

# HSTU-BLaIR: Lightweight Contrastive Text Embedding for Generative Recommender

Yijun Liu

yijunl@usc.edu

Ming Hsieh Department of Electrical and Computer Engineering  
University of Southern California  
Los Angeles, CA, United States

## Abstract

Recent advances in recommender systems have underscored the complementary strengths of generative modeling and pretrained language models. We propose HSTU-BLaIR, a hybrid framework that augments the Hierarchical Sequential Transduction Unit (HSTU)-based generative recommender with BLaIR, a lightweight contrastive text embedding model. This integration enriches item representations with semantic signals from textual metadata while preserving HSTU’s powerful sequence modeling capabilities.

We evaluate HSTU-BLaIR on two e-commerce datasets: three subsets from the Amazon Reviews 2023 dataset and the Steam dataset. We compare its performance against both the original HSTU-based recommender and a variant augmented with embeddings from OpenAI’s state-of-the-art text-embedding-3-large model. Despite the latter being trained on a substantially larger corpus with significantly more parameters, our lightweight BLaIR-enhanced approach—pretrained on domain-specific data—achieves better performance in nearly all cases. Specifically, HSTU-BLaIR outperforms the OpenAI embedding-based variant on all but one metric, where it is marginally lower, and matches it on another. These findings highlight the effectiveness of contrastive text embeddings in compute-efficient recommendation settings. Code is available at <https://github.com/snapfinger/HSTU-BLaIR>.

## Keywords

Generative recommendation, textual embedding, contrastive learning, personalization

## ACM Reference Format:

Yijun Liu. 2025. HSTU-BLaIR: Lightweight Contrastive Text Embedding for Generative Recommender. In *Proceedings of LLM4ECommerce Workshop at the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (LLM4ECommerce Workshop at KDD ’25)*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

## 1 Introduction

Recent progress in recommender systems has been fueled by advances in deep learning [3], and increasingly by techniques from

natural language processing [1, 5–7, 12, 13, 17, 23]. While the dominant industrial paradigm has long relied on multi-stage recommendation pipelines—where two-tower architectures perform efficient candidate retrieval from a massive corpus, followed by a ranking model, often an MLP, for fine-grained scoring [3, 8, 19, 20]—these systems often require extensive feature engineering and may struggle to capture complex temporal and semantic dependencies. In contrast, generative recommenders jointly model user interaction sequences and directly generate item predictions in an autoregressive fashion [15, 21, 22]. This paradigm enables richer temporal modeling and improved personalization. A key subtask within this space is sequential recommendation, which focuses on capturing the temporal dynamics of user interactions [10, 16].

Although recent work has begun to explore the integration of generative recommenders with pretrained language models, most existing approaches rely on prompting or fine-tuning large language models (LLMs) [5, 9, 11]. These models are often not specifically tailored for recommendation tasks and tend to be computationally inefficient—especially given that state-of-the-art LLMs typically contain billions of parameters [18]. As a result, the broader incorporation of pretrained language models into generative recommendation systems remains in its early stages, with critical challenges around scalability, domain adaptation, and efficient utilization of textual signals still largely unresolved.

On the other hand, self-supervised contrastive learning has shown strong potential for producing high-quality text embeddings that capture fine-grained semantic alignment [4, 6]. Generative recommenders, in parallel, excel at modeling structured user behavior over time. Combining these complementary strengths offers a promising direction to unify semantic understanding with behavioral dynamics, enabling more expressive, adaptable, and semantically informed recommendation systems.

In this work, we integrate two recent state-of-the-art models: BLaIR [6], a contrastive text encoder pretrained on user reviews and item metadata from the Amazon Reviews 2023 dataset, and HSTU [22], a generative model that formulates recommendation as sequential transduction, unifying retrieval and ranking through autoregressive modeling. BLaIR generates semantically rich textual representations from item metadata, while HSTU provides a scalable, end-to-end architecture for sequential recommendation. By fusing BLaIR’s semantic embeddings with HSTU’s behavioral modeling, our framework enhances personalization while maintaining scalability and interpretability—leveraging human-readable metadata and avoiding the computational overhead of large language models.

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).

LLM4ECommerce Workshop at KDD ’25, Toronto, ON, Canada

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>



## 2 Method

### 2.1 Base Model: Hierarchical Sequential Transduction Unit (HSTU)

We build on the generative recommender of Zhai et al. [22], which encodes a user’s interaction history with a *Hierarchical Sequential Transduction Unit* (HSTU). Stacked Transformer layers with hierarchical attention and relative positional bias map the sequence into autoregressive hidden states, which are subsequently scored against candidate items by a joint retrieval–ranking interaction module. We apply its lightweight base configuration, four transformer blocks with four attention heads, to balance efficiency and capacity.

### 2.2 Semantic Item Enrichment with BLAIR<sub>BASE</sub>

Each item is paired with a fixed textual embedding taken from the publicly released BLAIR<sub>BASE</sub> checkpoint<sup>1</sup> (125 M parameters). BLAIR is contrastively pretrained on the first 80 % of Amazon reviews (timestamp-sorted), ensuring strict temporal separation from our evaluation sets [6].

Let  $\mathbf{e}_{id} \in \mathbb{R}^d$  be the trainable ID embedding of an item and  $\mathbf{e}_{text} \in \mathbb{R}^{d_t}$  its corresponding BLAIR textual embedding. A learnable projection  $\mathbf{W}_{text} \in \mathbb{R}^{d \times d_t}$  maps textual features into the ID-embedding space, after which element-wise addition yields a unified representation

$$\mathbf{e}_{combined} = \mathbf{e}_{id} + \mathbf{W}_{text}\mathbf{e}_{text} \quad (1)$$

The fused vector is combined with positional encodings and dropout before entering the HSTU encoder, preserving dimensionality while injecting semantic context (Figure 1).

### 2.3 Text-Aware Local Negative Sampling

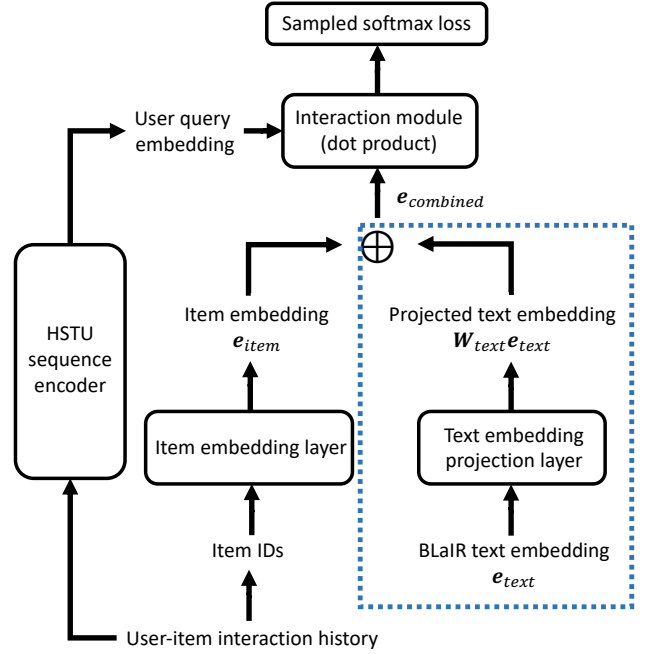
To supply semantically informed negatives, we extend the original local negative sampler by augmenting each sampled item with text embedding while leaving the sampling distribution unchanged. Specifically, for any sampled item the returned embedding is

$$\mathbf{e}_{sampled} = \frac{\mathbf{e}_{id} + \mathbf{W}_{neg}\mathbf{e}_{text}}{\|\mathbf{e}_{id} + \mathbf{W}_{neg}\mathbf{e}_{text}\|_2} \quad (2)$$

where  $\mathbf{W}_{neg} \in \mathbb{R}^{d \times d_t}$  is a projection matrix dedicated to the sampler. This modification injects semantic structure into the contrastive loss without altering the uniform ID-selection strategy or incurring additional sampling overhead.

### 2.4 Datasets

**Amazon Reviews 2023 Dataset.** We train and evaluate HSTU-BLAIr on the *Video Games*, *Office Products*, and *Musical Instruments* subsets of the Amazon Reviews 2023 dataset [6]. Each subset contains timestamped user purchase histories, product metadata (e.g., title and category), and review texts. We use the 5-core version of the dataset, filtering out users and items with fewer than five interactions, following the offline evaluation protocol in [22].



**Figure 1: HSTU-BLAIr architecture. Pre-computed BLAIR embeddings are projected and fused with trainable ID embeddings ( $\oplus$ ). The dotted box highlights the fusion module.**

**Steam Dataset.** To evaluate the generalization ability of our approach beyond traditional retail platforms, we further experiment on the Steam dataset [10], which represents a digital goods marketplace. We construct each item’s metadata by concatenating its app name, publisher, developer, genres, tags, and specifications into a single input sentence. User histories are constructed based on the user’s review timestamp, and we filter out users with fewer than five interactions and discarding items without metadata.

Steam is entirely disjoint from the Amazon corpus and was not seen by the BLAIR encoder during pretraining, thus serving as a robust cross-domain benchmark that eliminates potential overlap. Dataset statistics are summarized in Table 1.

	Items	Users	Interactions
<b>Video Games</b>	25,612	94,762	814,585
<b>Office Products</b>	77,551	223,308	1,800,877
<b>Musical Instruments</b>	24,587	57,439	511,835
<b>Steam</b>	11,808	139,259	1,538,285

**Table 1: Statistics of the three Amazon Reviews 2023 subsets and the Steam dataset used in our sequential recommendation experiments, including the number of users, items, and interactions.**

### 2.5 Tasks

We evaluate our model on next-item prediction in a sequential recommendation setting. Following Zhai et al. [22], user interactions

<sup>1</sup><https://huggingface.co/hyp1231/blair-roberta-base>



**Table 2: Evaluation metrics on the Video Games, Office Products, and Musical Instruments subsets of the Amazon Reviews 2023 dataset, as well as the Steam dataset. Each cell shows absolute values (top row) and percentage improvements over HSTU / SASRec (bottom row). Best values per column are bolded.**

Dataset	Model	HR@10	HR@50	HR@200	NDCG@10	NDCG@200	MRR
Video Games	SASRec	.1028	.2317	.3941	.0573	.1097	.0518
	HSTU	.1315	.2765	.4565	.0741	.1327	.0658
	HSTU-OpenAI (TE3L)	.1328	.2821	.4645	.0742	.1341	.0658
		(+1.0% / +29.2%)	(+2.0% / +21.8%)	(+1.8% / +17.9%)	(+0.1% / +29.5%)	(+1.1% / +22.2%)	(0.0% / +27.0%)
	<b>HSTU-BLaIR</b>	<b>.1353</b>	<b>.2852</b>	<b>.4684</b>	<b>.0760</b>	<b>.1361</b>	<b>.0674</b>
		(+2.9% / +31.6%)	(+3.1% / +23.1%)	(+2.6% / +18.9%)	(+2.6% / +32.6%)	(+2.6% / +24.1%)	(+2.4% / +30.1%)
Office Products	SASRec	.0281	.0668	.1331	.0153	.0335	.0143
	HSTU	.0395	.0880	.1649	.0223	.0443	.0207
	HSTU-OpenAI (TE3L)	.0477	.1050	.1940	.0269	.0526	.0247
		(+20.8% / +69.8%)	(+19.3% / +57.2%)	(+17.6% / +45.8%)	(+20.6% / +75.8%)	(+18.7% / +57.0%)	(+19.3% / +72.7%)
	<b>HSTU-BLaIR</b>	<b>.0484</b>	<b>.1068</b>	<b>.1946</b>	<b>.0271</b>	<b>.0529</b>	<b>.0248</b>
		(+22.5% / +72.2%)	(+21.4% / +59.9%)	(+18.0% / +46.2%)	(+21.5% / +77.1%)	(+19.4% / +57.9%)	(+19.8% / +73.4%)
Musical Instruments	SASRec	.0643	.1488	.2784	.0356	.0731	.0326
	HSTU	.0700	.1599	.2910	.0392	.0783	.0359
	HSTU-OpenAI (TE3L)	.0708	.1635	.3005	.0393	.0798	.0360
		(+1.1% / +10.1%)	(+2.3% / +9.9%)	(+3.3% / +7.9%)	(+0.3% / +10.4%)	(+1.9% / +9.2%)	(+0.3% / +10.4%)
	<b>HSTU-BLaIR</b>	<b>.0733</b>	<b>.1681</b>	<b>.3066</b>	<b>.0406</b>	<b>.0818</b>	<b>.0371</b>
		(+4.7% / +14.0%)	(+5.1% / +13.0%)	(+5.4% / +10.1%)	(+3.6% / +14.0%)	(+4.5% / +11.9%)	(+3.3% / +13.8%)
Steam	SASRec	.0881	.2277	.4312	.0462	.1068	.0426
	HSTU	.1038	.2575	.4704	.0544	.1195	.0492
	HSTU-OpenAI (TE3L)	.1089	.2657	.4806	.0579	.1241	.0525
		(+4.9% / +23.6%)	(+3.2% / +16.7%)	(+2.2% / +11.5%)	(+6.4% / +25.3%)	(+3.8% / +16.2%)	(+6.7% / +23.2%)
	<b>HSTU-BLaIR</b>	<b>.1100</b>	<b>.2668</b>	<b>.4812</b>	<b>.0581</b>	<b>.1241</b>	.0523
		(+6.0% / +24.9%)	(+3.6% / +17.2%)	(+2.3% / +11.6%)	(+6.8% / +25.8%)	(+3.8% / +16.2%)	(+6.3% / +22.8%)

are sorted chronologically, and a leave-one-out evaluation strategy is applied: for each user, the most recent interaction is held out for testing, while all prior interactions are used for training. We also adopt the training and evaluation protocol from Zhai et al. [22], which includes full data shuffling and multi-epoch training. All models, including our HSTU-BLaIR and the baselines, are trained for 100 epochs on the Amazon Reviews 2023 dataset as well as the Steam dataset.

## 2.6 Baselines

We compare our proposed framework against the following baselines:

- (1) **SASRec** [10], a widely used transformer-based sequential recommender that models item sequences using self-attention over ID embeddings;
- (2) **HSTU** [22], the original generative recommender, which encodes user histories using a hierarchical transformer architecture with only learned item ID embeddings and no textual input;

- (3) **HSTU-OpenAI (TE3L)**, a variant of HSTU-based GR in which item textual embeddings are derived from OpenAI’s text-embedding-3-large model [14], serving as an alternative to the BLaIR embeddings.

While OpenAI has not released architectural or training details for text-embedding-3-large, prior disclosures suggest it likely contains billions of parameters and was trained on broad corpora including Wikipedia, books, and large-scale web data [2]. In contrast, BLaIR<sub>BASE</sub> [6] is a domain-specific contrastive model with only 125M parameters, pretrained on Amazon Reviews 2023.

Notably, text-embedding-3-large returns unit-normed embeddings by default, whereas BLaIR imposes no explicit normalization constraint. We preserve the native output behavior of each encoder to reflect realistic usage and maintain their respective representational properties.

The text-embedding-3-large model supports a maximum input length of 8,179 tokens. In our datasets, only a small number of items—e.g., 2 in Video Games and 4 in Office Products—exceed this limit. For such cases, we retain only the first 8,179 tokens. In contrast, BLaIR was pretrained with a much shorter context



window of 64 tokens. However, following the official usage example provided on its Hugging Face model page,<sup>2</sup> we tokenize item texts using a maximum length of 512 tokens. While this exceeds the model’s original pretraining window, we found it performs robustly in practice.

## 2.7 Evaluation Metrics

We assess model performance using two standard ranking metrics commonly used in recommendation systems:

**HR@K (Hit Rate at K)** measures whether the ground-truth next item appears within the top- $K$  predicted items. It reflects the model’s ability to include the correct item in its top- $K$  recommendations, regardless of position.

**NDCG@K (Normalized Discounted Cumulative Gain at K)** extends this by considering the position of the correct item in the ranked list, assigning higher scores to items that appear earlier. It is defined as  $NDCG@K = \frac{DCG@K}{IDCG@K}$ , where  $DCG@K = \sum_{i=1}^K \frac{rel_i}{\log_2(i+1)}$  and  $IDCG@K$  denotes the maximum possible DCG for an ideal ranking. Here,  $rel_i$  is a binary relevance indicator (1 if the item at position  $i$  is correct, 0 otherwise).

**MRR (Mean Reciprocal Rank)** computes the average of the reciprocal ranks of the first relevant item across all queries. It is defined as  $MRR = \frac{1}{N} \sum_{i=1}^N \frac{1}{rank_i}$ , where  $rank_i$  is the position of the first relevant item for the  $i$ -th query. MRR rewards models that rank the correct item higher and is useful for single-ground-truth settings.

## 3 Results

Table 2 reports the performance of our proposed HSTU-BLaIR method in comparison to several baselines across four benchmarks: the Video Games, Office Products, and Musical Instruments subsets of the Amazon Reviews 2023 dataset, as well as the Steam dataset. As shown in Table 1, these datasets vary in scale and sparsity, with Office Products featuring a significantly larger user and item set, resulting in a sparser interaction matrix.

Across all datasets and evaluation metrics, HSTU-BLaIR consistently outperforms the original HSTU model, demonstrating the effectiveness of incorporating semantic item information via BLaIR’s pretrained textual embeddings. Compared to HSTU-OpenAI (TE3L), which leverages OpenAI’s text-embedding-3-large, HSTU-BLaIR achieves stronger performance on the majority of metrics. Specifically, it surpasses the OpenAI variant on 22 of 24 comparisons across datasets, with one metric slightly lower (MRR on Steam) and one metric equal (NDCG@200 on Steam). These results suggest that domain-specific contrastive encoders like BLaIR can provide more effective and compute-efficient representations for recommendation than general-purpose large language models.

Improvements are especially pronounced in the sparser Office Products dataset, where HSTU-BLaIR achieves up to a 77% gain in NDCG@10 over SASRec, and up to 21.5% over HSTU. These results highlight the effectiveness and generalizability of our approach across different dataset scales and sparsity levels.

## 4 Conclusion and Future Work

This work investigates the integration of contrastively learned textual embeddings into a generative sequential recommender system. Our findings demonstrate that enriching user sequence models with semantically grounded text representations yields consistent improvements in recommendation performance across multiple benchmarks. Specifically, we combine BLaIR-generated text embeddings with HSTU’s item ID embeddings via element-wise addition, allowing the model to jointly leverage both structured (ID-based) and unstructured (textual) item information.

To explore more flexible fusion strategies, we also implemented a learnable vector gating mechanism, where each embedding dimension independently balances ID and textual contributions. Surprisingly, this approach consistently underperformed compared to simple addition, likely due to optimization instability or insufficient signal early in training. As future work, we plan to investigate context-aware fusion strategies, such as user-conditioned or attention-based gating, as well as fine-tuning the text encoder to improve alignment with downstream objectives.

## Limitations

While our approach demonstrates the effectiveness of integrating contrastive textual signals into generative sequential recommenders, it has several limitations.

First, the BLaIR text embeddings are kept frozen during training. While this design improves efficiency and stability, it restricts the model’s ability to adapt text features to the specific downstream task. Second, the fusion of item ID embeddings and textual representations is relatively shallow (e.g., element-wise summation), which may limit the capacity of the model to capture rich interactions between the two modalities. Finally, the effectiveness of BLaIR embeddings depends heavily on the availability and quality of item metadata. In domains with sparse, noisy, or inconsistent textual content, the benefits of leveraging pretrained text representations may diminish.

## References

- [1] Kegin 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] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [3] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep neural networks for youtube recommendations. In *Proceedings of the 10th ACM conference on recommender systems*. 191–198.
- [4] Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- [5] 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.
- [6] Yupeng Hou, Jiacheng Li, Zhankui He, An Yan, Xiushi Chen, and Julian McAuley. 2024. Bridging Language and Items for Retrieval and Recommendation. *arXiv preprint arXiv:2403.03952* (2024).
- [7] Yupeng Hou, Junjie Zhang, Zihan Lin, Hongyu Lu, Ruobing Xie, Julian McAuley, and Wayne Xin Zhao. 2024. Large language models are zero-shot rankers for recommender systems. In *European Conference on Information Retrieval*. Springer, 364–381.

<sup>2</sup><https://huggingface.co/hyp1231/blair-roberta-base>



- [8] Jui-Ting Huang, Ashish Sharma, Shuying Sun, Li Xia, David Zhang, Philip Pronin, Janani Padmanabhan, Giuseppe Ottaviano, and Linjun Yang. 2020. Embedding-based retrieval in facebook search. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2553–2561.
- [9] Jianchao Ji, Zelong Li, Shuyuan Xu, Wenyue Hua, Yingqiang Ge, Juntao Tan, and Yongfeng Zhang. 2024. Genrec: Large language model for generative recommendation. In *European Conference on Information Retrieval*. Springer, 494–502.
- [10] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
- [11] 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).
- [12] Jiayi Liao, Sihang Li, Zhengyi Yang, Jiancan Wu, Yancheng Yuan, Xiang Wang, and Xiangnan He. 2024. LLaRA: Large language-recommendation assistant. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1785–1795.
- [13] Hanjia Lyu, Song Jiang, Hanqing Zeng, Yinglong Xia, Qifan Wang, Si Zhang, Ren Chen, Christopher Leung, Jiajie Tang, and Jiebo Luo. 2023. LLM-Rec: Personalized recommendation via prompting large language models. *arXiv preprint arXiv:2307.15780* (2023).
- [14] OpenAI. 2024. New embedding models and API updates. <https://openai.com/index/new-embedding-models-and-api-updates/>.
- [15] 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.
- [16] Fei Sun, Jun Liu, Jian Wu, Changhua Pei, Xiao Lin, Wenwu Ou, and Peng Jiang. 2019. BERT4Rec: Sequential recommendation with bidirectional encoder representations from transformer. In *Proceedings of the 28th ACM international conference on information and knowledge management*. 1441–1450.
- [17] Shuyuan Xu, Wenyue Hua, and Yongfeng Zhang. 2024. OpenP5: An open-source platform for developing, training, and evaluating llm-based recommender systems. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 386–394.
- [18] Wujiang Xu, Qitian Wu, Zujie Liang, Jiaojiao Han, Xuying Ning, Yunxiao Shi, Wenfang Lin, and Yongfeng Zhang. 2025. SLMRec: Distilling Large Language Models into Small for Sequential Recommendation. In *International Conference on Learning Representations (ICLR 2025)*.
- [19] Tiansheng Yao, Xinyang Yi, Derek Zhiyuan Cheng, Felix Yu, Ting Chen, Aditya Menon, Lichan Hong, Ed H Chi, Steve Tjoa, Jieqi Kang, et al. 2021. Self-supervised learning for large-scale item recommendations. In *Proceedings of the 30th ACM international conference on information & knowledge management*. 4321–4330.
- [20] Xinyang Yi, Ji Yang, Lichan Hong, Derek Zhiyuan Cheng, Lukasz Heldt, Aditee Kumthekar, Zhe Zhao, Li Wei, and Ed Chi. 2019. Sampling-bias-corrected neural modeling for large corpus item recommendations. In *Proceedings of the 13th ACM conference on recommender systems*. 269–277.
- [21] Fajie Yuan, Alexandros Karatzoglou, Ioannis Arapakis, Joemon M Jose, and Xiangnan He. 2019. A simple convolutional generative network for next item recommendation. In *Proceedings of the twelfth ACM international conference on web search and data mining*. 582–590.
- [22] Jiaqi Zhai, Lucy Liao, Xing Liu, Yueming Wang, Rui Li, Xuan Cao, Leon Gao, Zhaojie Gong, Fangda Gu, Jiayuan He, Yinghai Lu, and Yu Shi. 2024. Actions speak louder than words: trillion-parameter sequential transducers for generative recommendations. In *Proceedings of the 41st International Conference on Machine Learning (ICML'24)*.
- [23] 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.