, then append the corresponding contextualized representation to the input of scoring function.

(3) From the view of causal analysis:

- **PD** [47] removes the confounding popularity bias in model training and adjusts the recommendation score with desired popularity bias via causal intervention. The major consideration is to decouple the user-item matching with item popularity: $f_{\theta}(u, i, m_i^t) = f_{\theta}(u, i) \cdot g(m_i^t)$, where m_i^t is the item i 's present popularity calculated by a statistical method. We can draw an analogy that the popularity bias in PD corresponds to the data distribution bias in our scenario. Thus, in our experiment, we replace m_i^t with our distribution vector and use a $MLP(\cdot)$ to calculate the refining score.

- **DecRS** [38] explicitly models the causal relations of user representations during training, and leverages backdoor adjustment to eliminate the impact of the confounder which causes the bias amplification. In our experiment, we replace the group-level user representation in original model with a category-level representation of the candidate list.

(4) From the view of re-ranking:

- **DLCM** [1] employs a RNN to sequentially encode the top results using their feature vectors and learn a local context model to re-rank the top results. We apply DLCM to re-rank the ranking list of base models.

- **PRM** [29] employs a Transformer structure to model the global relationships between any pair of items in the whole list. We apply *PRM-BASE* to re-rank base model's results.

4.3.2 Results. The overall results are shown in Table 3, where we can make the following observations:

- In general, Ada-Ranker performs best on three datasets and with two base models (with an exception for GAUC metric in Taobao dataset), which demonstrates the effectiveness of our method.
- DNS hurts the performances of base models, which indicates that our assumption – DNS can encode and generalize to various types of data distribution – does not hold.

Table 4: Results of ablation Study.

Methods	ML10M		Taobao		Xbox	
	GAUC	NDCG	GAUC	NDCG	GAUC	NDCG
Base (GRU4Rec)	0.9187	0.7781	0.8734	0.7208	0.9388	0.8302
(1) avg	0.9238	0.7853	0.8792	0.7349	0.9423	0.8384
w/o FiLM	0.9261	0.7899	0.8861	0.7463	0.9444	0.8416
(2) add_bias	0.9273	0.7921	0.8876	0.7472	0.9437	0.8422
diff γ, β	0.9260	0.7889	0.8881	0.7481	0.9429	0.8392
w/o mem_net	0.9243	0.7863	0.8855	0.7459	0.9407	0.8357
w/o global_para	0.9209	0.7819	0.8879	0.7448	0.9437	0.8411
free_para	0.9269	0.7907	0.8851	0.7445	0.9442	0.8408
(3) add_bias (1 layer)	0.9257	0.7887	0.8843	0.7409	0.9438	0.8409
add_bias (2 layers)	0.9263	0.7894	0.8849	0.7415	0.9433	0.8396
#slots $L=5$	0.9260	0.7896	0.8881	0.7483	0.9444	0.8417
#slots $L=15$	0.9268	0.7909	0.8880	0.7480	0.9447	0.8423
Ada-Ranker	0.9279	0.7932	0.8888	0.7490	0.9451	0.8438

Table 5: Comparisons of different training strategies.

Methods	Training Strategies	GRU4Rec					
		ML10M		Taobao		Xbox	
		GAUC	NDCG	GAUC	NDCG	GAUC	NDCG
Base	Train Θ	0.9187	0.7781	0.8734	0.7208	0.9388	0.8302
	$\emptyset \Rightarrow \Theta + \Phi$	0.9279	0.7905	0.8805	0.7314	0.9434	0.8399
Ada-Ranker	$\Theta \Rightarrow \Theta + \Phi$	0.9288	0.7943	0.8906	0.7462	0.9488	0.8511
	$\Theta \Rightarrow \Phi$	0.9279	0.7932	0.8888	0.7490	0.9451	0.8438

- Methods based on GSF explicitly model distribution of candidates through concatenation/average/Transformer, and the results show that these ways can boost GRU4Rec and SASRec. As for GSF_{concat}, although it achieves surprising performances on Taobao dataset, it is unstable on other two datasets. GSF_{trm} almost outperforms GSF_{avg} on all datasets, and this owes to the high-quality context-aware representations calculated by the multi-head self-attention.
- PD and DecRS are causal recommendation models which revise the scores of candidates by considering data distribution bias. Both methods can improve the performance of GRU4Rec and SASRec. Note that PD performs best on Taobao dataset in terms of GAUC metric, but it is much worse than Ada-Ranker with NDCG on three datasets. To some extent, this result shows that PD is not robust to adapt model to a given list of candidates only by decoupling the user preference and data distribution bias.
- DLCM and PRM are two re-ranking methods, which decouple the process of training base model and the post-processing of the ranking lists. However, the decoupled design can't deeply fuse the base model with the candidate distribution information, so Ada-Ranker can outperform both DLCM and PRM.

4.4 Ablation Study

4.4.1 Effect of the Components. Key components in Ada-Ranker include (1) distribution extractor, (2) input modulation, and (3) parameter modulation. To verify each component's impact, we disable one component or replace it with some variant each time while keeping the other settings unchanged, then test how the performance will be affected. The result is reported in Table 4.

- For (1), we replace the neural process (Section 3.2) by averaging all candidates' embeddings (denoted as avg). The decline of performance indicates it is hard to fully capture the distribution information of the candidate list by simply averaging. Moreover, we provide visual analysis in Section 4.5.
- For (2), we try different variants: remove the input modulation (w/o FiLM); or replace FiLM with element-wise addition parameter (add_bias), which means that replacing Eq.(11) and Eq.(10) with $\tilde{\mathbf{q}}_t = \mathbf{q}_t + \tilde{\mathbf{b}}$ and $\tilde{\mathbf{b}} = f'_b(\mathbf{z})$; or let FiLM generate different γ, β variables for each item in the item sequence (diff γ, β), by taking item representation as additive input in Eq.(10). Removing FiLM block leads to a performance drop, which verifies the necessity of input modulation. diff γ, β does not performs well, which shows that learning different affine transformations on each item in the sequence is difficult. add_bias's performance on Xbox dataset is poor. In comparison, FiLM's performance is robust and consistently good.
- For (3), we consider these variants: removing parameter modulation (w/o mem_net); removing global parameters of predictive layer (w/o global_para), which means replacing Eq.(13) with $\tilde{\mathbf{W}}_k = \tilde{\mathbf{W}}_k, \tilde{\mathbf{b}}_k = \tilde{\mathbf{b}}_k$; generating adaptation parameters by Eq.(12) (free_para); adding distribution bias vector after a linear projection to the first layer of prediction layer (add_bias (1 layer)) or to both layers (add_bias (2 layers)); changing the number of slots in memory network (#slots $L = 5$, and #slots $L = 15$). No matter removing memory network or replacing it with other configurations, the performances drop much compared with Ada-Ranker.

4.4.2 Effect of Different Training Strategies. We compare different training strategies (see Section 3.5) on **global parameters** Θ and **adaptation parameters** Φ in Ada-Ranker. Table 5 shows that all methods can boost base models significantly. Specially, jointly training Θ and Φ from scratch performs worse than the other two methods in most cases, which shows that the two-stage training procedure is of great importance. Besides, fine-tuning both Θ and Φ outperforms fine-tuning Φ only (except a result in NDCG metric on Taobao dataset). This is in expectation. Θ has far more parameters than Φ . Making Θ fixed in the finetuning stage is a strict constraint, which limits the expressive ability to some extent. However, as discussion in Section 3.5, only fine-tuning Φ makes Ada-Ranker more flexible for real-world applications.

4.5 Qualitative Study

In this section, we explore whether our Ada-Ranker can distinguish different distributions of candidates. We choose 11 single categories (such as *Animation, Children*) and 2 mixed categories (including *Mystery&Thriller* and *Thriller&Crime*) of ML10M dataset. For each category, we generate 100 groups of candidate lists in which the items are all from the current category. For mixed category, we sample items from two categories with the proportion of 50% and 50%, respectively. Then, we apply the trained Ada-Ranker to infer on all candidate lists. Figure 4 shows t-SNE plots of the distribution vector \mathbf{z} in Eq. (9) and the coefficients α of memory network in Eq. (16). We compare two types of Ada-Ranker: generating \mathbf{z} by neural process (denoted as NP) and by averaging all candidates

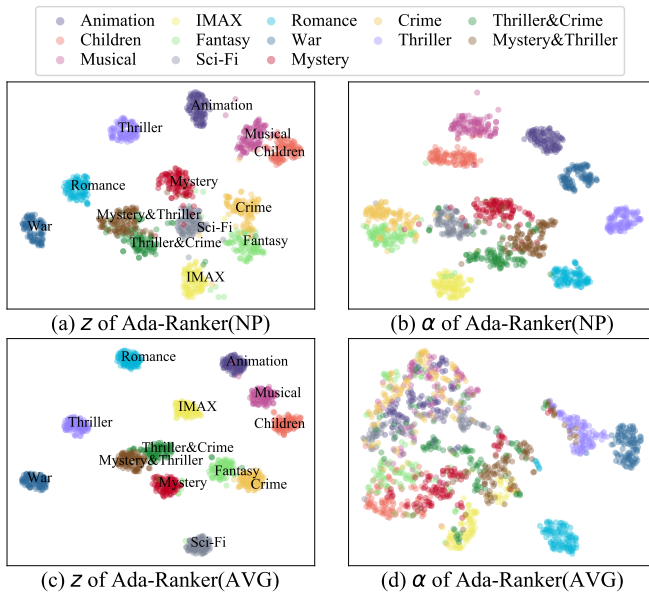


Figure 4: t-SNE plots for distribution representation z (a and c) and coefficients α of memory network (b and d), based on two different types of distribution extractor: NP and AVG. The dataset is ML10M.

embeddings (denoted as AVG). Each color represents a category. We have the following observations:

- In Figure 4(a) and Figure 4(c), distributions z generated by both NP or AVG can be clearly separated. However, we observe that z generated from NP has more flexibility. For example, in terms of mixed distribution like *Thriller&Crime*, the AVG method simply places $z\text{-Thriller\&Crime}$ in the middle of $z\text{-Thriller}$ and $z\text{-Crime}$, while the NP method can locate $z\text{-Thriller\&Crime}$ to a better place to avoid overlapping with irrelevant categories. Similar conclusions can be drawn for $z\text{-Mystery\&Thriller}$.
- For parameter modulation, the AVG method fails to distinguish α in memory network effectively (see Figure 4 (d)). It shows that our neural process encoder can model diverse distributions of different candidate lists well and improve the process of subsequent parameter modulation.

4.6 Adaptation to New Distributions

In Section 4.1.1 we introduce the distribution-mixer sampling as our default experiment setting, which ensures that both training set and test set contain diverse data distributions across different ranking requests. However, the prior distribution (i.e., Step 1 and Step 2 in the distribution-mixer sampling, which determine the distribution of item categories and item popularity) for both training set and test set are the same, which is a strong assumption and causes some gap between research and practice. In this section, we would like to answer this question: is Ada-Ranker generalized sufficiently to various scenarios, even when the test set’s prior distribution differs from the training set’s (which means that we are handling new emerging requests)?

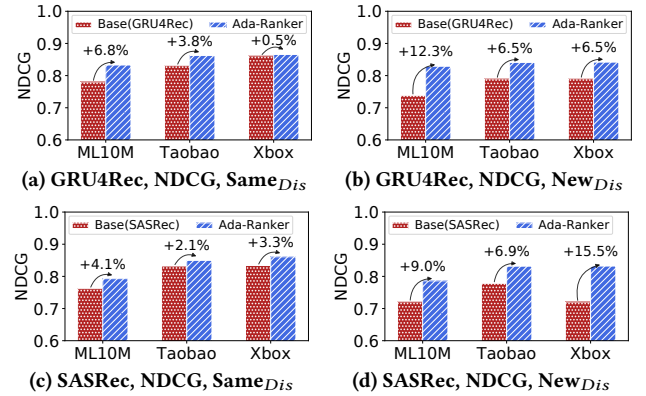


Figure 5: NDCG scores of Ada-Ranker with two base models {GRU4Rec, SASRec} on three datasets {ML10M, Taobao, Xbox} under different distribution settings. Same $_{Dis}$ means test set is under the same distribution as training set. New $_{Dis}$ means test set is generated from a new data distribution.

4.6.1 Settings. We revisit the classical recall-and-rank setting. The candidates that a ranker model need to handle are proposed by some recall models. However, for an industrial recommender system, there are several reasons to cause the inference distribution not identical to the training distribution, such as (1) the intermediate candidate set from recall models are not displayed to users, so we are not able to reconstruct the recall distribution from impression logs; (2) the recall models themselves are evolving from time to time. Therefore, we design another kind of setting, where the major difference lies in negative instances selection, with the same positive instances in Section 4.1.1.

We choose three widely-used recall methods (Popular, MF and Item2Item) to generate negative instances. Specifically, when constructing the training set, for each positive instance we generate 19 negative samples from the three recall models with a budget allocation of 20%, 50% and 30%, respectively. Then, we generate two different test sets in order to compare the performances of base models and Ada-Ranker under the same data distribution or a new data distribution: (1) Same $_{Dis}$: sampling negatives with the same distribution as the training set, i.e., 20%, 50% and 30% from the three recall methods. (2) New $_{Dis}$: sampling negatives with new distribution, i.e., sample negative items from Popular, MF, Item2Item with the proportion of 40%, 10% and 50%, respectively. The second setting increases the difficulty of inference which is closer to the actual scenario. For more details please refer to the Appendix A.1.2.

4.6.2 Results. From Figure 5 and Figure A1, we can observe that no matter the data distribution between training set and test set is the same or different, Ada-Ranker can improve the performance of two base models (GRU4Rec and SASRec) effectively and consistently, on three datasets and two metrics (results on GAUC metric is reported in Figure A1). Besides, when the distribution of test set is new, the performances of base models drop a lot compared with the setting under the same data distribution (e.g., GRU4Rec performs 0.78 on ML10M dataset over NDCG metric under Same $_{Dis}$, while it only performs 0.73 on the new test set). Such a performance decrease

is mainly caused by the inconsistent data distribution. Therefore, the relative improvement of the Ada-Ranker over the base model is more significant in the New_{Dis} setting than in the $Same_{Dis}$ setting. For example, in the Xbox dataset, with SASRec as the base model, Ada-Ranker’s NDCG gain in $Same_{Dis}$ is 3.3%, while in New_{Dis} , it is as high as 15.5%. This observation indicates that Ada-Ranker is capable of modeling the distribution of candidate samples and adjust parameters of the network accordingly, which alleviates the issue of inconsistent data distribution.

5 RELATED WORKS

Neural Recommender Systems. A web-scale recommender system usually comprises of recall modules and ranking modules to response users’ requests in (near) real-time [9, 26]. Recall modules propose a small set (such as thousands) of potentially relevant items as candidates with lightweight models. Ranking modules will score the given item candidates by adopting more advanced models to ensure the final recommendations are accurate.

One important recommendation scenario is the sequential recommendation, where DNNs have been successfully applied, such as GRU4Rec [15], NextItNet [45], and SASRec [18]. However, all these approaches follow the same traditional train-freeze-test paradigm, where a global optimal model is learned from the training dataset and then is frozen for testing. The testing environment in recommender systems is highly dynamic with diverse users’ requests, thus, it needs an approach that can perceive the current task’s peculiarity and conduct suitable parameter adaptation.

Domain Generalization. Domain generalization has attracted increasing interests in many application fields (e.g., NLP and CV). It aims at learning a model from several different but related domains, then generalize to unseen testing domain [37]. Related techniques can be roughly divided into four categories.

Multi-task Learning: Industrial recommender systems usually serve diverse application situations. For example, the *banner* and *Guess What You Like* channels on Taobao app homepage [33]. Multi-task learning (MTL) optimizes a model by jointly training several related tasks simultaneously. Typical models include parameter sharing of hidden layers [6], Multi-gate Mixture-of-Experts (MMoE)-based methods [28, 35], and Star Topology Adaptive Recommender (STAR) [33]. Note that all these MLT methods can only work when the tasks/domains are fixed and well-defined, they can hardly generalize to new domains or ad-hoc tasks with dynamic distributions.

Model Patch: Finetuning a pre-trained model with target task-specific data is the easiest way to transfer knowledge from source domains to a target domain. In this line of research, this work [17] proposes a parameter-efficient tuning paradigm called *model patch*, which adapts a large-scale pre-trained model to a specific downstream task by finetuning only a small number of new parameters. PeterRec [44] aims to adapt a user representation model for multiple downstream tasks. All these models require that the tasks/domains are fixed and known beforehand, they cannot directly apply to our scenario, but the idea motivates the design of our Ada-Ranker.

Meta-learning: Meta-learning is a learning-to-learn paradigm, aiming to learn experience from a distribution of related tasks so as to improve its future learning performance [16]. In particular,

MAML [12] based methods is frequently applied recently in cold-start recommendations for fast adaptation to new users [11, 23, 40, 43]. However, it requires a support set with supervised labels in the adaptation stage, which is not feasible in our setting, where all the items in candidate set C are unlabeled.

Feature Modulation: Another popular approach to adaptively altering a neural network’s behavior with dynamic conditions is FiLM [30], which carries out a feature-wise linear modulation on a neural network’s intermediate features conditioned on an arbitrary given query, so that the neural network computation over an input can be influenced by a dynamic condition. TaNP [27] introduces an improved version of FiLM, named Gating-FiLM, to enhance the ability of filtering the information which has the opposite effects during adaptation. AdaSpeech [7] adapts few adaptation data to personalized voices, which modulates hidden state vectors in network by categorizing the acoustic conditions in different levels and conditional layer normalization. In this paper, FiLM serves as one of the fundamental components in Ada-Ranker.

List-wise Ranking and Re-Ranking. List-wise learning-to-rank methods, which assign scores to items in a list by maximizing the utility of the entire list, are also related to our work. Most of them train models with list-wise loss functions, however, in the inference stage, they still use a point-wise scoring function (aka univariate scoring function) which is unaware of the context of candidate list [3–5]. GSFs [2] proposes multivariate scoring functions, in which the score of an item is determined jointly by multiple items in the candidate list. However, GSFs is hard to converge to a satisfying result because that by jointly modeling all items in the list, a large amount of new parameters are involved. Another line of related research is re-ranking [1, 29, 52], which refines a ranker’s suggested top- k items with another model by taking the top top- k items as ranking context. Typically, they are post-processing methods which are decoupled with ranker models. Since they take the whole candidate items as inputs, we can expect that they naturally have the ability to capture data distributions in candidate set. So we use them as baselines in Section 4.3.

6 CONCLUSION

In this paper, we propose a novel ranking paradigm, Ada-Ranker, which can perceive the data distribution of a specific candidate list and learn to adjust the ranker model before making predictions. In contrast to traditional rankers’ *train-freeze-test* paradigm, Ada-Ranker uses a new *train-adapt-test* paradigm. Ada-Ranker is lightweight, model-agnostic, and flexible for plug-and-play usage. We have conducted extensive experiments on three real-world datasets. Results show that Ada-Ranker can effectively enhance various types of base sequential recommender models, as well as outperform a comprehensive set of competitive baselines.

ACKNOWLEDGMENTS

This work was partially supported by National Natural Science Foundation of China under Grant No. 61872369, Beijing Natural Science Foundation under Grant No. 4222027, and Beijing Outstanding Young Scientist Program under Grant No. BJJWZYJH012019100020098. This work is also partially supported by Beijing Academy of Artificial Intelligence (BAAI). Xin Zhao is the corresponding author.

REFERENCES

- [1] Qingyao Ai, Keping Bi, Jiafeng Guo, and W. Bruce Croft. 2018. Learning a Deep Listwise Context Model for Ranking Refinement. In *SIGIR*. ACM, 135–144.
- [2] Qingyao Ai, Xuanhui Wang, Sebastian Bruch, Nadav Golbandi, Michael Bendersky, and Marc Najork. 2019. Learning Groupwise Multivariate Scoring Functions Using Deep Neural Networks. In *ICTIR*. ACM, 85–92.
- [3] Christopher JC Burges. 2010. From ranknet to lambdarank to lambdamart: An overview. *Learning* 11, 23-581 (2010), 81.
- [4] Christopher J. C. Burges, Robert Ragno, and Quoc Viet Le. 2006. Learning to Rank with Nonsmooth Cost Functions. In *NIPS*. MIT Press, 193–200.
- [5] 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.
- [6] Rich Caruana. 1998. Multitask Learning. In *Learning to Learn*. Springer, 95–133.
- [7] Mingjian Chen, Xu Tan, Bohan Li, Yanqing Liu, Tao Qin, Sheng Zhao, and Tie-Yan Liu. 2021. AdaSpeech: Adaptive Text to Speech for Custom Voice. In *ICLR*. OpenReview.net.
- [8] Kyunghyun Cho, Bart van Merriënboer, Çaglar Gülçehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *EMNLP*. ACL, 1724–1734.
- [9] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In *RecSys*. ACM, 191–198.
- [10] Zhenhua Dong, Hong Zhu, Pengxiang Cheng, Xinhua Feng, Guohao Cai, Xiquang He, Jun Xu, and Jirong Wen. 2020. Counterfactual learning for recommender system. In *RecSys*. ACM, 568–569.
- [11] Zhengxiao Du, Xiaowei Wang, Hongxia Yang, Jingren Zhou, and Jie Tang. 2019. Sequential Scenario-Specific Meta Learner for Online Recommendation. In *KDD*. ACM, 2895–2904.
- [12] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *ICML (Proceedings of Machine Learning Research, Vol. 70)*. PMLR, 1126–1135.
- [13] Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Jimenez Rezende, and S. M. Ali Eslami. 2018. Conditional Neural Processes. In *ICML (Proceedings of Machine Learning Research, Vol. 80)*. PMLR, 1690–1699.
- [14] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J. Rezende, S. M. Ali Eslami, and Yee Whye Teh. 2018. Neural Processes.
- [15] B. Hidasi, A. Karatzoglou, L. Baltrunas, and D. Tikk. 2016. Session-based Recommendations with Recurrent Neural Networks. In *ICLR 2016*.
- [16] Timothy M. Hospedales, Antreas Antoniou, Paul Micaelli, and Amos J. Storkey. 2020. Meta-Learning in Neural Networks: A Survey. *CoRR* abs/2004.05439 (2020).
- [17] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-Efficient Transfer Learning for NLP. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 2790–2799.
- [18] W.-C. Kang and J. J. McAuley. 2018. Self-Attentive Sequential Recommendation. In *ICDM 2018*. 197–206.
- [19] Diederik P. Kingma and Max Welling. 2014. Auto-Encoding Variational Bayes. In *ICLR*.
- [20] Yehuda Koren, Robert M. Bell, and Chris Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (2009), 30–37.
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *NIPS*. 1106–1114.
- [22] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems* 25 (2012), 1097–1105.
- [23] Hoyeop Lee, Jinbae Im, Seongwon Jang, Hynsounk Cho, and Sehee Chung. 2019. MeLU: Meta-Learned User Preference Estimator for Cold-Start Recommendation. In *KDD*. ACM, 1073–1082.
- [24] J. Li, P. Ren, Z. Chen, Z. Ren, T. Lian, and J. Ma. 2017. Neural Attentive Session-based Recommendation. In *CIKM 2017*. 1419–1428.
- [25] Jianxun Lian, Iyad Batal, Zheng Liu, Akshay Soni, Eun Yong Kang, Yajun Wang, and Xing Xie. 2021. Multi-Interest-Aware User Modeling for Large-Scale Sequential Recommendations. *arXiv preprint arXiv:2102.09211* (2021).
- [26] Jianxun Lian, Fuzheng Zhang, Xing Xie, and Guangzhong Sun. 2018. Towards Better Representation Learning for Personalized News Recommendation: a Multi-Channel Deep Fusion Approach. In *IJCAI*. ijcai.org, 3805–3811.
- [27] Xixun Lin, Jia Wu, Chuan Zhou, Shirui Pan, Yanan Cao, and Bin Wang. 2021. Task-adaptive Neural Process for User Cold-Start Recommendation. In *WWW*. ACM / IW3C2, 1306–1316.
- [28] Jiaqi Ma, Zhe Zhao, Xinyang Yi, Jilin Chen, Lichan Hong, and Ed H. Chi. 2018. Modeling Task Relationships in Multi-task Learning with Multi-gate Mixture-of-Experts. In *KDD*. ACM, 1930–1939.
- [29] Changhua Pei, Yi Zhang, Yongfeng Zhang, Fei Sun, Xiao Lin, Hanxiao Sun, Jian Wu, Peng Jiang, Junfeng Ge, Wenwu Ou, and Dan Pei. 2019. Personalized re-ranking for recommendation. In *RecSys*. ACM, 3–11.
- [30] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. 2018. FiLM: Visual Reasoning with a General Conditioning Layer. In *AAAI*. AAAI Press, 3942–3951.
- [31] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. 2019. Do ImageNet Classifiers Generalize to ImageNet?. In *ICML (Proceedings of Machine Learning Research, Vol. 97)*. PMLR, 5389–5400.
- [32] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy P. Lillicrap. 2016. Meta-Learning with Memory-Augmented Neural Networks. In *ICML (JMLR Workshop and Conference Proceedings, Vol. 48)*. JMLR.org, 1842–1850.
- [33] Xiang-Rong Sheng, Liqin Zhao, Guorui Zhou, Xinyao Ding, Binding Dai, Qiang Luo, Siran Yang, Jingshan Lv, Chi Zhang, Hongbo Deng, and Xiaoqiang Zhu. 2021. One Model to Serve All: Star Topology Adaptive Recommender for Multi-Domain CTR Prediction. In *CIKM*. ACM, 4104–4113.
- [34] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei A. Efros, and Moritz Hardt. 2020. Test-Time Training with Self-Supervision for Generalization under Distribution Shifts. In *ICML (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 9229–9248.
- [35] Hongyan Tang, Junning Liu, Ming Zhao, and Xudong Gong. 2020. Progressive Layered Extraction (PLE): A Novel Multi-Task Learning (MTL) Model for Personalized Recommendations. In *RecSys*. ACM, 269–278.
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *NIPS*. 5998–6008.
- [37] Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, and Tao Qin. 2021. Generalizing to Unseen Domains: A Survey on Domain Generalization. In *IJCAI*. ijcai.org, 4627–4635.
- [38] Wenjie Wang, Fuli Feng, Xiangnan He, Xiang Wang, and Tat-Seng Chua. 2021. Deconfounded Recommendation for Alleviating Bias Amplification. In *KDD*. ACM, 1717–1725.
- [39] Xuejian Wang, Lantao Yu, Kan Ren, Guanyu Tao, Weinan Zhang, Yong Yu, and Jun Wang. 2017. Dynamic Attention Deep Model for Article Recommendation by Learning Human Editors’ Demonstration. In *KDD*. ACM, 2051–2059.
- [40] Tianxin Wei, Ziwei Wu, Ruirui Li, Ziniu Hu, Fuli Feng, Xiangnan He, Yizhou Sun, and Wei Wang. 2020. Fast Adaptation for Cold-start Collaborative Filtering with Meta-learning. In *ICDM*. IEEE, 661–670.
- [41] Shu Wu, Yuyuan Tang, Yanqiao Zhu, Liang Wang, Xing Xie, and Tieniu Tan. 2019. Session-Based Recommendation with Graph Neural Networks. In *AAAI*. AAAI Press, 346–353.
- [42] Haochao Ying, Fuzhen Zhuang, Fuzheng Zhang, Yanchi Liu, Guandong Xu, Xing Xie, Hui Xiong, and Jian Wu. 2018. Sequential Recommender System based on Hierarchical Attention Networks. In *IJCAI*. ijcai.org, 3926–3932.
- [43] Runsheng Yu, Yu Gong, Xu He, Yu Zhu, Qingwen Liu, Wenwu Ou, and Bo An. 2021. Personalized Adaptive Meta Learning for Cold-start User Preference Prediction. In *AAAI*. AAAI Press, 10772–10780.
- [44] Fajie Yuan, Xiangnan He, Alexandros Karatzoglou, and Liguang Zhang. 2020. Parameter-Efficient Transfer from Sequential Behaviors for User Modeling and Recommendation. In *SIGIR*. ACM, 1469–1478.
- [45] Fajie Yuan, Alexandros Karatzoglou, Ioannis Arapakis, Joemon M. Jose, and Xiangnan He. 2019. A Simple Convolutional Generative Network for Next Item Recommendation. In *WSDM*. ACM, 582–590.
- [46] Weinan Zhang, Tianqi Chen, Jun Wang, and Yong Yu. 2013. Optimizing top-n collaborative filtering via dynamic negative item sampling. In *SIGIR*. ACM, 785–788.
- [47] Yang Zhang, Fuli Feng, Xiangnan He, Tianxin Wei, Chonggang Song, Guohui Ling, and Yongdong Zhang. 2021. Causal Intervention for Leveraging Popularity Bias in Recommendation. In *SIGIR*. ACM, 11–20.
- [48] Wayne Xin Zhao, Junhua Chen, Pengfei Wang, Qi Gu, and Ji-Rong Wen. 2020. Revisiting Alternative Experimental Settings for Evaluating Top-N Item Recommendation Algorithms. In *CIKM*. ACM, 2329–2332.
- [49] Wayne Xin Zhao, Shanlei Mu, Yupeng Hou, Zihan Lin, Yushuo Chen, Xingyu Pan, Kaiyuan Li, Yujie Lu, Hui Wang, Changxin Tian, Yingqian Min, Zhichao Feng, Xinyan Fan, Xu Chen, Pengfei Wang, Wendi Ji, Yaliang Li, Xiaoling Wang, and Ji-Rong Wen. 2021. RecBole: Towards a Unified, Comprehensive and Efficient Framework for Recommendation Algorithms. In *CIKM*. ACM, 4653–4664.
- [50] Guorui Zhou, Na Mou, Ying Fan, Qi Pi, Weijie Bian, Chang Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Deep Interest Evolution Network for Click-Through Rate Prediction. In *AAAI*. AAAI Press, 5941–5948.
- [51] Kun Zhou, Hui Wang, Wayne Xin Zhao, Yutao Zhu, Sirui Wang, Fuzheng Zhang, Zhongyuan Wang, and Ji-Rong Wen. 2020. S3-Rec: Self-Supervised Learning for Sequential Recommendation with Mutual Information Maximization. In *CIKM*. ACM, 1893–1902.
- [52] Tao Zhuang, Wenwu Ou, and Zhirong Wang. 2018. Globally Optimized Mutual Influence Aware Ranking in E-Commerce Search. In *IJCAI*. ijcai.org, 3725–3731.

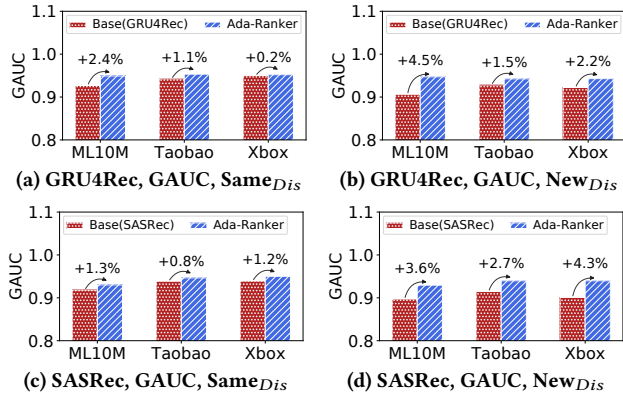


Figure A1: GAUC scores of Ada-Ranker with two base models {GRU4Rec, SASRec} on three datasets {ML10M, Taobao, Xbox} under different distribution settings. Same_{Dis} means test set is under the same distribution as training set. New_{Dis} means test set is generated from a new data distribution.

A APPENDIX

A.1 Settings of sampling negative instances from recall models

In section 4.6, we briefly introduce our method of constructing test sets with different distributions. In this section, we will introduce more details to help readers understand how we sample negative instances from recall models.

A.1.1 Three Recall Models. We choose Pop, MF, Item2Item as basic recall models to generate negative instances.

- **Pop.** We recall 1000 candidates from all items according to their popularity.
- **MF.** We first train a simple collaborative filtering model (e.g., MF) and obtain the user embedding table and the item embedding table. For each user, we will product its embedding with all items'

representations, and recall candidates ranked between 1000 and 2000.

- **Item2Item.** We take the historical behaviors of users as training data, and train item embeddings with the word2vec algorithm⁴. For each item, we will product its embedding with all items' representations, and recall candidates ranked between 500 and 1500.

A.1.2 Sampling According to Different Distributions. We assume the distribution of sampling is $\mathbf{d} = [d_1, d_2, d_3]$, where $\sum_{i=1}^3 d_i = 1.0$. Each element in \mathbf{d} represents the probability of sampling from the three recall models (Pop, MF and Item2Item, respectively).

For each positive instance (u, v) , we firstly use the function `numpy.random.multinomial` to obtain an instantiated number of items to be sampled by each recall model: $\mathbf{x} = [x_1, x_2, x_3]$, with $\sum_{i=1}^3 x_i = 19$. Then, we carry out the following three steps to generate the whole 19 negatives.

- Randomly sample x_1 items from the candidates in the range $[0, 1000]$ sorted by the item's popularity score.
- Randomly sample x_2 items from the candidates in range $[1000, 2000]$ sorted by the MF score (item i 's score is the dot product of embedding vectors of user u and item i).
- Randomly select one of the items that user has recently interacted with, say t . Randomly sample x_3 items from the candidates in the range $[500, 1500]$ sorted by the Item2Item score (item i 's score is the dot product of embedding vectors of item t and item i , and here the embedding vector is not from MF, but from a word2vec algorithm).

For training set, we set \mathbf{d} as $[0.2, 0.5, 0.3]$. For test set, we designed two settings: Same_{Dis} ($\mathbf{d} = [0.2, 0.5, 0.3]$) and New_{Dis} ($\mathbf{d} = [0.4, 0.1, 0.5]$). In section 4.6, we train the base models and Ada-Ranker on the training set, and infer on two different test sets.

A.2 Additional experimental results for Section 4.6

See Figure A1.

⁴<https://radimrehurek.com/gensim/models/word2vec.html>