

# Towards Sustainable Growth: A Multi-Value-Aware Retrieval Framework for E-Commerce Search

Yifan Wang\*  
Taobao & Tmall Group of Alibaba  
HangZhou, China  
wangyifan.ivan@alibaba-inc.com

Yixuan Wang\*  
Taobao & Tmall Group of Alibaba  
HangZhou, China  
suqin.wyx@alibaba-inc.com

Yidan Liang†  
Taobao & Tmall Group of Alibaba  
HangZhou, China  
liangyidan.lyd@alibaba-inc.com

Qiang Liu  
Taobao & Tmall Group of Alibaba  
HangZhou, China  
liuqiang.lq@alibaba-inc.com

Fei Xiao  
Taobao & Tmall Group of Alibaba  
HangZhou, China  
guren.xf@alibaba-inc.com

## Abstract

New item growth is critical for maintaining a healthy ecosystem in large-scale e-commerce platforms. However, existing systems tend to prioritize presenting users with already popular items—a phenomenon often referred to as the “Matthew effect”. In the context of search retrieval, current cold-start models suffer from the misalignment between training objectives and online business metrics, and they lack effective mechanisms to measure an item’s growth potential. In this paper, we propose a Multi-Value-Aware retrieval framework tailored for e-commerce search, designed to better align with the cascaded online values across different stages of the search system while balancing immediate conversion and long-term item growth. Our framework **GrowthGR** consists of two key components: an *Item Long-term Transaction Value Prediction (ItemLTV)* module and a *Multi-Value-Aware Generative Retrieval (MultiGR)* module. First, in the ItemLTV module, we employ counterfactual inference to quantify the long-term value increment attributable to a single user interaction. Second, in the MultiGR module, building upon a semantic-ID-based generative retrieval architecture, we leverage structured samples with the search cascade signals and adopt a *Multi-Value-Aware Policy Optimization (MoPO)* training paradigm to align with multi-stage online values, while explicitly balancing short-term transactional value and long-term growth potential estimated by ItemLTV. We successfully deployed GrowthGR on Taobao’s production platform, achieving a substantial 5.3% lift in new item GMV while delivering a non-trivial 0.3% gain in overall search GMV. Extensive online analysis and A/B testing demonstrate its positive impact on the overall ecosystem value.

## CCS Concepts

• **Information systems** → **Recommender systems**.

\*Both authors contributed equally to this paper.

†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](mailto:permissions@acm.org).

Conference acronym 'XX, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXXX.XXXXXXX>

## Keywords

E-Commerce Search, Item Cold-start, Growth Potential Prediction, Multi-Value-Aware Generative Retrieval

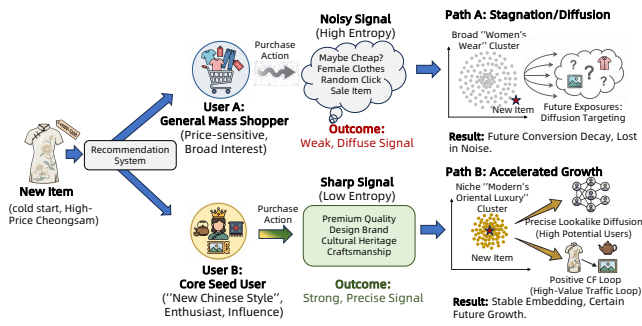
## ACM Reference Format:

Yifan Wang, Yixuan Wang, Yidan Liang, Qiang Liu, and Fei Xiao. 2018. Towards Sustainable Growth: A Multi-Value-Aware Retrieval Framework for E-Commerce Search. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

## 1 Introduction

In the era of rapidly evolving e-commerce, large-scale platforms like Taobao serve not only as transaction hubs but also as dynamic ecosystems where the continuous influx of new items is vital for long-term sustainability [28, 30, 36]. New items are the lifeblood of sustainable growth; they represent the latest consumer trends, foster brand innovation, and prevent the platform’s inventory from stagnating. However, effectively distributing these new items within the search scenario remains a non-trivial challenge. Most industrial search engines are inherently optimized for historical engagement metrics, such as Click-Through Rate (CTR) and Conversion Rate (CVR) [24]. This optimization logic tends to induce the “Matthew Effect” (or the “rich-get-richer” phenomenon), where the search system disproportionately favors head items with abundant historical data and high exposure [6]. Consequently, new items characterized by their massive scale and extreme data sparsity often struggle to break through this feedback loop [30]. The lack of prior interactions makes it difficult for traditional retrieval models to accurately estimate their potential value, resulting in a “cold-start” dilemma that undermines the sellers’ motivation and hinders the discovery experience for users.

Existing works have attempted to tackle the cold-start challenge from two primary directions. The first focuses on representation enhancement, leveraging multi-modal information, knowledge graphs, or Large Language Models (LLMs) to inject prior knowledge and mitigate popularity bias [5, 10, 31]. The second direction involves heuristic exposure mechanisms, such as dedicated traffic quotas or manual boosting interventions, to accelerate the growth of new items. While these methods provide temporary relief, they often suffer from significant efficiency trade-offs and lack a fundamental understanding of growth. We identify two critical limitations in these prior approaches:



**Figure 1: The Cold-start Dilemma: Immediate Conversion vs. Long-term Item Growth Value.**

- (1) **Lack of quantitative measurement for the long-term value of new items.** Most systems equate growth with simplistic metrics like current exposure share or immediate CTR/CVR, failing to capture the ripple effect an early-stage interaction has on an item’s future lifecycle.
- (2) **Limitations in balancing immediate conversion efficiency with long-term ecosystem health.** Existing strategies often focus too heavily on immediate gains, favoring high-conversion popular items to protect current revenue. This narrow focus sacrifices the growth of new items.

Inspired by these limitations, we posit that an effective item cold-start mechanism necessitates two fundamental capabilities:

- (1) **Quantifying the marginal growth value of individual interactions:** Rather than focusing solely on immediate conversion, the system should evaluate how a single interaction contributes to the long-term transaction of a new item. This enables fine-grained traffic allocation decisions.
- (2) **Precision-targeted initial distribution:** New items suffer from a lack of posterior feedback and insufficiently trained ID embeddings. Consequently, the initial trajectory is critical. Precise delivery to high-affinity users mitigates signal dilution and accelerates the acquisition of informative feedback.

To illustrate, as shown in Figure 1, consider a high-end designer dress in Taobao’s cold-start scenario: The system favors a niche enthusiast over a mass consumer, even if their predicted CVRs are identical. The interaction from an enthusiast provides a significantly sharper signal, allowing the system to decode the item’s precise market positioning early on. By establishing this high-fidelity initial seed, the system can leverage lookalike algorithms to accurately expand to similar high-potential users, carving out a sustainable growth path instead of being submerged in noisy, irrelevant traffic.

To address these challenges, we propose **GrowthGR**, a novel Multi-Value-Aware Generative Retrieval framework designed for sustainable new item growth. GrowthGR consists of two synergistic modules: the *Item Long-term Transaction Value (ItemLTV)* module and the *Multi-Value-Aware Generative Retrieval (MultiGR)* module. Recognizing that a single early-stage interaction for a new item can significantly shift its future trajectory—both in terms of system-level visibility and consumer perception—we formulate the long-term value estimation as a counterfactual causal inference problem. The ItemLTV module quantifies the uplift effect of an

individual interaction on future transaction increments, providing a rigorous metric for long-term potential. Building upon this, the MultiGR module introduces a generative modeling paradigm with dense semantic priors, leveraging the semantic IDs to transcend the limitations of traditional ID-based retrieval. To bridge the gap between multiple objectives, we employ preference alignment techniques *Multi-Value-Aware Policy Optimization (MoPO)* that integrate the item’s cascaded values within the search funnel (from candidates, impressions, and clicks to final conversions) with its predicted long-term potential. By doing so, GrowthGR aims to make the retrieval process not only aware of immediate gains but also strategically aligned with the platform’s sustainable development.

The main contributions of this paper are summarized as follows:

- **A Novel Perspective on Sustainable Growth:** We shift the focus from traditional short-term efficiency to the long-term ecological value of new items. We provide a quantitative measure the ItemLTV to estimate the uplift effect of individual interactions on an item’s future transaction trajectory.
- **A Multi-Value-Aware Generative Retrieval Framework:** We propose GrowthGR, the first generative retrieval framework tailored for new item growth.
- **Strategic Value Alignment:** We introduce a preference alignment mechanism MoPO that harmonizes multiple objectives within the search system. This allows the model to balance immediate conversion value with the predicted long-term potential.
- **Large-Scale Industrial Validation:** We have deployed GrowthGR on Taobao, one of the world’s largest e-commerce platforms. The results of large-scale online A/B tests conducted over two months indicate that our approach achieves a substantial 5.3% lift in new item GMV while delivering a non-trivial 0.3% gain in overall search GMV.

## 2 Related Works

### 2.1 Cold-Start Recommendation

To mitigate the data sparsity issue in cold-start scenarios, recent literature has shifted from heuristic approaches to deep learning paradigms that exploit auxiliary knowledge transfer, graph structures, and generative capabilities [2, 14, 21]. In the domain of meta-learning, FORM [29] innovates with an online regularized meta-leader to adapt to streaming user preferences, while PNMTA [25] and PAM [23] refine this by dynamically modulating network parameters based on user support sets and popularity biases, respectively. Concurrently, contrastive learning has emerged as a powerful tool to align disparate feature spaces; for instance, CCFCRec [39] bridges the gap between content and collaborative views via mutual information maximization, Aligning Distillation [12] employs ranking-based distillation to transfer knowledge from warm to cold items, and CMCLRec [32] extends these principles to cross-modal sequential recommendation. Addressing topological limitations, graph neural networks have been adapted for isolated cold nodes, with CC-GNN [33] introducing semantic phrase nodes to enhance e-commerce search connectivity, and CGRC [15] utilizing content-based graph reconstruction to predict missing links. Furthermore, cross-domain recommendation has seen breakthroughs like UEDMCF [22], which leverages optimal transport

to map user distributions across non-overlapping domains. Most recently, the field is embracing Large Language Models (LLMs), where ColdLLM [11] and LLMTreeRec [37] demonstrate the efficacy of zero-shot reasoning and tree-structured retrieval in overcoming extreme cold-start scenarios. Aliboot [28] first introduced the boosting principles and recommender techniques to ensure continuous growth of new items. However, these methods either only focus on immediate conversion or fail to accurately quantify the long-term value of new items.

## 2.2 Generative Recommendation

Over the past decade, industrial recommendation has relied on Deep Learning Recommendation Models (DLRM) focused on discriminative ranking. Recently, Generative AI and LLMs have catalyzed a shift toward Generative Recommendation [3, 7, 26, 38]. In terms of Generative Retrieval, TIGER [26] pioneered the use of RQ-VAE for semantic ID quantization, enabling autoregressive item prediction in cold-start scenarios, a direction advanced by OneRec [4] and TBGRecall [19], which integrate preference alignment and session-wise generation for industrial-scale, alongside GPT4Rec [18] which leverages search query generation. Regarding LLM-based Paradigms, P5 [7] established a unified text-to-text framework for multi-task recommendation via personalized prompts, inspiring subsequent works like VIP5 [8] for multimodal tasks, TALLRec [1] for efficient instruction tuning, and NoteLLM [35] and RecFormer [17] for content-centric representation generation. To support the computational demands of these generative models, Model Architectures are undergoing significant evolution; notably, Meta’s HSTU [34] introduced a hierarchical sequential transduction unit that replaces standard attention to validate scaling laws in trillion-parameter recommendation models, a high-efficiency direction echoed by KuaiFormer [20] and MTGR [9] in real-time industrial applications. Inspired by these works, we adopt a generative approach integrated with multi-value alignment.

## 3 Methodology

In this section, we present the technical details of GrowthGR, a comprehensive framework designed to foster sustainable new item growth through multi-value-aware generative retrieval, as illustrated in Figure 2. The methodology is structured into two primary components: 1) Item Long-term Transaction Value Prediction (ItemLTV): We describe our causal inference approach to quantify the uplift effect of new items. This module estimates the potential transaction increment triggered by user interactions. 2) Multi-Value-Aware Generative Retrieval (MultiGR): We detail the generative architecture and our proposed training strategy to balance immediate conversion value with the predicted long-term potential.

### 3.1 Item Long-term Transaction Value Prediction

The ItemLTV module identifies the growth potential of new items by modeling the causal impact of user interactions. In this context, we posit that user interactions facilitate item conversion from two perspectives. First, at the system level, these interactions enhance the item’s distribution efficiency in a personalized manner. For instance, if a new dress is clicked by a college student, the system

leverages this interaction to distribute the item to more students with similar user profiles. Second, at the user level, interactions such as conversions directly influence the item’s front-end display, which in turn affects the purchasing decisions of users in future exposures. Thus, we evaluate the impact of these two effects by employing a Counterfactual Causal Inference framework to estimate the long-term transaction increment triggered by specific user engagements.

**3.1.1 Problem Definition.** In our framework, we define the Treatment Point as the occurrence of a user click ( $W_i = 1$ ). Let  $X_i = \{x_i^I, x_i^C\}$  denote the covariates, where  $x_i^I$  represents item-specific attributes and  $x_i^C$  captures the user/query context. We define the potential outcomes for an item  $i$  as:

- $Y_i(1)$ : The average daily order count during the subsequent 7-day window following the initial 30-day "New Item Period" if a click occurs ( $W_i = 1$ ).
- $Y_i(0)$ : The base order count over the same 7-day period without the specific click ( $W_i = 0$ ).

Based on them, we further define the Uplift Growth as  $Y_i(1) - Y_i(0)$  and the Base Growth as  $Y_i(0)$ . We specifically choose user click over system exposure as the treatment variable for three critical reasons: 1) **Enhancing Signal Reliability**: System exposure alone has a limited effect on an item’s inherent growth trajectory and introduces high predictive variance in large-scale environments. 2) **Mitigating Negative Feedback**: In industrial search engines, exposure without a subsequent click is interpreted as negative feedback, leading the system to suppress the item in future pre-rank and rank stages. 3) **Reinforcing Success Cycles**: User clicks followed by downstream actions, such as favoriting or purchasing, provide high-fidelity labels. This feedback is returned to the model architecture, enhancing future efficiency estimations and naturally increasing the item’s future exposure opportunities.

To account for the heavy-tailed distribution of transaction values, we perform estimation in the logarithmic space. Our objective is to estimate the Conditional Average Treatment Effect (CATE), denoted as  $\tau(X_i)$ , representing the uplift growth attributable to the click:

$$\tau(X_i) = E[\log(Y_i(1) + 1) - \log(Y_i(0) + 1) | X_i]. \quad (1)$$

**3.1.2 Two-Tower Model Architecture.** We employ a two-tower structure to estimate these components:

- **Item Tower (Base Growth)**: This branch processes Item Features (e.g., semantic attributes) through an Embedding Layer to generate an Item Embedding. This embedding is passed through a dedicated MLP to output the predicted score for the item’s base growth  $G_{base}$ :

$$G_{base}(X_i) = f_1(x_i^I). \quad (2)$$

- **Uplift Prediction Tower (Uplift Growth)**: This tower integrates User History, User/Query Features, and Item Embedding. It utilizes an Attention mechanism to capture fine-grained interactions between user intent and item attributes, followed by a concatenation operation and an MLP to estimate the incremental value  $\tau(x_i)$ :

$$G_{uplift}(X_i) = f_2([g(x_i^C), x_i^I]). \quad (3)$$

During the training phase, the model is optimized based on the observed interaction. The predicted log-outcome  $\hat{y}_i$  is formulated

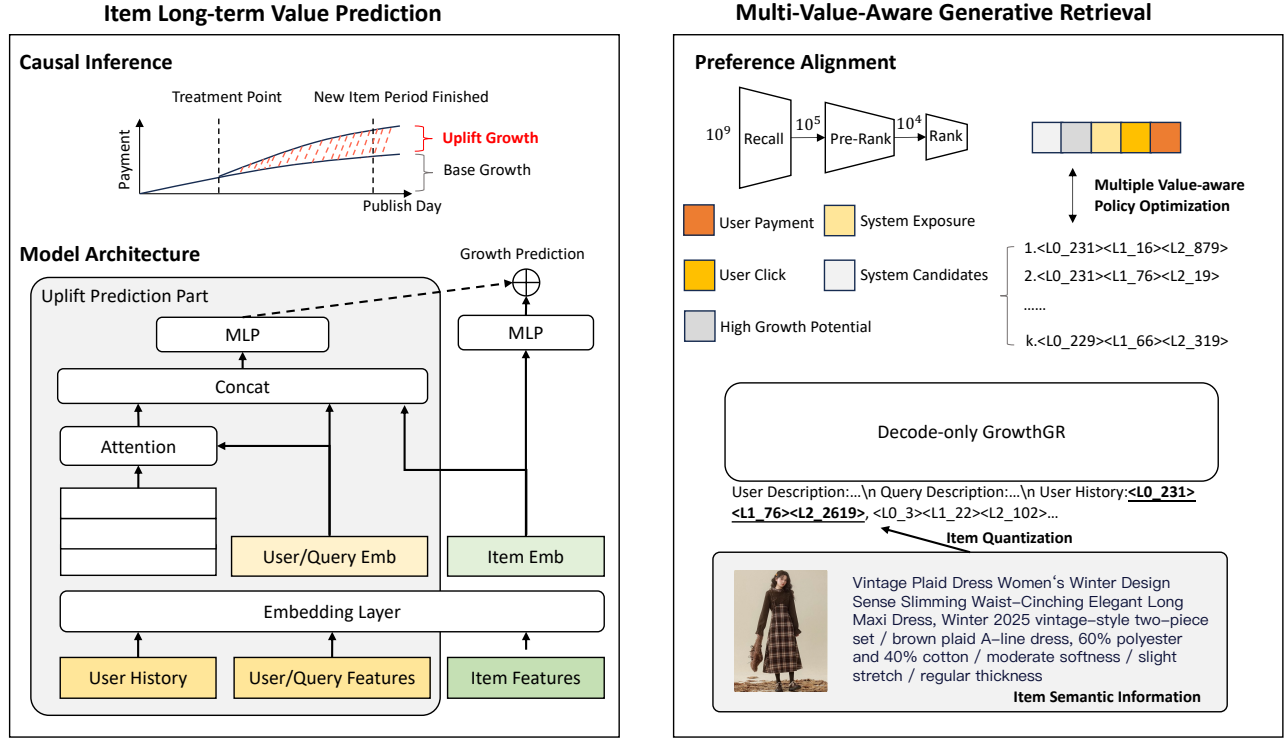


Figure 2: The overview of the proposed GrowthGR framework.

based on the treatment assignment  $W_i$ :

$$\hat{y}_i = G_{base}(X_i) + W_i \cdot G_{uplift}(X_i). \quad (4)$$

Then, we employ Mean Squared Error (MSE) as the loss function for model optimization:

$$L_{ItemLTV} = \sum_i \|\hat{y}_i - \log(Y_i + 1)\|_2^2, \quad (5)$$

where  $Y_i$  is the ground truth order count. By supervising the combined output against actual purchase data, the model learns to decouple the item's base growth value from the value added by user interaction.

### 3.2 Multi-Value-Aware Generative Retrieval

The MultiGR module is designed to identify high-potential new items from a billion-scale candidate pool by shifting the retrieval paradigm from traditional ID-matching to semantic ID generation, with an emphasis on balancing immediate conversion value and long-term value, to overcome data sparsity and ensure sustainable item discovery.

**3.2.1 Item Quantization.** To represent items in a discrete semantic space, we employ a hierarchical tokenization strategy inspired by the TIGER framework. We first utilize a pre-trained e-commerce foundation model to generate unified multimodal representations of items, capturing fine-grained item semantic information such as titles, properties, and images. We apply a Residual Quantized Variational AutoEncoder (RQ-VAE) [16] to these representations to

derive a three-layer semantic ID. Each item is assigned a hierarchical path (e.g.,  $\langle L0_{231} \rangle \langle L1_{16} \rangle \langle L2_{879} \rangle$ ), which preserves structural relationships between items via shared hierarchical prefixes and allows the model to generalize across semantically similar items even without historical interaction data.

**3.2.2 Generative Architecture.** The core of MultiGR is a Decoder-only Transformer-based model that formulates retrieval as a sequence-to-sequence task. The details are:

- **Input Composition:** The model takes a concatenated natural language string as input, including User/Query Features, description queries, and the user's historical interaction sequence represented as semantic IDs.
- **Autoregressive Generation:** The model learns to generate the semantic ID of the next item, effectively modeling the underlying user interests and transition patterns.
- **Constrained Decoding:** During inference, we implement a trie-based constrained decoding to ensure that the generated tokens follow a valid path within the pre-defined SID hierarchy, effectively filtering candidate items within the generation process.

**3.2.3 Multi-Objective Training Strategy.** To balance immediate conversion efficiency with long-term ecological value, we propose a two-stage training strategy:

**Supervised Pre-training.** The model is first trained using the Next Token Prediction (NTP) task on historical transaction logs.

The training loss function is as follows:

$$\mathcal{L}_{NTP} = -\frac{1}{\sum_{k=1}^N |o_k|} \sum_{i=1}^N \sum_{t=1}^{|o_i|} \log \mathcal{P}(o_{i,t} | o_{i,<t}, x), \quad (6)$$

where  $N$  is the number of samples,  $x$  denotes the input context (e.g., behavioral sequences), and  $\mathcal{P}$  denotes the probability of decoding the semantic ID  $o_i = (o_{i,1}, \dots, o_{i,|o_i|})$  corresponding to the purchased item. This stage provides a foundational capability for well-formed generation and aims to help the model learn the collaborative filtering signals of the platform while aligning latent representations of semantic IDs with natural language descriptions.

**Preference Alignment.** In the prior training stage, reliance on random negative sampling and the Maximum Likelihood Estimation (MLE) objective makes it difficult for the model to distinguish fine-grained relative rankings across outcomes in the causal funnel (e.g., exposure  $\rightarrow$  click  $\rightarrow$  purchase). Furthermore, since the inference paradigm of generative retrieval is typically beam search, it is crucial to focus not only on the scoring capability for top candidates but also on the collective value of the retrieved set. We propose MoPO, which inherits the efficient group-relative optimization framework of GRPO [27] while extending it with a Multi-value Reward Engine. Unlike vanilla GRPO that treats rewards as independent feedback, MoPO calibrates the reward signal using a clipped importance weighting (CIW) mechanism to counteract popularity bias and models the cascaded search funnel explicitly.

We consider the following objectives:

- **Cascaded Value Labels:** We utilize a hierarchy of labels including user purchase, user click, system exposure, and system rank candidates to model the search funnel.
- **Long-term Labels:** Since the ItemLTV module estimates the uplift triggered by user interactions, it is essential to ensure that the items exhibit a high predicted click-through rate. Thus, we leverage the ItemLTV module to identify high growth potential items by selecting those clicked items whose predicted uplift exceeds the global mean.

Different scores are assigned to different interaction labels, and the total reward is a calibrated weighted sum of these scores:

$$r_i = \text{Clip}(-\log \pi_{\theta_{old}}(o_i|x), 1, M) \sum_k w_k s_k, \quad (7)$$

where  $s_k \in \{0, 1\}$  is the indicator for the  $k$ -th objective (e.g., purchase),  $w_k$  is its corresponding weight, and  $\pi_{\theta_{old}}(o_i|x)$  is the probability of generating the semantic ID  $o_i$  under the behavior policy. The hyperparameter  $M > 1$  serves as an upper bound. To align the generated probability distribution with the ground-truth value distribution, we utilize the label distribution of the samples to determine the weights  $\{w_k\}$ . Furthermore, we incorporate a clipped importance weight, defined as  $\text{Clip}(-\log \pi_{\theta_{old}}(o_i|x), 1, M)$ , to effectively re-weight the reward signals. By employing the negative log-likelihood, we naturally assign higher importance to rare, long-tail semantic IDs that are assigned low probability by the behavior policy under the current context, while compressing the dynamic range of weights to ensure training stability. The clipping mechanism further ensures that head IDs retain a baseline reward to prevent vanishing gradients, while capping the influence of extreme outliers. This strategy effectively prevents the model from

being dominated by head samples, encouraging it to capture rare but high-value patterns and ultimately leading to a more diverse and optimal retrieval set. Let  $\rho_{i,t}(\theta) = \frac{\pi_{\theta}(o_{i,t}|x, o_{i,<t})}{\pi_{\theta_{old}}(o_{i,t}|x, o_{i,<t})}$  denote the importance sampling ratio. The final loss function is as follows:

$$\mathcal{L}_{MoPO} = -\frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \sum_{t=1}^{|o_i|} [\mathcal{J}_{i,t}^{CLIP}(\theta) - \beta D_{KL}(\pi_{\theta} || \pi_{ref})], \quad (8)$$

where we define the clipped surrogate objective as:

$$\mathcal{J}_{i,t}^{CLIP}(\theta) = \min[\rho_{i,t}(\theta) \hat{A}_{i,t}, \text{clip}(\rho_{i,t}(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_{i,t}], \quad (9)$$

and  $\hat{A}_{i,t} = \frac{r_i - \text{mean}(r)}{\text{std}(r)}$  is the normalized advantage.

**3.2.4 Inference.** At inference time, MultiGR leverages beam search to retrieve the top- $k$  sequences with the highest likelihood. Furthermore, a constrained decoding strategy is applied to guarantee that the generated semantic ID paths strictly conform to the pre-defined valid candidates. To resolve semantic ID collisions where multiple items share the same identifier, we utilize a dense retrieval model for internal re-ranking. Please refer to Appendix B for more details.

## 4 Experiments

In this section, we conduct a series of experiments on Taobao's real-world e-commerce data to comprehensively validate the effectiveness, the component contributions, the scaling properties, and the industrial applicability of GrowthGR. We aim to address the following research questions:

- **RQ1:** How does GrowthGR perform compared to competitive retrieval baselines?
- **RQ2:** How do the individual components contribute to capturing long-term value and improving retrieval accuracy?
- **RQ3:** Does the proposed generative retrieval framework adhere to scaling laws?
- **RQ4:** How does GrowthGR balance immediate conversion efficiency with the long-term growth of new items on Taobao?

### 4.1 Experimental Setup

**4.1.1 Datasets.** We conduct extensive experiments on Taobao's real-world e-commerce search dataset, which comprises billions of user-item interactions. Taobao is a leading global platform where intelligent retrieval solutions must handle massive-scale data to serve millions of users and items daily.

We utilize two distinct datasets tailored for our Multi-Value-Aware framework: 1) Uplift Prediction Dataset: Given that our ItemLTV module focuses on long-term transaction value growth triggered by specific user engagements, we utilize a specialized click-oriented dataset. It contains 2.4 billion interactions collected from historical search logs, covering 0.1 billion users and 0.16 billion items including 3 million new items. It is primarily used to train the uplift prediction model. 2) Full-chain Conversion Dataset: This dataset is designed to train the foundational retrieval and collaborative filtering capabilities within MultiGR. This set contains 5.0 billion multi-behavior interactions spanning 0.17 billion users and 0.11 billion items including 2 million new items. This large-scale data provides a solid foundation for the model to learn the basic semantic and behavioral alignments between users and items.

**Table 1: Statistics of the Datasets used in GrowthGR.**

Dataset	Interactions	Users	Items	New Items
Uplift	2.4B	0.1B	0.16B	3M
Full-chain	5.0B	0.17B	0.11B	2M
Online Daily	1.6B	0.15B	0.3B	3M

All online experiments are conducted on Taobao’s primary online search system. This environment supports large-scale validation of sustainable growth metrics. The daily traffic includes approximately 1.6 billion interactions involving 0.15 billion active users and 0.3 billion unique items, of which 3 million are new items.

**4.1.2 Evaluation Metrics.** To verify the effectiveness of GrowthGR in balancing immediate conversion efficiency with sustainable new item growth, we utilize a comprehensive set of offline and online metrics.

**Offline Metrics.** Our evaluation considers multiple ground-truth labels, including In-scenario Conversions marked as *search* (to assess the retrieval of head, mid, and tail items within the search context), Platform-wide Conversions marked as *all-net* (to measure the model’s ranking capability beyond the specific search scene, capturing incremental cross-scenario relevance) and Long-term Value marked as *long-term* (to assess the fidelity of the retrieval results to the ItemLTV module’s uplift predictions). Specifically, all-net labels are derived by linking a user’s site-wide transactions based on query relevance. Empirically, the performance on search labels reflects the model’s ability to fit the distribution of the current system, serving as an indicator of its alignment with the existing production pipeline. Recall performance on all-net labels exhibits higher consistency with final online immediate outcomes, thereby providing a more unbiased reflection of the model’s search effectiveness, while long-term labels follow the same definition as in Preference Alignment where they are empirically shown to be positively correlated with the future conversion of new items online. We employ two standard ranking metrics to evaluate the retrieval performance on historical logs:

- **Recall@k:** This metric measures the model’s ability to recall relevant items within the top  $k$  candidates. The formula is defined as:

$$\text{Recall@}k = \frac{1}{N} \sum_{i=1}^N \mathbb{I}(\text{rank}_i \leq k), \quad (10)$$

where  $N$  is the total number of relevant items,  $\text{rank}_i$  is the predicted rank and  $\mathbb{I}(\cdot)$  is the indicator function.

- **Normalized Discounted Cumulative Gain (NDCG) [13]:** To evaluate the quality of the ranked list and the model’s alignment with user preferences, we use NDCG.

**Online Metrics.** The online performance is measured via online A/B testing on Taobao, focusing on transaction impact, traffic structure, and long-term growth:

- **Gross Merchandise Volume (GMV):** The total sales revenue generated within the search scenario. This is the core metric for measuring immediate conversion efficiency and the platform’s overall economic health.

**Table 2: Performance Comparison of ItemLTV.**

Metric	MSE (↓)	NDCG (↑)
base	1.348	0.842
+uplift part	<b>1.329</b>	<b>0.853</b>

- **Page View Ratio (PVR):** Defined as the proportion of total exposures allocated to new items. This metric reflects the traffic structure and the system’s success in distributing traffic to new items to mitigate the “Matthew effect”.
- **Transaction Impact at T+30 (TI@30):** We adopt daily GMV during the subsequent 7-day window following the initial 30-day period as a proxy for the successful transition of items from the exploration phase to the exploitation phase. A higher TI@30 indicates that our model can accurately identify potential items during their early exposure window and provide them with an optimized growth trajectory. This aims to ensure that the platform does not merely provide temporary traffic boosts but fosters genuine market fit and sustained transaction volume for new items.

**4.1.3 Implementation Settings.** The Dense Retrieval (DR) model employed in our experiments consists of 40B sparse embedding parameters and 10M dense network parameters, which also serves as the current state-of-the-art production model. All generative models are configured with a total size of 0.5B, including 0.15B embedding parameters.

## 4.2 Main Results (RQ1)

**4.2.1 ItemLTV Effectiveness.** We first conduct an offline evaluation of the ItemLTV module independently. Since the uplift effect of user interactions on items is inherently counterfactual and lacks direct ground truth for evaluation, we assess the module’s performance by its ability to fit the aforementioned average daily order count of new items using MSE for regression accuracy and NDCG for ranking performance. Specifically, the “Base” model relies solely on the item-side tower for label prediction, while the “+uplift part” configuration incorporates the uplift prediction component for joint estimation. As shown in Table 2, the results demonstrate that the overall prediction reaches a high performance level. Notably, the “+uplift” variant exhibits superior fitting and ranking capabilities, confirming that the uplift module effectively captures the impact of user interactions.

Further analysis demonstrates that new items are significantly more sensitive to user interactions in their early stages (see Appendix C.1 for a comprehensive discussion).

**4.2.2 MultiGR Effectiveness.** Table 3 summarizes the offline performance of our proposed models compared to state-of-the-art baselines across three distinct labels: search, all-net, and long-term. Specifically, DR is the state-of-the-art dense retrieval model currently deployed in Taobao Search’s production system. TIGER [26] serves as a generative retrieval baseline, implemented using semantic IDs and Supervised Fine-Tuning. GrowthGR-twoStage denotes a two-stage retrieval pipeline where GrowthGR serves as the initial retriever to produce top-2,000 candidates, followed by a re-ranking

**Table 3: Experimental performance comparison on the Taobao real-world dataset. Bold indicates the best results while underline denotes the second-best results.**

Method	search				all-net				long-term			
	Recall@10	Recall@100	Recall@1000	NDCG	Recall@10	Recall@100	Recall@1000	NDCG	Recall@10	Recall@100	Recall@1000	NDCG
DR	0.4115	0.6775	0.8820	0.3628	0.2463	0.4260	0.6194	0.2422	0.3990	0.6704	0.8772	0.4119
TIGER	0.2950	0.5909	0.8215	0.2842	0.2159	0.4430	0.6969	0.2307	0.3810	0.7209	0.8883	0.3375
GrowthGR	0.3419	0.6306	0.8438	0.3160	0.2568	0.4981	0.7147	0.2555	0.4440	0.7523	0.8991	0.3820
GrowthGR-twoStage	0.513	0.8062	0.8820	0.4102	0.3187	0.5702	0.7578	0.2970	0.6221	0.8558	0.9261	0.4801

**Table 4: Ablation Study of GrowthGR Components.**

Model	all-net		long-term	
	Recall@10	Recall@1000	Recall@10	Recall@1000
<b>GrowthGR (Ours)</b>				
w/o ItemLTV	-0.0006	0.0004	-0.0532	-0.0165
w/o CIW	-0.0010	-0.0038	-0.0004	-0.0060
w/o MoPO	-0.0220	-0.0368	-0.0539	-0.0325

stage using the DR model to generate the final recommendation list. We observe several key findings as follows.

GrowthGR demonstrates a significant competitive advantage over traditional and generative baselines. GrowthGR consistently outperforms TIGER across all label types, particularly in all-net and long-term metrics. For instance, in the all-net category, GrowthGR achieves a Recall@1000 of 0.7147 compared to TIGER’s 0.6969. This highlights that our multi-value-aware approach and preference alignment effectively address the limitations of standard supervised fine-tuning in capturing cross-scenario relevance and growth potential. While DR shows strong performance on search labels reflecting its heavy optimization for the existing system’s distribution, GrowthGR achieves the second-best results in all-net and long-term recall. Notably, on long-term Recall@1000, GrowthGR (0.8991) outperforms DR (0.8772). This underscores GrowthGR’s superior ability to identify items with high growth potential that the current production system might overlook.

The GrowthGR-twoStage variant, which utilizes GrowthGR as the primary retriever followed by DR for re-ranking, achieves the best overall performance across every single metric. These results indicate that GrowthGR provides a high-quality candidate set rich in both immediate relevance and long-term potential. By serving as a robust initial retriever, it allows the ranking stage to operate on a pool of candidates that is fundamentally better aligned with both user intent and platform sustainability.

The performance on long-term labels is particularly telling for the MultiGR module’s effectiveness. GrowthGR-twoStage achieves an NDCG of 0.4801 and a Recall@1000 of 0.9261. This exceptionally high fidelity to uplift-based labels demonstrates that our framework successfully translates Preference Alignment into actual retrieval behavior, ensuring that items predicted to have the highest marginal return on exposure are prioritized. Empirically, the strong performance on these labels is positively correlated with the future conversion of new items online, validating our strategy for sustainable ecosystem growth.

### 4.3 Ablation Study (RQ2)

To further investigate the contribution of each component in GrowthGR, we conduct an ablation study. As shown in Table 4, removing either the ItemLTV module, the CIW or the MoPO strategy leads to a noticeable performance decline across all metrics.

**4.3.1 w/o ItemLTV.** We conducted an ablation study by removing the scores generated by the ItemLTV module from the MoPO framework. This allows us to investigate its specific impact on the recall performance for both immediate conversion and new items with high long-term value uplift. Specifically, the absence of ItemLTV results in a drop of 1.65 pt in Recall@1000 and 5.32 pt in Recall@10 for long-term labels, validating its critical role in capturing long-term value. Meanwhile, we observed that the impact on the Recall@1000 for immediate conversion was negligible, with a marginal decrease of only 0.04 pt.

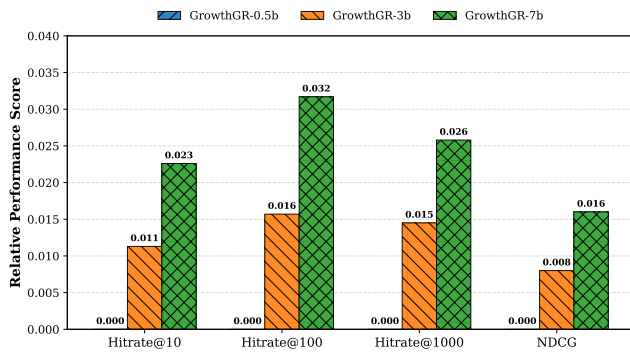
**4.3.2 w/o CIW.** We further investigated the impact of the clipped importance weighting mechanism by excluding the weighting term. The absence of CIW leads to a consistent performance degradation across all metrics. Specifically, we observed a decrease of 0.38 pt in Recall@1000 for all-net labels and a more pronounced drop of 0.60 pt in Recall@1000 for long-term labels. These results confirm that CIW is instrumental in mitigating popularity bias. By calibrating the rewards of frequently sampled head items, CIW encourages the model to explore and prioritize high-value yet sparse patterns in the long tail.

**4.3.3 w/o MoPO.** Similarly, we replaced the MoPO module with vanilla GRPO targeting immediate conversion as a baseline to evaluate the superiority of our proposed training paradigm. The results demonstrate that MoPO significantly boosts Recall@10 by 2.2 pt and Recall@1000 by 3.68 pt for all-net labels. Meanwhile, an increase of 5.39 pt in Recall@10 and 3.25 pt in Recall@1000 is observed for long-term labels. The performance degradation observed in the variant without MoPO confirms that MoPO is essential for achieving optimal retrieval accuracy.

To examine the impact of aligning training-time sampling with inference-stage decoding, we further explore different rollout strategies in Appendix C.2.

### 4.4 Scaling Analysis (RQ3)

To evaluate the scalability of the proposed GrowthGR framework, we conduct a systematic investigation by varying the model capacity across three parameter scales: 0.5B, 3B, and 7B. This analysis aims to determine whether our architectural design can effectively leverage increased computational power and parameters.



**Figure 3: Performance improvement on all-net labels across model scales relative to the 0.5B baseline.**

As illustrated in Figure 3, GrowthGR exhibits a clear and consistent scaling trend. We observe a monotonic improvement across all evaluation metrics, including  $\text{Recall}@k$  and NDCG as the parameter count increases. Notably, shifting from 0.5B to 7B yields substantial relative gains (e.g., over +2.0 pt in  $\text{Recall}@1000$ ), suggesting that the model has not yet reached a performance plateau. These empirical results underscore that GrowthGR adheres to the scaling law, demonstrating its potential for further performance breakthroughs as model capacity continues to expand in large-scale industrial applications.

#### 4.5 Online A/B Testing (RQ4)

We deployed GrowthGR on Taobao, a leading e-commerce platform, with billion-scale daily traffic for an observation period exceeding two months. In the Taobao search system, new items retrieved by the GrowthGR framework undergo a lightweight re-ranking process and are then output as a dedicated new item channel. This stream is subsequently blended with other streams (e.g., advertisement, mainstream) before being presented to users. Specific implementation details of the online deployment are available in Appendix A.

Our experimental design considers both immediate conversion efficiency and the long-term ecological impact on new items through two distinct A/B testing methodologies:

- **User-side A/B Testing:** We employ a standard user-level partitioning approach where incoming traffic is hashed by user IDs into orthogonal buckets. This setup allows us to compare the immediate conversion efficiency of different retrieval strategies by observing real-time user-item interactions.
- **Item-side A/B Testing:** Recognizing that the growth of new items is a long-term process, we implement item-level partitioning based on a uniform hash of item IDs. Different strategies are applied to specific item buckets to observe the differential growth trajectories over an extended period. To ensure that a specific group of new items is influenced by only one strategy and to prevent cross-contamination between policies, we utilize a coupled bucket mechanism.

Each item bucket is exclusively mapped to a corresponding user bucket.

**4.5.1 User-side A/B Testing.** To evaluate the real-world impact of GrowthGR, we conducted a large-scale online A/B experiment. The results demonstrate that GrowthGR significantly enhances both the exposure and conversion efficiency of new items, yielding a 5.39% increase in new item GMV and a 1.54% lift in PVR. From an ecosystem perspective, GrowthGR effectively orchestrates the dynamics of our two-sided marketplace: it satisfies the evolving novelty-seeking behavior of users while simultaneously providing tangible incentives for merchants to launch innovative items. Notably, the overall Search GMV grew by 0.31%, a statistically significant result. This confirms that our strategy does not lead to traffic cannibalization or a zero-sum redistribution. Instead, it generates genuine incremental value, thereby bolstering the platform’s aggregate revenue.

We also conducted a further analysis of the impact across different categories, with detailed results provided in Appendix C.3

**4.5.2 Item-side A/B Testing.** To further investigate the long-term impact on the lifecycle of new items, we conducted an item-side A/B test. Specifically, we tracked the performance of new items and compared their transaction scales after the 30th day post-launch period between the treatment group and the control group.

The experimental results reveal that new items supported by the GrowthGR strategy achieved a +20.0% uplift in  $\text{TI}@30$ . This significant improvement demonstrates that our framework does not merely provide a transient “injection” of traffic during the initial cold-start phase; instead, it fundamentally enhances the long-term traffic-carrying capacity and conversion efficiency of new items within the search ecosystem. From a strategic perspective, by successfully nurturing high-potential items into mature, high-performing ones, GrowthGR ensures a continuous supply of fresh content that aligns with evolving user preferences. This fosters a more dynamic and sustainable merchant ecosystem on Taobao, effectively bridging the gap between niche user needs and emerging item trends.

## 5 Conclusion

In this paper, we addressed the critical challenge of fostering sustainable growth for new items in large-scale e-commerce search scenarios. Industrial search systems are often plagued by the “Matthew Effect”, where traditional retrieval models favor head items with abundant historical data, leaving new items characterized by massive scale and extreme data sparsity struggling for visibility. To overcome these limitations, we proposed GrowthGR, a novel Multi-Value-Aware Generative Retrieval framework which can accurately quantify the incremental transaction value an item gains relative to its base growth and harmonize immediate conversion efficiency with predicted long-term potential. Extensive offline experiments and online A/B tests on Taobao involving billions of interactions demonstrate that GrowthGR significantly improves the discovery and growth rate of new items without sacrificing overall platform efficiency. Our experimental design, which integrates both user-side and item-side A/B testing, confirms the framework’s capability to mitigate popularity bias and support long-term ecosystem health.

## References

- [1] Keqin Bao, Jizhi Zhang, Yang Zhang, Wenjie Wang, Fuli Feng, and Xiangnan He. 2023. Tallrec: An effective and efficient tuning framework to align large language model with recommendation. In *Proceedings of the 17th ACM conference on recommender systems*. 1007–1014.
- [2] Hao Chen, Zefan Wang, Feiran Huang, Xiao Huang, Yue Xu, Yishi Lin, Peng He, and Zhoujun Li. 2022. Generative adversarial framework for cold-start item recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2565–2571.
- [3] Zeyu Cui, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. M6-rec: Generative pretrained language models are open-ended recommender systems. *arXiv preprint arXiv:2205.08084* (2022).
- [4] Jiaxin Deng, Shiyao Wang, Kuo Cai, Lejian Ren, Qigen Hu, Weifeng Ding, Qiang Luo, and Guorui Zhou. 2025. Onerec: Unifying retrieve and rank with generative recommender and iterative preference alignment. *arXiv preprint arXiv:2502.18965* (2025).
- [5] Ali Mamdouh Elkahky, Yang Song, and Xiaodong He. 2015. A multi-view deep learning approach for cross domain user modeling in recommendation systems. In *Proceedings of the 24th international conference on world wide web*. 278–288.
- [6] Francesco Fabbri, Maria Luisa Croci, Francesco Bonchi, and Carlos Castillo. 2022. Exposure inequality in people recommender systems: The long-term effects. In *Proceedings of the international AAAI conference on web and social media*, Vol. 16. 194–204.
- [7] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In *Proceedings of the 16th ACM conference on recommender systems*. 299–315.
- [8] Shijie Geng, Juntao Tan, Shuchang Liu, Zuohui Fu, and Yongfeng Zhang. 2023. Vip5: Towards multimodal foundation models for recommendation. *arXiv preprint arXiv:2305.14302* (2023).
- [9] Ruidong Han, Bin Yin, Shangyu Chen, He Jiang, Fei Jiang, Xiang Li, Chi Ma, Mincong Huang, Xiaoguang Li, Chunzhen Jing, et al. 2025. Mtgr: Industrial-scale generative recommendation framework in meituan. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*. 5731–5738.
- [10] Zhankui He, Zhouhang Xie, Rahul Jha, Harald Steck, Dawen Liang, Yesu Feng, Bodhisattwa Prasad Majumder, Nathan Kallus, and Julian McAuley. 2023. Large language models as zero-shot conversational recommenders. In *Proceedings of the 32nd ACM international conference on information and knowledge management*. 720–730.
- [11] Feiran Huang, Yuanchen Bei, Zhenghang Yang, Junyi Jiang, Hao Chen, Qijie Shen, Senzhang Wang, Fakhri Karray, and Philip S Yu. 2025. Large Language Model Simulator for Cold-Start Recommendation. In *Proceedings of the Eighteenth ACM International Conference on Web Search and Data Mining*. 261–270.
- [12] Feiran Huang, Zefan Wang, Xiao Huang, Yufeng Qian, Zhetao Li, and Hao Chen. 2023. Aligning distillation for cold-start item recommendation. In *Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval*. 1147–1157.
- [13] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)* 20, 4 (2002), 422–446.
- [14] Hao Jiang, Chuanzhen Li, Juanjuan Cai, Runyu Tian, and Jingling Wang. 2023. Self-supervised contrastive enhancement with symmetric few-shot learning towers for cold-start news recommendation. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 945–954.
- [15] Jinri Kim, Eungi Kim, Kwangeun Yeo, Yujin Jeon, Chanwoo Kim, Sewon Lee, and Joonseok Lee. 2024. Content-based graph reconstruction for cold-start item recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1263–1273.
- [16] Doyup Lee, Chiheon Kim, Saehoon Kim, Minsu Cho, and Wook-Shin Han. 2022. Autoregressive image generation using residual quantization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 11523–11532.
- [17] Jiacheng Li, Ming Wang, Jin Li, Jinmiao Fu, Xin Shen, Jingbo Shang, and Julian McAuley. 2023. Text is all you need: Learning language representations for sequential recommendation. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1258–1267.
- [18] Jinming Li, Wentao Zhang, Tian Wang, Guanglei Xiong, Alan Lu, and Gerard Medioni. 2023. GPT4Rec: A generative framework for personalized recommendation and user interests interpretation. *arXiv preprint arXiv:2304.03879* (2023).
- [19] Zida Liang, Changfa Wu, Dunxian Huang, Weiqiang Sun, Ziyang Wang, Yuliang Yan, Jian Wu, Yuning Jiang, Bo Zheng, Ke Chen, et al. 2025. Tbgrecall: A generative retrieval model for e-commerce recommendation scenarios. In *Proceedings of the 34th ACM International Conference on Information and Knowledge Management*. 5863–5870.
- [20] Chi Liu, Jiangxia Cao, Rui Huang, Kai Zheng, Qiang Luo, Kun Gai, and Guorui Zhou. 2024. KuaiFormer: Transformer-Based Retrieval at Kuaishou. *arXiv preprint arXiv:2411.10057* (2024).
- [21] Taichi Liu, Chen Gao, Zhenyu Wang, Dong Li, Jianye Hao, Depeng Jin, and Yong Li. 2023. Uncertainty-aware Consistency Learning for Cold-Start Item Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2466–2470.
- [22] Weiming Liu, Chaochao Chen, Xinting Liao, Mengling Hu, Jiajie Su, Yanchao Tan, and Fan Wang. 2024. User distribution mapping modelling with collaborative filtering for cross domain recommendation. In *Proceedings of the ACM Web Conference 2024*. 334–343.
- [23] Yunze Luo, Yuezihan Jiang, Yinjie Jiang, Gaode Chen, Jingchi Wang, Kaigui Bian, Peiyi Li, and Qi Zhang. 2025. Online item cold-start recommendation with popularity-aware meta-learning. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*. 927–937.
- [24] Xiao Ma, Liqin Zhao, Guan Huang, Zhi Wang, Zelin Hu, Xiaoqiang Zhu, and Kun Gai. 2018. Entire space multi-task model: An effective approach for estimating post-click conversion rate. In *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. 1137–1140.
- [25] Haoyu Pang, Fausto Giunchiglia, Ximing Li, Renchu Guan, and Xiaoyue Feng. 2022. PNMTA: A pretrained network modulation and task adaptation approach for user cold-start recommendation. In *Proceedings of the ACM Web Conference 2022*. 348–359.
- [26] Shashank Rajput, Nikhil Mehta, Anima Singh, Raghunandan Hulikal Keshavan, Trung Vu, Lukasz Heldt, Lichan Hong, Yi Tay, Vinh Tran, Jonah Samost, et al. 2023. Recommender systems with generative retrieval. *Advances in Neural Information Processing Systems* 36 (2023), 10299–10315.
- [27] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Yang Wu, et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300* (2024).
- [28] Qijie Shen, Yuanchen Bei, Zihong Huang, Jialin Zhu, Keqin Xu, Boya Du, Jiawei Tang, Yuning Jiang, Feiran Huang, Xiao Huang, et al. 2025. AliBoost: Ecological Boosting Framework in Alibaba Platform. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 2*. 4827–4838.
- [29] Xuehan Sun, Tianyao Shi, Xiaofeng Gao, Yanrong Kang, and Guihai Chen. 2021. FORM: follow the online regularized meta-leader for cold-start recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1177–1186.
- [30] Jizhe Wang, Pipei Huang, Huan Zhao, Zhibo Zhang, Binqiang Zhao, and Dik Lun Lee. 2018. Billion-scale commodity embedding for e-commerce recommendation in alibaba. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 839–848.
- [31] Xuansheng Wu, Huachi Zhou, Yucheng Shi, Wenlin Yao, Xiao Huang, and Ninghao Liu. 2024. Could small language models serve as recommenders? towards data-centric cold-start recommendation. In *Proceedings of the ACM Web Conference 2024*. 3566–3575.
- [32] Xiaolong Xu, Hongsheng Dong, Lianying Qi, Xuyun Zhang, Haolong Xiang, Xiaoyu Xia, Yanwei Xu, and Wanchun Dou. 2024. Cmlrec: Cross-modal contrastive learning for user cold-start sequential recommendation. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1589–1598.
- [33] Guipeng Xu, Chen Lin, Wanxian Guan, Jinping Gou, Xubin Li, Hongbo Deng, Jian Xu, and Bo Zheng. 2023. E-commerce search via content collaborative graph neural network. In *Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining*. 2885–2897.
- [34] Jiaqi Zhai, Lucy Liao, Xing Liu, Yueming Wang, Rui Li, Xuan Cao, Leon Gao, Zhaojie Gong, Fangda Gu, Michael He, et al. 2024. Actions speak louder than words: Trillion-parameter sequential transducers for generative recommendations. *arXiv preprint arXiv:2402.17152* (2024).
- [35] Chao Zhang, Shiwei Wu, Haoxin Zhang, Tong Xu, Yan Gao, Yao Hu, and Enhong Chen. 2024. Notellm: A retrievable large language model for note recommendation. In *Companion Proceedings of the ACM Web Conference 2024*. 170–179.
- [36] Weizhi Zhang, Yuanchen Bei, Liangwei Yang, Henry Peng Zou, Peilin Zhou, Aiwei Liu, Yinghui Li, Hao Chen, Jianling Wang, Yu Wang, et al. 2025. Cold-start recommendation towards the era of large language models (llms): A comprehensive survey and roadmap. *arXiv preprint arXiv:2501.01945* (2025).
- [37] Wenlin Zhang, Chuhan Wu, Xiangyang Li, Yuhao Wang, Kuicai Dong, Yichao Wang, Xinyi Dai, Xiangyu Zhao, Huifeng Guo, and Ruiming Tang. 2025. Llmtrerec: Unleashing the power of large language models for cold-start recommendations. In *Proceedings of the 31st International Conference on Computational Linguistics*. 886–896.
- [38] Zhaoqi Zhang, Haolei Pei, Jun Guo, Tianyu Wang, Yufei Feng, Hui Sun, Shaowei Liu, and Aixun Sun. 2025. OneTrans: Unified Feature Interaction and Sequence Modeling with One Transformer in Industrial Recommender. *arXiv preprint arXiv:2510.26104* (2025).
- [39] Zhihui Zhou, Lilin Zhang, and Ning Yang. 2023. Contrastive collaborative filtering for cold-start item recommendation. In *Proceedings of the ACM web conference 2023*. 928–937.



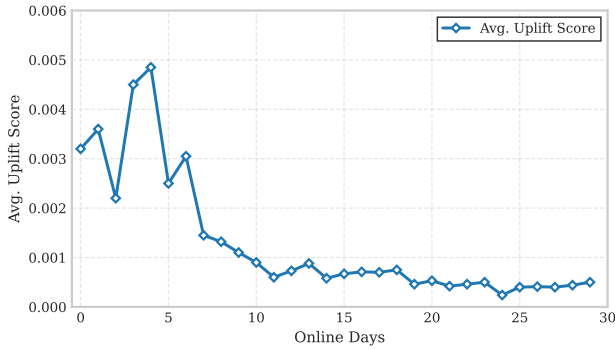


Figure 6: Average Uplift Score Across Different Online Days.

Table 5: Performance of different rollout strategies.

Strategy	all-net		long-term	
	Recall@10	Recall@1000	Recall@10	Recall@1000
Top-p Sampling	0.2568	<b>0.7147</b>	0.4440	0.8991
Beam Search	<b>0.2722</b>	0.7138	<b>0.4459</b>	<b>0.9064</b>

in their early stages, which aligns with empirical intuition in the e-commerce domain.

## C.2 Analysis of Rollout Strategies

In this section, we investigate the impact of different rollout strategies during the Preference Alignment of GrowthGR. Since the generative retrieval process relies on a robust exploration-exploitation balance to discover new items with high growth potential, the choice of decoding strategy during training is critical. We specifically compare two widely adopted strategies: Top-p Sampling and Beam Search. For Top-p Sampling, we set  $top-p = 0.95$ ,  $temperature = 1.0$ , and  $n = 5$ . For Beam Search, the beam size  $n$  is set to 5. The performance comparison across all-net and long-term metrics is summarized in Table 5.

Our experimental results reveal a clear trade-off between top-tier retrieval performance and long-tail coverage. In the all-net scenario, Beam Search demonstrates superior performance in top-tier metrics, achieving a Recall@10 of 0.2722, outperforming Top-p Sampling (0.2568). This suggests that Beam Search is highly effective at identifying the most probable candidates, which typically align with head items already favored by the current system distribution. Conversely, Beam Search exhibits comparable yet slightly lower performance in broader retrieval metrics, with Recall@1000 dropping to 0.7138, compared to Top-p Sampling’s 0.7147. This decline indicates that Beam Search may marginally suffer from likelihood bias, often neglecting the long-tail candidates that characterize new item distributions.

Interestingly, Beam Search maintains a slight lead in most long-term metrics, such as a Recall@1000 of 0.9064. This is likely because the ItemLTV module provides a strong prior for high-potential items, allowing even exploitative decoding to capture these growth signals effectively.

Table 6: Performance comparison across different categories.

Categories	PV (%)	GMV (%)
Apparel & Fashion	1.49	2.72
Sports & Outdoors	1.45	11.98
Home & Living	2.09	3.83
Fast Moving Consumer Goods	1.47	4.05
3C & Digital	1.12	22.18
Stationery, Education & Fragrance	2.90	0.22
Health	1.21	0.46
Food & Fresh	2.20	-0.67
Jewelry & Accessories	1.50	8.21
Automobiles	1.36	-0.81
Pets	1.20	-2.62
Toys & Hobbies	0.50	-5.48
Home Furnishing	1.38	5.54
Flowers & Gardening	0.65	8.11
Major Appliances	1.79	1.72
Small Appliances	1.34	21.15

## C.3 Category-wise Analysis

The category-wise increments for new items are summarized in Table 6. These results provide a granular view of how GrowthGR influences the platform’s ecosystem. The experimental data reveals several key insights into the behavior of the Multi-Value-Aware retrieval framework:

- GrowthGR achieved positive PV increments across all 16 major categories. This demonstrates that the MultiGR module, powered by hierarchical semantic identifiers and generative modeling, effectively bypasses the historical engagement barriers that usually trap new items. Even in categories like Food & Fresh (+2.20%) and Stationery, Education & Fragrance (+2.90%), the system successfully injected significant new item traffic into the search results.
- In several high-value categories, the growth in GMV significantly outpaced the growth in PV. For instance, 3C & Digital saw a +22.18% GMV increase with only a +1.12% PV increase, while Small Appliances achieved +21.15% GMV growth. By aligning retrieval with these value-rich candidates, the framework does not just blindly boost new items but strategically surfaces those that provide the highest marginal return for the platform.
- In high-frequency, low-deliberation categories such as Food & Fresh and Toys & Hobbies, we observed a slight decline in immediate GMV (−0.67% and −5.48% respectively) despite notable PV increases. This phenomenon reflects the inherent challenge of new item distribution: in categories where consumers have extremely strong brand or “head-item” loyalty, introducing new candidates can lead to a temporary dilution of immediate conversion efficiency. However, from a sustainable growth perspective, this discovery cost is a necessary investment to prevent category stagnation and maintain long-term seller diversity.

By providing a PV increment ranging from +0.50% to +2.90%, the framework ensures that new items are no longer invisible in the

shadow of established head items. The substantial GMV gains in over 70% of the categories further prove that a significant portion of new items are high-quality candidates that were previously underexposed due to data sparsity. By balancing immediate efficiency

with multiple value dimensions, GrowthGR creates a more dynamic and inclusive search environment on Taobao.