

From Head to Tail: LLM-Augmented Sequential Recommendation and Trade-off Optimization

Adeep Hande¹, Shaghayegh Agah¹, Neeraj Sharma¹, Karina Halevy², Sardar Hamidian¹
¹Applied AI Research, Comcast, USA ²Department of Computer Science, Carnegie Mellon University, USA
{aadeep_hande2, shaghayegh_agah, neeraj_sharma, sardar_hamidian}@comcast.com, khalevy@cmu.edu

Abstract

Recommender systems exhibit popularity bias, often under-serving long-tail and cold-start items with sparse interaction data. We propose a group-specific data augmentation approach that leverages LLMs to generate synthetic interactions, formulated as a constrained optimization problem, maximizing group-specific NDCG@20 and coverage while bounding overall NDCG@20 degradation within an acceptable threshold. Items are partitioned by interaction frequency to find the optimal augmentation strategy for each group. We further study combining the augmentation across groups simultaneously and find that while individual optimal strategies navigate the multi-group augmentation design, they are not simply additive. We further ablate on the LLM-based augmentation itself, comparing LLM scales, augmentation configuration under token constraints, finding that augmentation quality scales with contextual richness, albeit at a proportionally higher cost. Experiments with SASRec demonstrate the effectiveness of the proposed approach, achieving over 100% and 60% improvement in group-specific NDCG@20 for extreme and regular long-tail items, respectively, on our industrial dataset, and 40%, 25% for MovieLens-32M, while increasing their coverage without significant overall degradation.

CCS Concepts

• **Computing methodologies** → Natural language generation.

Keywords

LLM, Data Augmentation, Popularity Bias, Recommendations

ACM Reference Format:

Adeep Hande¹, Shaghayegh Agah¹, Neeraj Sharma¹, Karina Halevy², Sardar Hamidian¹. 2018. From Head to Tail: LLM-Augmented Sequential Recommendation and Trade-off Optimization. In *Proceedings of GRAIL 2026: Workshop on Generative, Retrieval-augmented, and Agentic Intelligence for Personalization (GRAIL @ CIKM '26)*. ACM, New York, NY, USA, 7 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Collaborative filtering models learn latent item representations from user-item interactions, but perform poorly when items have inherently sparse interaction data [10, 27]. Popularity bias creates

a feedback loop that systematically suppresses exposure for items with insufficient co-occurrence signals for both of long-tail and cold-start content [19, 37]. Within our industrial dataset (PRD) and ML-32M [23], the top 15% of items account for 95% and 98% of all interactions, respectively. Baseline SASRec [13] exhibits 92-95% drop from head to tail (Fig 1), with cold-start items receiving zero exposure.

Traditional approaches such as self-supervised learning [36], contrastive sequence augmentation [30], and adaptive negative sampling [4] address the interaction sparsity while operating within the existing interaction distribution without introducing new exposure for underrepresented items. More recently, LLM-based approaches have been introduced within the recommendation paradigm, generating pairwise sequences [28], augmenting interaction graphs [31, 33], metadata-based item enrichment [3], and serving LLMs as recommenders for cold-start items [22]. Kim et al. [14] align LLM token spaces with item embeddings, while Huang et al. [11] simulate user item interactions from filtered candidate sets, and Wang et al. [26] cluster LLM recommendations. Other optimizations include training an RL policy for efficient user selection in LLM augmentation [24] and post-training LLMs for cold-start reasoning [15]. In-context learning has also been applied to build synthetic interaction matrices for cold-start users [2, 17].

Recent work has focused more on addressing long-tail recommendations, improving collaborative signals via self-distillation with semantic embeddings [18], and retrieval-augmented LLMs [29], both enhancing existing representations. Towards generating synthetic interactions, TADA [5] introduces tail-aware augmentation with heuristic sequence updates. Synthetic interactions generated with LLM have proven to serve as a reliable augmentation signal on a scale [8, 22, 26], reinforcing that LLMs possess sufficient reasoning to act as effective data augmenters. However, these approaches either treat the long-tail as a monolith or do not characterize the tradeoff between overall and group-specific ranking quality under augmentation. Wang et al. [25] underscores the importance of exploring cold-start items, but does not differentiate between popularity tiers. We build on this by introducing a percentile-based grouping strategy over the interaction distribution, augmenting user histories with LLM-generated interactions drawn from group-specific candidate pools, and analyzing resulting trade-off between overall and group-specific ranking quality.

Our contributions are summarized as follows:

- We propose a group-wise targeted augmentation strategy that classifies items by interaction frequency to tailor the augmenting approach for each group, adapting to the dataset's popularity distribution.
- We characterize the trade-off between group-specific and overall metrics by estimating the Pareto frontier and examining their

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.

GRAIL @ CIKM '26, Rome, Italy

© 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/XXXXXXX.XXXXXXX>

correlated regions, demonstrating that each group induces a distinct trade-off curve determined by its representation in the historical data. This highlights that a single strategy does not generalize across groups and that one-size-fits-all approaches are suboptimal.

- We demonstrate that individually optimal per-group strategies do not compose into a globally optimal solution, confirming cross-group augmentation as a multi-objective optimization problem requiring joint calibration.
- We perform a comprehensive ablation study over candidate selection, LLM scale, insertion position, augmentation percentage, and configurations, introducing a stepwise experimentation framework that narrows the design space without exhaustive search.

2 Methodology

2.1 Dataset and Grouping Strategies

We evaluate MovieLens-32M (ML-32M)[23], and our Proprietary Real-world Dataset (PRD) comprising user watch history (TV shows, movies, sports, etc.) from a large-scale entertainment platform. We partition the item catalog in each dataset $\mathcal{I}$ into item-interaction frequency groups at the 75th and 95th percentile computed in log-space over each dataset’s interaction distribution. We select these boundaries based on the baseline SASRec’s performance across the candidate splits, as this partition produces groups with measurably distinct NDCG@20: G_3 produce closer to the head, G_2 has close to a 95% drop, while G_0/G_1 score near-zero performance (Fig 1).




Figure 1: Original Group-wise NDCG@20 (GN@20) by item popularity group on ML-32M (left) and PRD (right)

Unlike fixed threshold approaches that define grouping by arbitrary interaction thresholds, these percentile-based groupings adapt to each dataset’s interaction distribution, as the same split yields different absolute thresholds on ML-32M and PRD, reflecting their distinct popularity curves, where G_1 (0-75th) and G_2 (75th - 95th) capture extreme and regular long-tail items respectively, while G_3 is the popular tier and G_0 are cold-start items.

2.2 Augmentation Pipeline

Llama-3.3-70B-Instruct [9] serves as the default synthetic interaction generator. For each target group $G_k \in \{G_0, G_1, G_2\}$, the pipeline randomly samples a proportion (p) of users and augments each with a single synthetic item ($i = 1$) via LLM, selected from among $c = 10$ candidates drawn from G_k . The augmentation per user is motivated by previous literature [24, 26], where a single synthetic interaction per user minimizes the risk of shifting the users’

Table 1: Group distribution of items (full labeled dataset) and training-data interaction ranges across datasets

Group	Percentile	ML-32M		PRD	
		Items	Range	Items	Range
G_0	N/A	71	0	2,755	0
G_1	0-75 th	62,645	[1-23]	225,806	[1-30]
G_2	75 th -95 th	16,414	[23 - 731]	60,930	[30-302]
G_3	95 th +	5,124	731+	16,671	302+




Figure 2: Overview of the data augmentation pipeline.

preference distribution while still providing a meaningful signal to the base recommender. Candidate selection is performed via cosine similarity between Linq-Embed-Mistral [12] item embeddings and mean-pooled users’ representation on PRD, determined by the encoder ablation in Table 2. On ML-32M, where item metadata is insufficient for meaningful embedding construction, random sampling is used to select candidates for augmentation. The LLM selects from a closed candidate pool represented in the prompt (Appendix A), eliminating out-of-catalog selections. The selected item is inserted at S_{n-2} , in the user’s interaction sequence $S_u = (s_1, s_2, \dots, s_n)$ immediately preceding the validation (s_{n-1}) and test (s_n) items. SASRec is retrained from scratch on the augmented interaction matrix. All metrics are averaged over three independent runs. The prompt templates and implementation details are provided in Appendix A. We formulate evaluation as a constrained optimization, group-specific NDCG@20 (GN@20) and group coverage GCov@100 that measures each group’s share of top-100 recommendations per user, are the objectives to maximize, subject to overall NDCG@20 (N@20) degradation remaining within an acceptable threshold.

2.3 Experiment Designs

Given the vast search space, optimizing all variables concurrently results in a combinatorial explosion. Therefore, we fixed specific variables such as the quantity and positioning of the augmented data, following prior work [17, 24, 26] and prioritized reporting the most impactful axes of variation. To systematically navigate this search space, we utilized a coordinate descent optimization strategy: sequentially fixing hyperparameters to isolate and optimize individual variables step-by-step. This greedy optimization pipeline guarantees practical reproducibility for any dataset and model. Sequential studies designed to isolate the effects of candidate strategy, augmentation proportion, and configuration are discussed below.

1) Candidate Selection Optimization: We compare embedding encoders from the MTEB leaderboard [6, 20]: Linq-Embed-Mistral (Mistral-7B) [12], qwen3-4b [32], gte-large-en [16, 34] on G_2 at $p = 5\%$, selecting the encoder that maximizes group-specific NDCG@20. We validate the best encoder against random sampling across several augmentation proportions and quantify the LLM’s marginal contribution by comparing paired settings with and without LLM reranking (Table 2 (b)). On ML-32M, we randomly sample the candidates due to lack of sufficient item metadata.

2) Pareto-front analysis via Augmentation-proportion: Given the best candidate selection strategy from (1), we vary p to identify both the group-specific Pareto front, finding the trade-off boundary, and a correlated region where targeted and guardrail metrics improve simultaneously.

3) Cross Augmentation Optimization: Depending on the industry’s objectives, optimizing performance across all groups may be necessary rather than focusing on a single specific group. we augment G_0 , G_1 , and G_2 simultaneously under varying augmentation splits, inspired by the previous step, to evaluate the optimal cross-group augmentation strategy.

4) LLM Design Space Optimization: We evaluate whether targeted solutions can be further improved using LLM optimization techniques under a cost constraint. We fix G_2 at $p = 5\%$, and perform ablation studies over LLM scale (8B, 70B, 405B), insertion position ($S_{n/2}$ vs S_{n-2}), user-profiles, and LLM-as-a-judge [7] to isolate each component’s individual contribution towards overall and group-wise metrics. We note that this specific setup serves as an illustrative benchmark to demonstrate the optimization framework, which can be similarly applied to other groups or extended across all groups.

3 Results

1) Candidate Selection Optimization: Table 2 (left) compares different embedding-based models for G_2 augmentation in PRD dataset at $p = 5\%$. The encoder determines the quality of the candidate pool presented to the LLM, providing it with an informed selection space by surfacing the most relevant items with respect to each user. Mistral-7B achieves the strongest balance between the overall and group-specific performance and consistently improves against random selection at all p values (Table 2 (right)), confirming embedding-based selections improve across all metrics. Paired with Llama-70B (Table-3), we further improve GN@20 by three times. Random selection degrades both metrics, as sampling from a larger vocabulary produces no semantic relevance to the user. *Mistral-7B* with Llama-70B is used as the default augmentation configuration for all downstream experiments. Embedding-based selection alone yields +5.59% N@20 and +15.97% GN@20 (Table 2), validating that cosine similarity already surfaces relevant candidates. Appending LLM to the pipeline boosts GN@20 to +48.28%. In addition, the first two rows in Table 3 shows that LLM augments better than random in selecting augmented data.

2) Pareto-front analysis via Augmentation-proportion: As shown in Figure 3-a, each group traces a distinct tradeoff curve and optimal point governed by both deployment targeted metrics objectives and acceptable degradation in overall performance. We

Table 2: Candidate selection optimization (PRD) (G_2): (left) encoder-only (right) $\Delta GN@20(\%)$: best encoder vs. random.

embeddings	N@20	GN@20	p	random	mistral-7b
mistral-7B	+6.36	+15.97	1%	-0.79	+ 8.14
Qwen-3B	+4.79	+6.42	5%	+24.87	+48.28
GTE-en	-6.47	+27.49	10%	+31.15	+62.42

Table 3: Comparing two baselines (no augmenters) vs LLM reranking across candidate selection strategies G_2 : $p = 5\%$.

Candidate Selection	Augmenter	$\Delta N@20(\%)$	$\Delta GN@20(\%)$
Random	-	-3.44	-13.04
Random	Llama-70B	-4.03	+24.87
mistral-7B	-	+5.59	+15.97
mistral-7B	Llama-70B	-2.09	+48.28

define a guardrail constraint of $\leq 5\%$ N@20 degradation for identifying viable configurations. Within PRD, both G_2 and G_1 show positive correlated and tradeoff regions. In G_2 , at $p = 1\%$, overall and group NDCG metrics improve by +3.8% and +8.14% respectively, targeted metrics continue to rise (GN. of +62% and GCov. of +5% in $p = 15\%$) while N@20 remains within -2% of the baseline throughout, suggesting that G_2 is the most tolerant group to augmentation. Augmenting items that already possess partial learned representations provides SASRec with additional informative signals, allowing the model to learn their representations better, without introducing a sufficient distributional shift that degrades the overall performance. G_1 items are extremely sparse in the training data, resulting in under-trained item embeddings sensitive to small volumes of synthetic augmentations. At lower p , synthetic interactions can be insufficient to overcome the noise, and their effect becomes volatile and smaller. At higher volumes ($p \geq 10\%$), G_1 shows consistent group-wise improvement, at the cost of sacrificing the overall performance.

The lack of any N@20 improvements for ML-32M in G_1 and G_2 plots can be attributed to the candidate selection strategy. For G_1 , ML-32M evaluates at GN@100 due to limited test users (660) as GN@20 yields near-zero across all configurations. Random candidate selection can introduce noise that prevents joint metrics optimization. Due to the varying distribution of long-tail groups across datasets, the gain in targeted metrics may differ; however, consistent trends are observed across them. For example, given a fixed p , G_1 shows a larger improvement in the metric than G_2 across both datasets, as it starts from a substantially lower baseline. We select $p = 3\%$ as G_1 ’s operating point for further cross-group augmentation experiments as it falls in the correlated region while having the lowest augmentation budget to minimize the overall augmentation proportions [5].

Cold-start items retain their random initialization throughout training, receiving no gradient updates from real interactions, and augmentation provides the only source of learning signal for these embeddings. At $p \leq 5\%$, the augmentation is small enough for the model to improve both targeted metrics while maintaining the N@20 within the acceptable threshold. However, beyond this,




Figure 3: N@20-GN@20 and N@20-GCov@100 tradeoffs across different (p). The black dotted line indicates the predefined constraint for the guardrail metric. With a zero baseline for G_k , improvements are reported relative to 1% (p). All reported numbers are averaged across three runs and are statistically significant.

Table 4: PRD Cross augmentation opt. via p distribution.

$[G_0(p), G_1(p), G_2(p)]$	N@20	G0=N@20	G1=N@20	G2=N@20
[5, 3, 15]	-24.49%	2.96x	+53.95%	+45.67%
[5, 5, 5]	-21.38%	5.74x	+11.31%	+19.86%
[1.66, 1.66, 1.66]	-1.06%	2.87x	-1.46%	+20.59%
[1.92, 1.16, 1.92]	-11.04%	2x	-7.69%	+38.41%
[3.84, 2.32, 3.84]	-14.95%	4.17x	-41.06%	+17.95%
[5, 3, 5]	-13.8%	5.57x	+32.24%	+19.24%
[5.77, 3.46, 5.77]	-20.6%	6.61x	+71.52%	13.44%

Table 5: PRD augmenter ablation performance (%)

Augmenter	$\Delta \text{N@20}(\%)$	$\Delta \text{GN@20}(\%)$	token (x)
Llama-70B-default (baseline)	-2.09	<u>+48.28</u>	1.0x
Llama-8B	-0.15	+43.45	1.0x
Llama-405B	<u>+3.11</u>	<u>+48.55</u>	1.0x
Llama-70B (early: $S_{n/2}$)	+0.06	+31.67	1.0x
LLM-as-a-judge [7]	+6.36	+34.93	2.06x
Llama-70B + user profiles	-0.16	+34.09	1.57x
Llama-70B + metadata	<u>+3.78</u>	+60.39	5.12x

the distribution shift steepens the tradeoff sharply. The optimal augmentation is $p = 5\%$, balancing a 4x GN@20 and 2x GCov@100 improvement against $\leq 5\%$ N@20 degradation. On ML-32M, the correlated regions can only be observed at G_0 for $p = 1\%$, owing to significantly smaller candidate pool for augmentation (71 vs 2.7k).

3) Cross Augmentation Optimization: We evaluate whether per-group augmentation strategies compose additively. Naively combining independent optimal proportions (23% total) degrades N@20 by -24.49%. This drop may stem from the total augmentation volume itself, their interaction effects, or both. We cap the overall augmentation proportion to $\leq 15\%$ as this represents the maximum

evaluated in our per group analysis (Figure 3), to keep the guardrail metric decline, N@20 below the certain threshold and measure additive effects in joint optimization.

The non-uniform p allocation across groups and their respective shares of the total augmentation budget is motivated by patterns observed in single-group trade-off optimization. At 5% total cross-augmentation with equal splits, both G_0 and G_2 improve across 3rd and 4th rows, while non-uniform allocation yields a greater overall gain. However, G_1 remains sensitive to the splits, improving at higher p that is consistent with our findings that G_1 requires a minimum augmentation volume to produce stable gains. No configuration simultaneously improves all groups without overall degradation. Applying group specific p inspired by previous results, last two rows, simultaneously improve all group specific metrics. However, to mitigate degradation in guardrail metrics, further calibration of the joint augmentation is required, framing it as a multi-objective optimization problem.

4) LLM Design Space Optimization We fix G_2 at 5% augmentation across all ablations for consistency and evaluate against several augmentation strategies as presented in Table 5. Insertion position significantly affects efficacy, with early-augmented diluting their influence as SASRec weighs heavily toward more recent items [13, 21]. User-profile generation [1] and llm-as-a-judge [7] adopt a least-to-most prompting paradigm [35], in which intermediate structure are leveraged to iteratively refine the final candidate selection, thereby enhancing overall ranking performance. Augmenting the prompt with additional metadata achieves the highest GN@20, indicating that richer contextual information improves the quality of the resulting augmentations. While these augmentation configurations yield stronger gains, the optimal configuration is itself a tradeoff between guardrail metrics and cost/computational

scalability. The base configuration (Mistral-7B + Llama-70B) offers the most favorable cost-quality ratio for production deployment, while metadata enrichment achieves the strongest group-specific improvement at significantly higher computational cost. This motivates task-specific fine-tuning of smaller LLMs, which may yield both stronger augmentation reasoning and lower inference cost compared to prompting a general purpose large language model [15].

4 Conclusion

We presented a group-wise LLM augmentation framework for sequential recommendation and evaluated it on ML-32M and a large proprietary dataset by separating items into cold-start (G0), extreme long-tail (G1), regular long-tail (G2), and popular tiers (G3). Across candidate selection, augmentation ratio, cross-group splits, and LLM-side design choices, the results show that augmentation is group dependent rather than governed by a single strategy. Embedding-based candidate selection is more reliable than random selection when semantic retrieval is available, and the underrepresented groups differ in how much augmentation they can absorb and in how they trade off with overall NDCG@20. The LLM ablations also show that augmentation quality depends not only on model scale but also on prompting and insertion choices, with clear effects on GN@20, coverage, cost, and guardrail performance. Just as importantly, the best single-group settings do not carry over directly to the joint case, where cross-group augmentation remains a constrained multi-objective problem. A natural next step is to make that joint calibration explicit and test whether the same patterns hold across other recommenders and datasets.

References

- [1] Seunghwan Bang and Hwanjun Song. 2025. Llm-based user profile management for recommender system. *arXiv preprint arXiv:2502.14541* (2025).
- [2] Shangkun Che, Minjia Mao, and Hongyan Liu. 2024. New community cold-start recommendation: A novel large language model-based method. (2024).
- [3] Lei Chen, Chen Gao, Xiaoyi Du, Hengliang Luo, Depeng Jin, Yong Li, and Meng Wang. 2025. Enhancing ID-based Recommendation with Large Language Models. *ACM Trans. Inf. Syst.* 43, 5, Article 112 (July 2025), 30 pages. doi:10.1145/3704263
- [4] Yongjun Chen, Jia Li, Zhiwei Liu, Nitish Shirish Keskar, Huan Wang, Julian McAuley, and Caiming Xiong. 2022. Generating Negative Samples for Sequential Recommendation. *arXiv:2208.03645 [cs.LG]* <https://arxiv.org/abs/2208.03645>
- [5] Yizhou Dang, Zhifu Wei, Minhan Huang, Lianbo Ma, Jianzhe Zhao, Guibing Guo, and Xingwei Wang. 2026. Tail-aware data augmentation for long-tail sequential recommendation. In *Proceedings of the ACM Web Conference 2026*. 5886–5897.
- [6] Kenneth Enevoldsen, Isaac Chung, Imene Kerboua, Márton Kardos, Ashwin Mathur, David Stap, Jay Gala, Wissam Sibli, Dominik Krzemiński, Genta Indra Winata, Saba Sturua, Saiteja Utpala, Mathieu Ciancone, Marion Schaeffer, Gabriel Sequeira, Diganta Misra, Shreeya Dhakal, Jonathan Rystrom, Roman Solomatin, Ömer Çağatan, Akash Kundu, Martin Bernstorff, Shitao Xiao, Akshita Sukhlecha, Bhavish Pawha, Rafał Poświata, Kranthi Kiran GV, Shawon Ashraf, Daniel Auras, Björn Plüster, Jan Philipp Harries, Loïc Magne, Isabelle Mohr, Mariya Hendriksen, Dawei Zhu, Hippolyte Gisserot-Boukhlef, Tom Aarsen, Jan Kostkan, Konrad Wojtasik, Taemin Lee, Marek Suppa, Crystina Zhang, Roberta Rocca, Mohammed Hamdy, Andrianos Michail, John Yang, Manuel Fayse, Aleksei Katolin, Nandan Thakur, Manan Dey, Dipam Vasani, Pranjal Chitale, Simone Tedeschi, Nguyen Tai, Artem Snegirev, Michael Günther, Mengzhou Xia, Weijia Shi, Xing Han Lu, Jordan Clive, Gayatri Krishnakumar, Anna Maksimova, Silvan Wehrli, Maria Tikhonova, Henil Panchal, Aleksandr Abramov, Malte Ostendorff, Zheng Liu, Simon Clematide, Lester James Miranda, Alena Fenogenova, Guangyu Song, Ruqiyah Bin Safi, Wen-Ding Li, Alessia Borghini, Federico Cassano, Hongjin Su, Jimmy Lin, Howard Yen, Lasse Hansen, Sara Hooker, Chenghao Xiao, Vaibhav Adlakha, Orion Weller, Siva Reddy, and Niklas Muennighoff. 2025. MMTEB: Massive Multilingual Text Embedding Benchmark. *arXiv preprint arXiv:2502.13595* (2025). doi:10.48550/arXiv.2502.13595
- [7] Francesco Fabbri, Gustavo Penha, Edoardo D'Amico, Alice Wang, Marco De Nadai, Jackie Doremus, Paul Giglioli, Andreas Damianou, Oskar Stål, and Mounia Lalmas. 2025. Evaluating podcast recommendations with profile-aware llm-as-a-judge. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems*. 1181–1186.
- [8] Natalija Glisovic, Martin Tegner, and Danica Kragic. 2026. Data Augmentation with LLMs for Cold Start Recommendation in E-Commerce. In *European Conference on Information Retrieval*. Springer, 94–99.
- [9] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783* (2024).
- [10] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web (Perth, Australia) (WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 173–182. doi:10.1145/3038912.3052569
- [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 (Hannover, Germany) (WSDM '25)*. Association for Computing Machinery, New York, NY, USA, 261–270. doi:10.1145/3701551.3703546
- [12] Jihoon Kwon Junseong Kim, Seolhwa Lee et al. 2024. Linq-Embed-Mistral:Elevating Text Retrieval with Improved GPT Data Through Task-Specific Control and Quality Refinement. Linq AI Research Blog. <https://getlinq.com/blog/linq-embed-mistral/>
- [13] Wang-Cheng Kang and Julian McAuley. 2018. Self-Attentive Sequential Recommendation. In *2018 IEEE International Conference on Data Mining (ICDM)*. 197–206. doi:10.1109/ICDM.2018.00035
- [14] Sein Kim, Hongseok Kang, Seungyeon Choi, Donghyun Kim, Minchul Yang, and Chanyoung Park. 2024. Large Language Models meet Collaborative Filtering: An Efficient All-round LLM-based Recommender System. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (Barcelona, Spain) (KDD '24)*. Association for Computing Machinery, New York, NY, USA, 1395–1406. doi:10.1145/3637528.3671931
- [15] Shijun Li, Yu Wang, Jin Wang, Ying Li, Joydeep Ghosh, and Anne Cocos. 2025. LLM Reasoning for Cold-Start Item Recommendation. *arXiv preprint arXiv:2511.18261* (2025).
- [16] Zehan Li, Xin Zhang, Yanzhao Zhang, Dingkun Long, Pengjun Xie, and Meishan Zhang. 2023. Towards general text embeddings with multi-stage contrastive learning. *arXiv preprint arXiv:2308.03281* (2023).
- [17] Xurong Liang, Vu Nguyen, Vuong Le, Paul Albert, and Julien Monteil. 2025. In-context Learning for Addressing User Cold-start in Sequential Movie Recommenders. In *Proceedings of the Nineteenth ACM Conference on Recommender Systems (RecSys '25)*. Association for Computing Machinery, New York, NY, USA, 979–982. doi:10.1145/3705328.3748109
- [18] Qidong Liu, Xian Wu, Yejing Wang, Zijian Zhang, Feng Tian, Yefeng Zheng, and Xiangyu Zhao. 2024. LLM-ESR: large language models enhancement for long-tailed sequential recommendation. In *Proceedings of the 38th International Conference on Neural Information Processing Systems (Vancouver, BC, Canada) (NIPS '24)*. Curran Associates Inc., Red Hook, NY, USA, Article 839, 27 pages.
- [19] Masoud Mansoury, Himan Abdollahpour, Mykola Pechenizkiy, Bamshad Mobasher, and Robin Burke. 2020. Feedback Loop and Bias Amplification in Recommender Systems. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (Virtual Event, Ireland) (CIKM '20)*. Association for Computing Machinery, New York, NY, USA, 2145–2148. doi:10.1145/3340531.3412152
- [20] Niklas Muennighoff, Nouamane Tazi, Loïc Magne, and Nils Reimers. 2022. MTEB: Massive Text Embedding Benchmark. *arXiv preprint arXiv:2210.07316* (2022). doi:10.48550/ARXIV.2210.07316
- [21] Jeonglyul Oh and Sungzoon Cho. 2024. Measuring recency bias in sequential recommendation systems. *arXiv preprint arXiv:2409.09722* (2024).
- [22] Scott Sanner, Krisztian Balog, Filip Radlinski, Ben Wedin, and Lucas Dixon. 2023. Large Language Models are Competitive Near Cold-start Recommenders for Language- and Item-based Preferences. In *Proceedings of the 17th ACM Conference on Recommender Systems (Singapore, Singapore) (RecSys '23)*. Association for Computing Machinery, New York, NY, USA, 890–896. doi:10.1145/3604915.3608845
- [23] Mark D. Smucker and Houman Chamani. 2025. Extending MovieLens-32M to Provide New Evaluation Objectives. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (Padua, Italy) (SIGIR '25)*. Association for Computing Machinery, New York, NY, USA, 3520–3529. doi:10.1145/3726302.3730328
- [24] Nachiket Subbaraman, Jaskinder Sarai, Aniruddh Nath, Lichan Hong, Lukasz Heldt, Li Wei, and Zhe Zhao. 2025. Selecting User Histories to Generate LLM Users for Cold-Start Item Recommendation. *arXiv:2511.21989 [cs.LG]* <https://arxiv.org/abs/2511.21989>
- [25] Dong Wang, Junyi Jiao, Arnab Bhadury, Yaping Zhang, Mingyan Gao, and Onkar Dalal. 2025. Item-centric Exploration for Cold Start Problem. In *Proceedings of the*

- Nineteenth ACM Conference on Recommender Systems (RecSys '25). Association for Computing Machinery, New York, NY, USA, 987–990. doi:10.1145/3705328.3748113
- [26] Jianling Wang, Haokai Lu, Yifan Liu, He Ma, Yueqi Wang, Yang Gu, Shuzhou Zhang, Ningren Han, Shuchao Bi, Lexi Baugher, Ed H. Chi, and Minmin Chen. 2024. LLMs for User Interest Exploration in Large-scale Recommendation Systems. In *Proceedings of the 18th ACM Conference on Recommender Systems (Bari, Italy) (RecSys '24)*. Association for Computing Machinery, New York, NY, USA, 872–877. doi:10.1145/3640457.3688161
- [27] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural Graph Collaborative Filtering. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval (Paris, France) (SIGIR'19)*. Association for Computing Machinery, New York, NY, USA, 165–174. doi:10.1145/3331184.3331267
- [28] Wei Wei, Xubin Ren, Jiabin Tang, Qinyong Wang, Lixin Su, Suqi Cheng, Junfeng Wang, Dawei Yin, and Chao Huang. 2024. LLMRec: Large Language Models with Graph Augmentation for Recommendation. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining (Merida, Mexico) (WSDM '24)*. Association for Computing Machinery, New York, NY, USA, 806–815. doi:10.1145/3616855.3635853
- [29] Junda Wu, Cheng-Chun Chang, Tong Yu, Zhankui He, Jianing Wang, Yupeng Hou, and Julian McAuley. 2024. CoRAL: Collaborative Retrieval-Augmented Large Language Models Improve Long-tail Recommendation. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (Barcelona, Spain) (KDD '24)*. Association for Computing Machinery, New York, NY, USA, 3391–3401. doi:10.1145/3637528.3671901
- [30] Xu Xie, Fei Sun, Zhaoyang Liu, Shiwen Wu, Jinyang Gao, Jiandong Zhang, Bolin Ding, and Bin Cui. 2022. Contrastive Learning for Sequential Recommendation. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*. 1259–1273. doi:10.1109/ICDE53745.2022.00099
- [31] Lanling Xu, Junjie Zhang, Bingqian Li, Jinpeng Wang, Sheng Chen, Wayne Xin Zhao, and Ji-Rong Wen. 2025. Tapping the Potential of Large Language Models as Recommender Systems: A Comprehensive Framework and Empirical Analysis. *ACM Trans. Knowl. Discov. Data* 19, 5, Article 105 (June 2025), 51 pages. doi:10.1145/3726871
- [32] An Yang, Anfeng Li, Baosong Yang, et al. 2025. Qwen3 Technical Report. arXiv:2505.09388 [cs.CL] <https://arxiv.org/abs/2505.09388>
- [33] Haobo Zhang, Qiannan Zhu, and Zhicheng Dou. 2025. Enhancing Reranking for Recommendation with LLMs through User Preference Retrieval. In *Proceedings of the 31st International Conference on Computational Linguistics*, Owen Rambow, Leo Wanner, Marianna Apidianaki, Hend Al-Khalifa, Barbara Di Eugenio, and Steven Schockaert (Eds.). Association for Computational Linguistics, Abu Dhabi, UAE, 658–671. <https://aclanthology.org/2025.coling-main.45/>
- [34] Xin Zhang, Yanzhao Zhang, Dingkun Long, Wen Xie, Ziqi Dai, Jialong Tang, Huan Lin, Baosong Yang, Pengjun Xie, Fei Huang, et al. 2024. mgte: Generalized long-context text representation and reranking models for multilingual text retrieval. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: Industry Track*. 1393–1412.
- [35] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc Le, et al. 2022. Least-to-most prompting enables complex reasoning in large language models. *arXiv preprint arXiv:2205.10625* (2022).
- [36] 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 *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (Virtual Event, Ireland) (CIKM '20)*. Association for Computing Machinery, New York, NY, USA, 1893–1902. doi:10.1145/3340531.3411954
- [37] Ziwei Zhu, Yun He, Xing Zhao, Yin Zhang, Jianling Wang, and James Caverlee. 2021. Popularity-Opportunity Bias in Collaborative Filtering. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining (Virtual Event, Israel) (WSDM '21)*. Association for Computing Machinery, New York, NY, USA, 85–93. doi:10.1145/3437963.3441820

A Appendix

A.1 Pipeline configuration

Table 6 shows the default configuration used across all experiments unless specified otherwise. All components are interchangeable, and our approach could be substituted with any base recommender, embedding encoder, or instruction-following LLM.

Table 6: Default pipeline configuration. All components are interchangeable.

Component	Implementation
Base recommender	SASRec [13]
LLM augmenter	Llama-3.3-70B-Instruct [9]
Candidate encoder	LinQ-Embed-Mistral [12]
Candidate selection	Cosine similarity (top-c)
Items per user	$i = 1$
Candidate pool size	$c = 10$
Insertion position	s_{n-2}

A.2 Percentile-Based Item Grouping

Algorithm 1 describes the grouping strategy. The procedure is agnostic to any dataset and could be applied to any interaction dataset without manual threshold selection. However, the optimal percentile boundaries may vary across domains and recommender architectures, and we recommend validating the split against the baseline group-level performance to ensure distinct tiers (Fig. 1).

Algorithm 1 Percentile-Based Item Grouping

Require: Training set $\mathcal{D}_{train}$, test/val items $\mathcal{I}_{test} \cup \mathcal{I}_{val}$

Ensure: Item groups $\{G_0, G_1, G_2, G_3\}$

Count interactions per item: $n_i \leftarrow |\{(u, i) \in \mathcal{D}_{train}\}|$

Compute log-counts: $\ell_i \leftarrow \log(n_i + 1)$

Remove outliers: exclude items where $|\ell_i - \mu| > 3\sigma$

Compute P_{75}, P_{95} over remaining ℓ_i

$G_1 \leftarrow \{i : \ell_i \leq P_{75}\}$

$G_2 \leftarrow \{i : P_{75} < \ell_i \leq P_{95}\}$

$G_3 \leftarrow \{i : \ell_i > P_{95}\}$

$G_0 \leftarrow \{i \in (\mathcal{I}_{test} \cup \mathcal{I}_{val}) : n_i = 0\}$

G_0 items are introduced into training histories via synthetic interactions, thus providing the model with its initial learning signal for these items. The evaluation measures ranking quality on held-out test sequences, and no test interaction is seen during data augmentation step.

A.3 Prompt Template

The following listwise prompt template is used for all LLM-based candidate generation experiments. The user’s interaction history (SHOWS), and candidate items (CANDIDATES) are populated per user.

I have previously watched the following shows: {SHOWS}. Given my taste, I have the following movies to choose from, written below in comma-separated fashion: {CANDIDATES}. Based on my watch history and preferences, please recommend which {i} movies you think I would enjoy the most, in descending order. Format your answer as a comma-separated list of {i} movie titles, with no other text or explanations.

Figure 4: Listwise prompt template for LLM-based candidate selection.

A.4 Coverage Analysis

Table 7 lists the shift in group-wise coverage when each group is augmented at its optimal proportion (p^*) from the per-group analysis. The PRD coverage analysis for the selected Pareto-optimal point (Figure 3-b) shows augmentation redistributes exposure from G_3 to the targeted group, increasing item diversity in recommendation outputs. Augmenting G_1 in PRD and G_0 in ML-32M also increases G_3 coverage; this happens due to the synthetic interaction that increases both overall and group metrics. Under the $\leq 5\%$ N@20 threshold, ML-32M's G_1 and G_2 , the optimal p^* is limited to $p = 1\%$, resulting in negligible coverage shifts.

Table 7: Group Coverage distribution shift for the selected Pareto-optimal (p) per Group ($G_k(p^*)$). Each row shows the coverage shift across all groups when augmenting the target group.

	$G_k(p^*)$	ΔG_0	ΔG_1	ΔG_2	ΔG_3
PRD	G_0 (5%)	+1.35	-0.38	+0.40	-1.38
	G_1 (3%)	0.0	+1.35	-2.77	+1.43
	G_2 (5%)	0.0	+0.46	+5.68	-6.14
ML	G_0 (3%)	+1.26	-3.92	-2.78	+5.43
	G_1 (1%)	0.0	+0.81	-5.16	+4.35
	G_2 (1%)	0.0	-0.90	+0.36	+0.54

A.5 Data Preprocessing

For both datasets, we follow SASRec style training and evaluation [13], with the penultimate for validation, last interaction for testing, and the rest for training. Interaction sequences are ordered chronologically and truncated or zero-padded to a maximum length of 150. Users with fewer than K interactions are excluded from both datasets.