

# Behavior Tokens Speak Louder: Disentangled Explainable Recommendation with Behavior Vocabulary

Xinshun Feng<sup>1</sup>, Mingzhe Liu<sup>1\*</sup>, Yi Qiao<sup>2</sup>, Tongyu Zhu<sup>2</sup>, Leilei Sun<sup>2</sup>, Shuai Wang<sup>1</sup>

<sup>1</sup>Hangzhou International Innovation Institute, Beihang University, Hangzhou, China

<sup>2</sup>State Key Laboratory of Complex & Critical Software Environment, Beihang University, Beijing, China  
{xinshunfeng, mzliu1997, qiaoy, zhutongyu, leileisun, shuaiwang}@buaa.edu.cn

## Abstract

Recent advances in explainable recommendation have explored the integration of language models to analyze natural language rationales for user–item interactions. Despite their potential, existing methods often rely on ID-based representations that obscure semantic meaning and impose structural constraints on language models, thereby limiting their applicability in open-ended scenarios. These challenges are intensified by the complex nature of real-world interactions, where diverse user intents are entangled and collaborative signals rarely align with linguistic semantics. To overcome these limitations, we propose BEAT, a unified and transferable framework that tokenizes user and item behaviors into discrete, interpretable sequences. We construct a behavior vocabulary via a vector-quantized autoencoding process that disentangles macro-level interests and micro-level intentions from graph-based representations. We then introduce multi-level semantic supervision to bridge the gap between behavioral signals and language space. A semantic alignment regularization mechanism is designed to embed behavior tokens directly into the input space of frozen language models. Experiments on three public datasets show that BEAT improves zero-shot recommendation performance while generating coherent and informative explanations. Further analysis demonstrates that our behavior tokens capture fine-grained semantics and offer a plug-and-play interface for integrating complex behavior patterns into large language models.

**Code** — <https://github.com/fxsxjtu/BEAT>

## Introduction

With the rapid growth of online content, users are confronted with information overload, which significantly hinders their ability to identify the most relevant items. Recommendation systems have been widely adopted to help users filter out irrelevant products from vast choices (Khusro, Ali, and Ullah 2016; Zhang et al. 2019a). Various approaches have been explored to enhance recommendation accuracy by modeling users’ latent interests and preferences, leading to a satisfying user experience (Ricci, Rokach, and Shapira 2010).

Despite high predictive accuracy, most current recommendation system methods rarely reveal explanatory behav-

ioral patterns. Explainable recommendations have been proposed, aiming to accompany each recommendation with a human-interpretable explanation (Zhang, Chen et al. 2020). Early solutions (Li et al. 2017; Dong et al. 2017) learned discrete ID-based embeddings for known users and items; however, this severely restricts their applicability to unseen users, who seldom possess extensive review histories and thus present a cold-start challenge. Recent work integrates advanced language models, including Transformer (Cheng et al. 2023), GPT-2 (Yang et al. 2024), and other large language models (LLMs) (Kim et al. 2025), to leverage powerful generative and understanding capabilities.

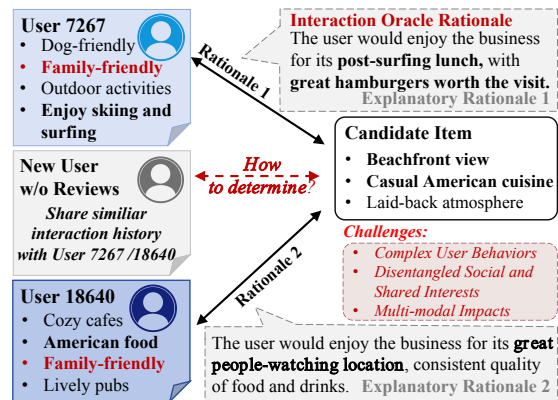


Figure 1: Illustration of challenges in real-world scenarios. Shared interests are highlighted in red, while the explanatory behavioral patterns are presented in bold.

Existing approaches, however, still face three primary challenges as illustrated in Figure 1. First, their reliance on ID-level representations (Li, Zhang, and Chen 2021) impedes generalization to cold-start users and items. While graph-based methods offer a partial remedy (Ma, Ren, and Huang 2024), they often suffer from the over-smoothing effect, which diminishes personalization (Chen et al. 2020). For instance, users with similar interaction histories (e.g., user 7267 and user 18640) may engage with the same item for entirely different reasons. Thus, effectively representing users and items while balancing their collective and individual characteristics remains a significant difficulty. Second, many frameworks incur substantial computational overhead,

\*Corresponding Authors: Mingzhe Liu  
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

as they either incorporate extensive profile texts into prompts or require fine-tuning LLMs (Zheng et al. 2024). A core challenge is failing to efficiently distill explanatory behavioral patterns behind the interaction, which is fundamentally driven by user interests and item characteristics. Third, most methods focus on either interaction history or review semantics in isolation, overlooking the valuable correlations between the two. Effectively unifying this heterogeneous information presents a persistent challenge.

Tackling these limitations can provide new insights into the design of explainable recommendation models. In this paper, we introduce **BEAT**, a lightweight **Behavior-Aware Tokenizer** that empowers a frozen LLM to reason about user-item interactions, eliminating the need for costly fine-tuning. Specifically, review-enriched users/items are mapped to discrete words drawn from a shared **Behavior Vocabulary**, which captures collective trends while preserving fine-grained individual intents. This condensed representation allows for harnessing the model’s reasoning power without the heavy computational overhead of lengthy textual profiles. We feed these behavior tokens into the LLM, enabling it to use its built-in language knowledge to link user and item features and explain why each recommendation is made. The vocabulary for these tokens is further refined using pre-trained textual correlations as supervision, thereby bridging collaborative filtering signals with natural language semantics. Finally, we introduce an auxiliary alignment objective that embeds these behavior tokens directly into the LLM’s semantic space, enabling seamless integration. To evaluate the effectiveness of the BEAT, we conduct extensive experiments on three real-world datasets in zero-shot settings, which is more common in practical applications. The results demonstrate that our approach effectively aligns linguistic and collaborative semantics through macro- and micro-level language supervision. Moreover, the proposed auxiliary alignment task further improves the LLM’s ability to interpret behavior tokens within a unified semantic space. The key contributions of this work are summarized:

- We propose to learn a `behavior` vocabulary that represents diverse users and items by capturing their mutual interests while preserving individual-level distinctions. The vocabulary is developed through a multi-level cross-modality learning framework, leveraging both collaborative and textual semantic signals into supervision.
- Our new training paradigm focuses on lightweight, explainable recommendations. It avoids computationally intensive processes and auxiliary model architectures.
- Extensive experiments validate the effectiveness of our approach. Across-LLM evaluations confirm the tokenizer’s transferability, while other analyses reveal how LLMs comprehend behavior tokens. Our behavior tokens effectively capture shared and distinct behavior patterns, transforming them into rich semantic representations.

## Related Work

### Explainable Recommendation

Early neural approaches generate review-style explanations based on explicit attribute inputs. Att2Seq (Dong et al. 2017)

encodes user-item pairs into an RNN decoder but relies on fixed attribute schemas. NRT (Li et al. 2017) jointly models rating prediction and explanation generation via shared latent factors, yet still depends on per-entity embeddings and historical ratings. PEPLER (Li, Zhang, and Chen 2023) and PETER (Li, Zhang, and Chen 2021) enhance interpretability by applying phrase-level attention to generate aspect-aware explanations without large-scale pretraining. With the rise of LLM, XRec (Ma, Ren, and Huang 2024) further advances this paradigm by injecting collaborative filter signals into the embedding space of LLM, establishing a structured framework. Recently, models such as Review-LLM (Peng et al. 2024), Reason4-Rec (Fang et al. 2025), and EXP3RT (Kim et al. 2025) rely on the retrieval or construction of user profiles. However, this approach is computationally intensive, making it impractical for review-scarce scenarios. Moreover, these methods are constrained by their reliance on collaborative filtering for attribute extraction and architectural modifications to the LLM. This thereby hinders their adaptability across evolving model backbones.

### Disentangled Recommendation Systems

Disentangled representations of users’ latent intents have proven effective in refining preferences and improving the precision of recommendations (Zhao et al. 2022). MacridVAE (Ma et al. 2019) models multiple user intents using a variational autoencoder in a structured latent space. DGCF (Wang et al. 2020a) and DisenHAN (Wang et al. 2020b) employ graph-based methods with GNNs and attention mechanisms to disentangle interests. Recent work (Chen et al. 2022) integrates contrastive learning to enhance intent separation and representation discriminability. The efficacy of existing disentangled intent models in explainable recommendation is constrained by their coarse-level preference modeling. Such limitations hinder their applicability in explainable recommendations, where interpretability and semantic alignment are crucial for decision-making. Our method encodes multi-modal information into a hierarchical structure of discrete tokens, capturing macro-level shared interests and micro-level individual preferences. This architecture provides LLMs with the nuanced representations required for reasoning in complex scenarios.

### Problem Definition

We aim to encode a user/item into  $K$  informative discrete tokens, originating from a unified `behavior` vocabulary  $\mathcal{B}$ . Although the user and item vocabularies are distinct in practice, we adopt a unified notation for simplicity. Given an LLM vocabulary  $\mathcal{T}$  consisting of  $N$  language tokens, our objective is to establish a semantic alignment between  $\mathcal{T}$  and the behavior vocabulary  $\mathcal{B}$ , allowing the LLM to perform explainable recommendations analogous to language generation based on given representations,

$$\text{Explanation}(u, i) = \text{LLM}(\text{Prompt}, \mathcal{B}_u^k, \mathcal{B}_i^j), \quad (1)$$

where  $\mathcal{B}_u^k$  and  $\mathcal{B}_i^j$  denote the top- $K$  behavior tokens selected for user  $u$  and item  $i$ , respectively, and  $k, j \in \{1, \dots, K\}$ . The explanation in the output text reveals the explanatory behavior patterns behind the user-item interactions.

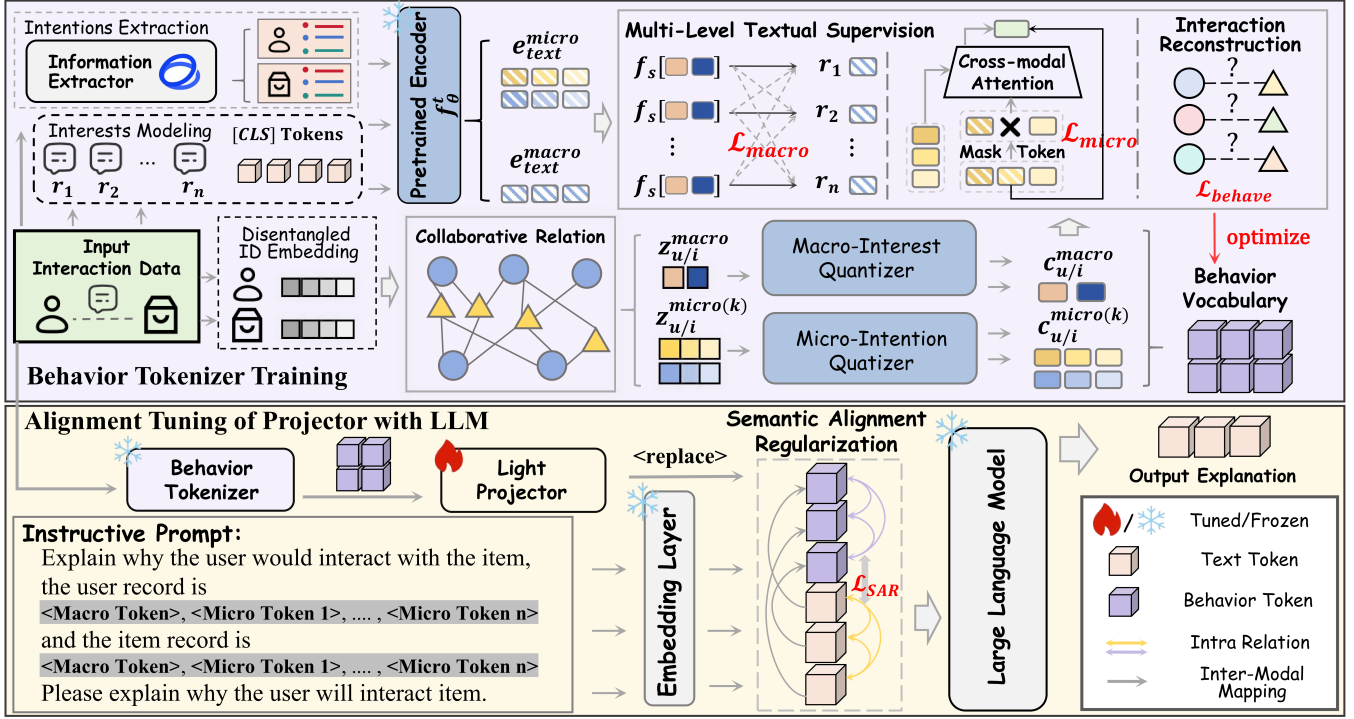


Figure 2: The BEAT framework operates in two sequential stages. It employs a collaborative language tokenizer, supervised under multi-level language semantics, to encode diverse interactions into meaningful behavior tokens. Semantic alignment regularization is then introduced to refine the model by leveraging token-level correlations from the LLM’s native vocabulary.

## Methodology

As illustrated in Figure 1, our two-stage method first constructs a disentangled user representation,  $\mathbf{z}_u$ , which captures a user’s macro interests,  $\mathbf{z}_u^{\text{macro}}$ , and micro intentions,  $\mathbf{z}_u^{\text{micro}}$ . Following a graph propagation step, we select the most suitable macro and micro behavior tokens,  $\mathbf{c}_u^{\text{macro}}$  and  $\mathbf{c}_u^{\text{micro}}$ , from the behavior vocabulary to align with the user’s representation. We then employ multi-level semantic supervision with explicit language signals  $\mathbf{e}^{\text{micro}}$  and  $\mathbf{e}^{\text{macro}}$  to enhance the precision of our user modeling. Finally, a multi-objective training strategy, which combines semantic alignment regularization with next-token prediction, is utilized to establish a semantic correlation between the learned behavior vocabularies and the LLM’s semantic space. An analogous procedure is followed for items.

### Disentangled Behavior Modeling

A key challenge in enhancing LLMs for recommendations is effectively representing users and items as tokens within the model framework (Hua et al. 2023). Existing methods assign unique IDs to each user and item, resulting in redundant vocabularies and the out-of-vocabulary (OOV) problem when adapting to new entities (Zheng et al. 2024). However, while users exhibit diverse preferences, they also share collective patterns shaped by trends or social factors. For instance, although preferences differ across product categories, common attributes such as affordability and brand reputation consistently remain priorities. To capture these collective and distinctive behaviors, we represent each user

as a sentence composed of multiple words, where each word reflects an aspect of the user’s preferences. The core idea is that similar users and items tend to share overlapping semantic tokens, allowing each unique token combination to correspond to a specific latent semantic pattern.

Formally, we aim to learn a factorized representation for each user  $u$ , consisting of macro interests and micro intentions:  $\mathbf{z}_u = [\mathbf{z}_u^{\text{macro}} \parallel \mathbf{z}_u^{\text{micro}(1)} \parallel \dots \parallel \mathbf{z}_u^{\text{micro}(N)}] \in \mathbb{R}^{d'}$ , where  $d' = (N + 1)d$ . Here,  $\mathbf{z}_u^{\text{macro}}$  captures user-specific macro-level preferences, while  $\mathbf{z}_u^{\text{micro}(i)}$  represents the shared micro-level intention across users. The operator  $\parallel$  denotes the concatenation of these components. To incorporate collaborative relations among users, we adopt a lightweight graph convolutional approach (He et al. 2020) that propagates information through the user–item interaction graph:

$$\mathbf{z}_u^{(l+1)} = \sum_{i \in \mathcal{N}_u} \frac{1}{\sqrt{|\mathcal{N}_u|} \sqrt{|\mathcal{N}_i|}} \mathbf{z}_i^{(l)}. \quad (2)$$

After deriving collaboration-enhanced representations, we apply an averaging operation across graph propagation layers to preserve multi-order collaborative information:

$$\mathbf{z}_u = \sum_{k=0}^K \frac{1}{K+1} \mathbf{z}_u^{(k)}. \quad (3)$$

To capture the collectiveness and individuality of user behaviors, we train a Vector Quantized Variational Autoencoder (VQ-VAE) (Van Den Oord, Vinyals et al. 2017) to

discretize continuous input vectors by mapping them to a set of codewords. For each user representation  $\mathbf{z}_u$ , we first decompose it into a set of interest- and intention-level vectors:  $\{\mathbf{z}_u^{\text{macro}}, \mathbf{z}_u^{\text{micro}(i)}\}_{i=1}^N$ . We then construct two separate codebooks of fixed size for macro-level interests and micro-level intentions. Together, these codebooks constitute the `behavior vocabulary` in our model, representing users in a discrete, interpretable token space, with codewords denoting `words`. Taking the  $i$ -th micro-intention representation  $\mathbf{z}_u^{\text{micro}(i)}$  as an example, we first project it into the vocabulary space  $\mathbb{R}^{d'}$  via a linear transformation. We then compute its distance to all words in the vocabulary  $\mathcal{C} = \{\mathbf{c}_k\}_{k=1}^K$ , and assign it to the nearest:

$$\mathbf{c}_j = \arg \min_k \left\| \mathbf{z}_u^{\text{micro}(i)} - \mathbf{c}_k \right\|_2^2, \quad (4)$$

where  $\mathbf{c}_j$  is the selected word assigned to the  $i$ -th input representation  $\mathbf{z}_u^{\text{micro}(i)}$ , and  $\mathbf{c}_k$  denotes the  $k$ -th word in the corresponding vocabulary. We derive a quantized representation for each user as  $\mathbf{q}_u = [\mathbf{c}_u^{\text{macro}} \parallel \mathbf{c}_u^{\text{micro}(1)} \parallel \dots \parallel \mathbf{c}_u^{\text{micro}(N)}] \in \mathbb{R}^{M'}$ , where  $M' = (N + 1)M$  and  $M$  denotes the codebook embedding dimension. The overall objective function for predicting user-item interaction using the quantized representations  $\mathbf{q}_u$  and  $\mathbf{q}_i$  is defined as follows:

$$\begin{aligned} \mathcal{L}_{\text{RECON}} &= \sqrt{\|\mathbf{I}_{ui} - \mathbf{q}_u^\top \mathbf{q}_i\|_2^2}, \\ \mathcal{L}_{\text{VQ-VAE}} &= \|\text{sg}[\mathbf{q}] - \mathbf{c}\|_2^2 + \eta \|\mathbf{q} - \text{sg}[\mathbf{c}]\|_2^2, \\ \mathcal{L}_{\text{behave}} &= \mathcal{L}_{\text{RECON}} + \mathcal{L}_{\text{VQ-VAE}}, \end{aligned} \quad (5)$$

where  $\mathbf{I}_{ui} \in \{0, 1\}$  is a binary indicator of interaction between user  $u$  and item  $i$ ,  $\text{sg}[\cdot]$  denotes the stop-gradient operator, and  $\eta$  is a balancing hyperparameter, setting to 0.5. The overall loss function comprises two components: a reconstruction loss  $\mathcal{L}_{\text{RECON}}$  that encourages accurate prediction of user-item interactions, and a symmetric vector quantization loss  $\mathcal{L}_{\text{VQ}}$  that minimizes the distance between the encoded representations and their assigned codebook embeddings.

### Multi-level Textual Semantic Supervision

While Vector Quantization (VQ) excels at capturing common user behaviors (Rajput et al. 2023; Lin et al. 2025), its effectiveness is often hampered by codebooks constructed from unimodal data. Methods relying solely on collaborative signals face challenges with behavioral ambiguity, whereas those using textual data are constrained by its inherent sparsity. These issues are compounded by data sparsity, as users in real-world systems typically interact with few items and provide limited explicit feedback (Zheng et al. 2019). To overcome these issues, we propose a multi-level textual supervision approach to construct a unified codebook that integrates rich textual semantics with collaborative signals, thereby bridging the modality gap.

**Macro Semantic Supervision** aims to guide the learning of users’ macro-level interests using semantic signals derived from reviews. A user’s macro interest is expected to be unique and distinguishable, as it reflects the fundamental and stable aspects of the user’s overall preferences. Such

macro-level interests may involve high-level goals, such as purchasing for broad categories of interest.

Given a review  $r_{ui}$  associated with user  $u$  and item  $i$ , we utilize a frozen, pretrained text encoder  $f_\theta^t$  to extract its semantic representation. Considering that real-world reviews often vary in length and may contain irrelevant information, we avoid naively averaging all token embeddings. Instead, we adopt the `[CLS]` token as a global contextual feature that summarizes the review. Specifically, we sample a batch of user-item pairs  $\mathcal{B}$  along with their corresponding reviews, then apply the text encoder  $f_\theta^t$  to obtain the semantic embeddings  $\mathbf{e}^{[\text{CLS}]}$ , which serve as supervision signals for learning the macro behaviors. Since a review reflects the interaction rationale between a user and an item, we introduce a fusion function  $f_s$  to approximate this process as  $\mathbf{c}^{\text{macro}} = f_s(\mathbf{c}_u^{\text{macro}}, \mathbf{c}_i^{\text{macro}})$ , where  $\mathbf{c}_u^{\text{macro}}$  and  $\mathbf{c}_i^{\text{macro}}$  are the quantized representations. The reviewed user-item pairs are treated as positive instances, while all other pairs in the batch are considered negatives. Then the InfoNCE (Oord, Li, and Vinyals 2018) loss is employed for bridging the semantic gap between reviews and behavior tokens, encouraging alignment between the review representation and the corresponding behavior token representations while distinguishing them from others,  $s(\cdot, \cdot)$  is cosine similarity:

$$\mathcal{L}_{\text{macro}} = - \sum_{i \in \mathcal{B}} \log \frac{\exp(s(\mathbf{e}_i^{\text{CLS}}, \mathbf{c}_i^{\text{macro}}))}{\sum_{j \in \mathcal{B}} \exp(s(\mathbf{e}_i^{\text{CLS}}, \mathbf{c}_j^{\text{macro}}))}. \quad (6)$$

**Micro Semantic Supervision** uses semantic cues from behavioral data to capture fine-grained user intentions. This approach complements macro-level supervision, which identifies broad preferences that are often unique to an individual. In contrast, micro-intentions are frequently shared across many users—for instance, a common preference for features like “durability” or “ease of use,” regardless of the product category. Modeling users’ micro-intentions toward specific items is achieved by extracting fine-grained intent expressions from their historical reviews. Given a user’s historical review list, denoted as  $\mathbf{R}_u = \{r_1, r_2, \dots, r_n\}$ , we prompt an LLM to perform an information extraction task over the review list. This process retrieves a set of interpretable micro-intentions:

$$\mathbf{e}_1^{\text{micro}}, \dots, \mathbf{e}_n^{\text{micro}} = f_\theta(\text{LLM}(\mathbb{P}_e, \mathbf{R}_u)), \quad (7)$$

where  $\mathbb{P}_e$  denotes the designed extraction prompt, and  $\mathbf{e}_i^{\text{micro}}$  is the  $i$ -th extracted micro-intention, encoded using a pretrained text encoder  $f_\theta^t$  by preserving `[CLS]` tokens.

Despite we obtain both the textual micro-intention embeddings  $\{\mathbf{e}_i^{\text{micro}}\}$  and the collaborative behavior token representations  $\{\mathbf{c}_u^{\text{micro}(i)}\}$ , establishing a direct one-to-one correspondence between them is intractable due to their latent and unordered characteristics. To address this, we propose a fine-grained masked reconstruction strategy that leverages the semantic relationships among micro-intentions for supervision. For every sequence of  $n$  textual micro-intention representations, we randomly sample  $t$  positions to be masked, resulting in a corrupted sequence  $\mathcal{S}^{\text{msk}} = \{\mathbf{e}_1^{\text{micro}}, \mathbf{m}_1, \mathbf{e}_3^{\text{micro}}, \dots, \mathbf{e}_n^{\text{micro}}\}$ , where  $\mathbf{m}_i$  denotes a mask token at the  $i$ -th position. We then employ a cross-attention

module  $f_{\theta}^c$  to incorporate the behavior token representations  $\mathbf{c}^{\text{micro}}$  and infer the embeddings of the masked micro-intentions based on the contextual cues from the unmasked ones. The objective is to reconstruct the original semantic embedding of each masked intention:

$$\mathcal{L}_{\text{micro}} = -\mathbb{E}_{(c_{\text{micro}}, S^{\text{msk}}) \sim \mathcal{B}} \phi(e_{\text{msk}}, f_{\theta}^c(c_{\text{micro}}, S^{\text{msk}})), \quad (8)$$

where  $e_{\text{msk}}$  denotes the ground-truth embedding of the masked micro-intentions, and  $\phi$  distance measuring function between distributions, which is the L2 distance in practice.

**Overall Training Objective** The behavior vocabulary is optimized by integrating CF signals with textual semantic supervision, and is balanced by  $\alpha$  and  $\beta$ :

$$\mathcal{L}_{\text{tokenizer}} = \alpha \cdot \mathcal{L}_{\text{macro}} + \beta \cdot \mathcal{L}_{\text{micro}} + \mathcal{L}_{\text{behave}}. \quad (9)$$

### Behavior Tokens Comprehension

Multi-level semantic supervision establishes semantic correlations among the learned behavior tokens; however, a distributional gap persists between the LLM and the behavior tokens. We first introduce a natural language task prompt shown in Figure 2 that contextualizes the behavior tokens. This prompt design allows the LLMs to interpret the tokens in a semantically meaningful context, leveraging the behavior patterns and their intrinsic language understanding (Liu et al. 2023). Previous methods typically tune LLMs to understand extensive tokens or explicitly modify LLM structures (Zheng et al. 2024; Ma, Ren, and Huang 2024), which limits their generalization in practical scenarios. We address this by introducing a projection module that maps behavior tokens directly into the LLM’s semantic space with multiple training objectives. The projected token embeddings are then used to replace the placeholders (Tokens), enabling the LLM to process the behavior tokens directly.

Since LLMs already possess rich semantic correlations among their native token spaces from massive pre-training, we propose to transfer these inherent semantic relationships to the newly introduced behavior tokens. For example, if a user is a history enthusiast, their behavior tokens and the behavior tokens of historical books should be associated with the “history” aspect. Furthermore, the relationship between these specific tokens should reflect the textual semantic relationship between phrases like ‘love’ and ‘historical books’, thus catering to the LLM’s intrinsic understanding. We then propose **Semantic Alignment Regularization** as an auxiliary training objective to align the projected behavior tokens with the LLM’s native semantic space.

For every word representation  $\epsilon_{w_i}$  encoded by the LLM from the explanation text, we aim to establish a semantic correlation between text attributes and behavior tokens. Since we have applied multi-level textual supervision, we map each text token to its corresponding behavior token:

$$\hat{c}_i = \arg \min_{c \in \mathcal{C}} \|\epsilon_{w_i} - \mathbf{c}\|_2, \quad (10)$$

where  $\mathcal{C}$  denotes the set of projected behavior token embeddings. We then introduce the following loss function:

$$\mathcal{L}_{\text{SAR}} = \sum_{(w_i, w_j) \in \mathcal{T}} (s(\hat{c}_i, \hat{c}_j) - s(\epsilon_{w_i}, \epsilon_{w_j}))^2, \quad (11)$$

where  $\mathcal{T}$  is a sampled set of word pairs from the explanation text, and  $s(\cdot, \cdot)$  denotes cosine similarity. This loss aligns the textual and behavioral semantic spaces by encouraging behavior tokens to reflect the semantic relationships found in the text. To enhance the coherence of the generated explanations, we train the model by minimizing the negative log-likelihood (NLL) of the ground-truth explanation sequence:

$$\mathcal{L}_{\text{NLL}} = -\frac{1}{N} \sum_{i=1}^N \sum_{c=1}^{C_i} y_{ic} \cdot \log(\hat{y}_{ic}), \quad (12)$$

where  $N$  denotes the number of explanations,  $C_i$  is token number in the  $i$ -th explanation, and  $y_{ic}$  and  $\hat{y}_{ic}$  denote the ground truth and predicted probabilities for the  $c$ -th token.

## Experiments

In this section, we conduct comprehensive experiments to evaluate the effectiveness of **BEAT**. These experiments are designed to address the following research questions:

- **RQ1:** How effective is BEAT at generating explainable recommendations in a zero-shot setting, and what are the respective contributions of its core components?
- **RQ2:** How robust and transferable is the proposed behavior tokenizer across different LLM backbones?
- **RQ3:** How effectively does the behavior tokenizer capture user behavior patterns, particularly in cold-start scenarios involving unseen users and items?
- **RQ4:** How does the model interpret the learned behavior tokens, and what semantic meaning do they convey?

### Experiment Settings

The complete experimental settings are detailed in the Appendix. For all subsequent experiments, we define the ‘zero-shot’ setting as users or items that have interaction data but lack any explicit textual reviews.

### Performance Comparison (RQ1)

Experiments conducted on three real-world datasets demonstrate the effectiveness of our proposed BEAT. We choose three different metrics to validate the performances, as BERTScore (Zhang et al. 2019b), BARTScore (Yuan, Neubig, and Liu 2021), and BLEU (Papineni et al. 2002). The results, presented in Table 1, show that BEAT achieves state-of-the-art or highly competitive performance across all three scenarios. Traditional ID-based methods, such as Attn2Seq or NRT, appear to struggle in zero-shot settings, indicating the limited expressiveness of ID representations. While LLM-based methods have usually achieved promising results, demonstrating the effectiveness of LLMs’ generation ability. Recently popular cross-modality alignment methods also show competitive results, showing the necessity of semantic alignment; our model effectively integrates these strengths. A detailed ablation study isolates the contributions of our model’s key components: the micro tokens, macro tokens, and the second-stage alignment. Removing the micro tokens led to a significant performance drop, confirming that fine-grained representations for users and items

Model type	Model	Amazon			Google			Yelp		
		BLEU↑	BART↑	BERT↑	BLEU↑	BART↑	BERT↑	BLEU↑	BART↑	BERT↑
ID Based	Attn2seq	0.2843	-4.9738	-0.1207	0.3846	-4.2112	0.3101	0.3834	-4.6040	0.1841
	NRT	0.2366	-4.3370	0.3300	0.3045	-4.1964	0.2965	0.3112	-5.0841	0.2652
	PETER	0.3682	-4.2300	0.1488	0.3533	<b>-3.6307</b>	0.3283	0.3329	-4.6469	0.1675
LLM Based	PLETER	0.3120	-4.0630	0.2816	0.3850	<u>-4.1243</u>	0.3631	0.3470	-4.5996	0.3252
	XRec	0.2999	-4.3210	<u>0.3598</u>	0.2785	-4.5901	0.2628	0.3565	-4.6226	0.2871
	XRec-Profile	0.3854	-4.1749	0.3259	0.3289	-4.3491	0.3251	0.2925	-4.5955	0.2579
	GraphGPT	0.3596	-4.0380	0.3494	0.3702	-4.4328	0.3650	0.3399	-4.5275	0.2958
	TEA-GLM	<u>0.3971</u>	-4.1348	0.3406	0.3689	-4.3574	0.3521	<u>0.3844</u>	-4.5452	0.3067
	Time-LLM	0.3936	-4.0780	0.3470	0.2727	-4.2873	0.2983	0.3368	-4.5933	0.3135
Our Proposed	<b>BEAT w/o Micro</b>	0.3781	-4.0600	0.3447	0.3027	-4.2686	0.3397	0.3136	-4.5545	0.2793
	<b>BEAT w/o Macro</b>	0.3795	-4.1382	0.3415	0.3834	-4.3384	0.3580	<b>0.3852</b>	-4.5691	0.3272
	<b>BEAT w/o SAR</b>	0.3720	-4.0945	0.3531	<b>0.3922</b>	-4.3939	0.3776	0.2967	<b>-4.5019</b>	0.3124
	<b>BEAT</b>	<b>0.4195</b>	<b>-3.9929</b>	<b>0.3821</b>	<u>0.3866</u>	-4.3027	<b>0.3781</b>	0.3771	<u>-4.5442</u>	<b>0.3302</b>

Table 1: Zero-shot accuracy on Amazon, Google, and Yelp datasets. Each dataset’s evaluation metrics include BLEU, BARTScore, and BERTScore. (**bold** highlights the best result across all methods, while underline highlights the second)

are essential. The impact of removing macro tokens, however, was dataset-dependent: it degraded performance on Google but improved it on Amazon and Yelp. We conjecture that a high-level summary token may occasionally distract the LLM from detailed user preferences. Disabling the second-stage alignment caused a substantial performance decrease, particularly on the complex Yelp dataset. This result underscores that for sophisticated domains, incorporating semantic relationships through alignment is necessary to grasp the underlying recommendation rationale fully.

### Robustness across LLM Backbones (RQ2)

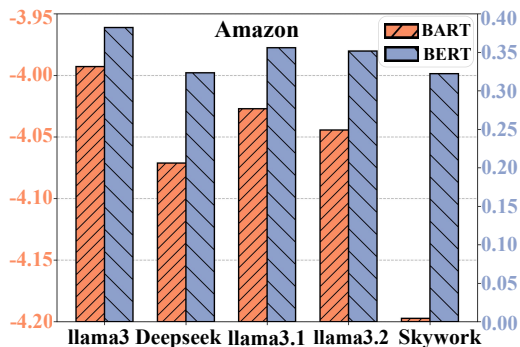


Figure 3: Robustness of BEAT Across LLM Backbones

To assess the robustness and transferability of BEAT, we evaluated its performance with several open-source LLM backbones: DeepSeek-8B (Liu et al. 2024), LLaMA3.1-8B, LLaMA3.2-3B (Grattafiori et al. 2024), and Skywork-8B (Wei et al. 2023). The evaluation was conducted on the Amazon dataset, with results reported in Figure 3 using BARTScore and BERTScore-F1 as metrics. The LLaMA models (LLaMA3.1-8B and LLaMA3.2-3B) achieved performance comparable to LLaMA3-8B, suggesting that models with similar architectures share compatible semantic

spaces. Notably, the much smaller LLaMA3.2-3B delivered strong results, demonstrating our method’s scalability and effectiveness in resource-constrained settings. In contrast, DeepSeek-8B and Skywork-8B showed lower fluency and fidelity, which we attribute to their known tendency for hallucination. While these backbones are less competitive, they still surpassed most baselines. Overall, these findings confirm the robustness of BEAT across various backbones and its straightforward adaptability to suit practical needs.

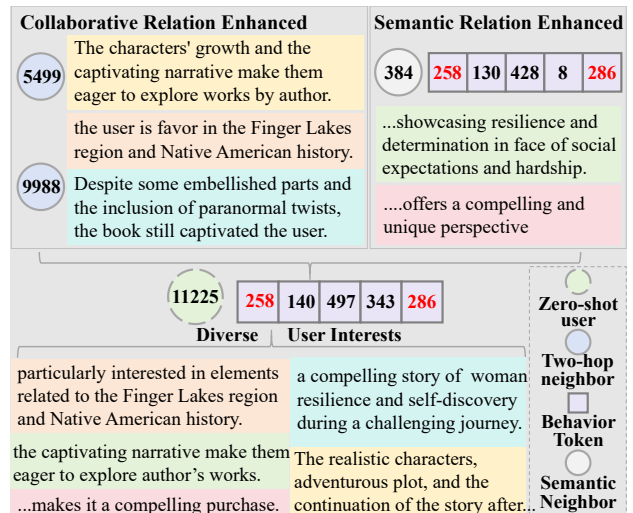


Figure 4: A Token-Level Visualization for Zero-Shot User.

### Modeling Behavior for Cold-Start Users (RQ3)

Figure 4 presents a case study on a “cold-start” user to demonstrate how our model assembles a token representation for them without any review history. The model first constructs an initial profile by borrowing from semantic neighbors; for instance, it uses tokens from User 384 to in-

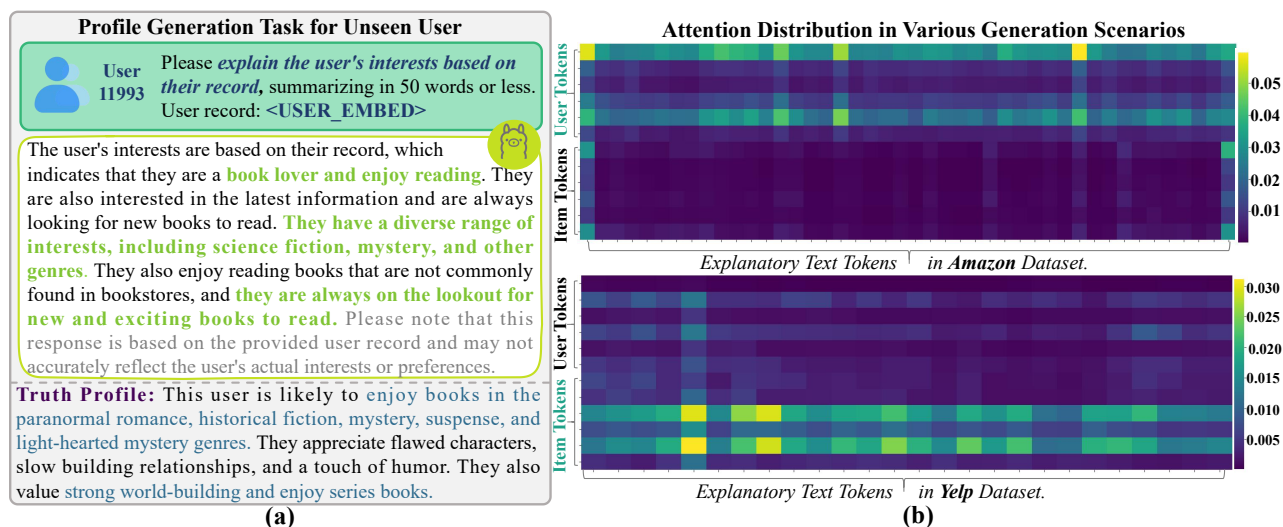


Figure 5: Illustration of the LLM’s Comprehension of Behavior Tokens. (a)The LLM’s generated response when conditioned on a batch of behavior tokens and an instruction prompt. (b)Visualization of the attention distribution over input behavior tokens.

fer an interest in themes like ”resilience and self-discovery” (Token 286). This profile is then refined with collaborative signals from users with overlapping interests, such as User 5499 and User 9988, who both share an interest in the Finger Lakes region and Native American history. As illustrated by the color-coding in the figure, this process weaves together semantic and collaborative information to bridge low-level symbolic tokens with high-level interests. Ultimately, the model produces a coherent and interpretable token set for a completely new user, showcasing its robustness in overcoming the cold-start limitations of traditional methods.

### Interpretability of Behavior Tokens (RQ4)

**User Profile Generation** Despite the effectiveness of the proposed behavior tokens, it remains unclear whether the LLM is simply memorizing input patterns or understanding their semantic correlations. To investigate this, we prompted the LLM to generate a user profile directly, without any task-specific training. Figure 5(a) presents the result for a zero-shot user, whose reviews are ablated from all training stages. The results demonstrate that the LLM successfully inferred the user’s preferences and interests. Notably, it also correctly deduced the book-purchasing context without explicit cues, indicating that the scenario’s semantics were properly encoded. The generated profile offered detailed explanations for the user’s interests and largely matched the ground-truth. Although some hallucinations (Yao et al. 2023) were observed, we posit that they can be mitigated through targeted fine-tuning. Given that the LLM was not trained on this task, these findings suggest that it has established a semantic correlation with the behavior tokens, rather than memorizing.

**Behavior Token Functioning** Although our experiments demonstrate that the LLM can holistically understand behavior tokens, the underlying working mechanism remains obscure. This prompted us to investigate the functional effectiveness of these individual behavior tokens. Figure 5(b)

visualizes the attention matrices for two distinct scenarios: the Amazon dataset (top) and the Yelp dataset (bottom). In each matrix, the color intensity reflects the attention score between the input tokens on the left (partitioned into user- and item-specific sets) and the generated words on the bottom. Specifically, the model’s attentional focus shifts from users on the Amazon dataset to items on the Yelp dataset. We conjecture that this is because the Amazon dataset covers a narrow domain: book reviews. In such a homogeneous context, the rationale for a user-item interaction is likely to depend on the user’s specific interests. In contrast, the Yelp dataset encompasses a diverse range of items, the unique attributes of the item itself become more critical for generating a meaningful rationale. This finding demonstrates the model’s ability to dynamically shift its focus between users and items to suit the dataset’s context. Furthermore, different tokens make distinct contributions during the generation process, revealing the complexity of user behavior. Our model leverages these varied signals to craft explanations that reflect the nuanced interaction between a user’s diverse preferences and an item’s specific attributes.

### Conclusion

In this work, we introduced a unified behavior tokenizer that translates user-item interactions into an LLM-comprehensible vocabulary. By leveraging a two-level textual supervision mechanism and incorporating semantic relational knowledge, our method enables LLMs to understand these behavioral tokens without extra fine-tuning, achieving state-of-the-art or highly competitive performance on benchmark datasets. Further analysis reveals how our approach captures diverse user interests, enhances the LLM’s capacity for behavioral reasoning, and how the LLM interprets given behavior tokens. Future work will focus on extending this approach to diverse recommendation scenarios and evaluating its cross-domain transferability.

## Acknowledgments

This work was supported by National Key R&D Program of China (2024YFB4506004), National Natural Science Foundation of China (62272023), and partly supported by Science and Technology Project of Beijing Municipal Commission of Transport (2025-KJC-03-003).

## References

- Chen, D.; Lin, Y.; Li, W.; Li, P.; Zhou, J.; and Sun, X. 2020. Measuring and relieving the over-smoothing problem for graph neural networks from the topological view. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 3438–3445.
- Chen, Y.; Liu, Z.; Li, J.; McAuley, J.; and Xiong, C. 2022. Intent contrastive learning for sequential recommendation. In *Proceedings of the ACM web conference 2022*, 2172–2182.
- Cheng, H.; Wang, S.; Lu, W.; Zhang, W.; Zhou, M.; Lu, K.; and Liao, H. 2023. Explainable recommendation with personalized review retrieval and aspect learning. *arXiv preprint arXiv:2306.12657*.
- Dong, L.; Huang, S.; Wei, F.; Lapata, M.; Zhou, M.; and Xu, K. 2017. Learning to generate product reviews from attributes. In *15th EACL 2017 Software Demonstrations*, 623–632. Association for Computational Linguistics.
- Fang, Y.; Wang, W.; Zhang, Y.; Zhu, F.; Wang, Q.; Feng, F.; and He, X. 2025. Reason4Rec: Large Language Models for Recommendation with Deliberative User Preference Alignment. *arXiv preprint arXiv:2502.02061*.
- Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- He, X.; Deng, K.; Wang, X.; Li, Y.; Zhang, Y.; and Wang, M. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, 639–648.
- Hua, W.; Xu, S.; Ge, Y.; and Zhang, Y. 2023. How to index item ids for recommendation foundation models. In *Proceedings of the Annual International ACM SIGIR Conference on Research and Development in Information Retrieval in the Asia Pacific Region*, 195–204.
- Khusro, S.; Ali, Z.; and Ullah, I. 2016. Recommender systems: issues, challenges, and research opportunities. In *Information science and applications (ICISA) 2016*, 1179–1189. Springer.
- Kim, J.; Kim, H.; Cho, H.; Kang, S.; Chang, B.; Yeo, J.; and Lee, D. 2025. Review-driven personalized preference reasoning with large language models for recommendation. CoRR, abs/2408.06276, 2024. doi: 10.48550. *arXiv preprint ARXIV.2408.06276*.
- Li, L.; Zhang, Y.; and Chen, L. 2021. Personalized transformer for explainable recommendation. *arXiv preprint arXiv:2105.11601*.
- Li, L.; Zhang, Y.; and Chen, L. 2023. Personalized prompt learning for explainable recommendation. *ACM Transactions on Information Systems*, 41(4): 1–26.
- Li, P.; Wang, Z.; Ren, Z.; Bing, L.; and Lam, W. 2017. Neural rating regression with abstractive tips generation for recommendation. In *Proceedings of the 40th International ACM SIGIR conference on Research and Development in Information Retrieval*, 345–354.
- Lin, G.; Hua, Z.; Feng, T.; Yang, S.; Long, B.; and You, J. 2025. Unified semantic and ID representation learning for deep recommenders. *arXiv preprint arXiv:2502.16474*.
- Liu, A.; Feng, B.; Xue, B.; Wang, B.; Wu, B.; Lu, C.; Zhao, C.; Deng, C.; Zhang, C.; Ruan, C.; et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Liu, H.; Li, C.; Wu, Q.; and Lee, Y. J. 2023. Visual instruction tuning. *Advances in neural information processing systems*, 36: 34892–34916.
- Ma, J.; Cui, P.; Kuang, K.; Wang, X.; and Zhu, W. 2019. Disentangled graph convolutional networks. In *International conference on machine learning*, 4212–4221. PMLR.
- Ma, Q.; Ren, X.; and Huang, C. 2024. XRec: Large Language Models for Explainable Recommendation. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 391–402.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Papineni, K.; Roukos, S.; Ward, T.; and Zhu, W.-J. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 311–318.
- Peng, Q.; Liu, H.; Xu, H.; Yang, Q.; Shao, M.; and Wang, W. 2024. Review-LLM: Harnessing Large Language Models for Personalized Review Generation. *arXiv preprint arXiv:2407.07487*.
- Rajput, S.; Mehta, N.; Singh, A.; Hulikal Keshavan, R.; Vu, T.; Heldt, L.; Hong, L.; Tay, Y.; Tran, V.; Samost, J.; et al. 2023. Recommender systems with generative retrieval. *Advances in Neural Information Processing Systems*, 36: 10299–10315.
- Ricci, F.; Rokach, L.; and Shapira, B. 2010. Introduction to recommender systems handbook. In *Recommender systems handbook*, 1–35. Springer.
- Van Den Oord, A.; Vinyals, O.; et al. 2017. Neural discrete representation learning. *Advances in neural information processing systems*, 30.
- Wang, X.; Jin, H.; Zhang, A.; He, X.; Xu, T.; and Chua, T.-S. 2020a. Disentangled graph collaborative filtering. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, 1001–1010.
- Wang, Y.; Tang, S.; Lei, Y.; Song, W.; Wang, S.; and Zhang, M. 2020b. Disenhan: Disentangled heterogeneous graph attention network for recommendation. In *Proceedings of the 29th ACM international conference on information & knowledge management*, 1605–1614.

Wei, T.; Zhao, L.; Zhang, L.; Zhu, B.; Wang, L.; Yang, H.; Li, B.; Cheng, C.; Lü, W.; Hu, R.; et al. 2023. Skywork: A more open bilingual foundation model. *arXiv preprint arXiv:2310.19341*.

Yang, M.; Zhu, M.; Wang, Y.; Chen, L.; Zhao, Y.; Wang, X.; Han, B.; Zheng, X.; and Yin, J. 2024. Fine-tuning large language model based explainable recommendation with explainable quality reward. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 9250–9259.

Yao, J.-Y.; Ning, K.-P.; Liu, Z.-H.; Ning, M.-N.; Liu, Y.-Y.; and Yuan, L. 2023. Llm lies: Hallucinations are not bugs, but features as adversarial examples. *arXiv preprint arXiv:2310.01469*.

Yuan, W.; Neubig, G.; and Liu, P. 2021. Bartscore: Evaluating generated text as text generation. *Advances in neural information processing systems*, 34: 27263–27277.

Zhang, S.; Yao, L.; Sun, A.; and Tay, Y. 2019a. Deep learning based recommender system: A survey and new perspectives. *ACM computing surveys (CSUR)*, 52(1): 1–38.

Zhang, T.; Kishore, V.; Wu, F.; Weinberger, K. Q.; and Artzi, Y. 2019b. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.

Zhang, Y.; Chen, X.; et al. 2020. Explainable recommendation: A survey and new perspectives. *Foundations and Trends® in Information Retrieval*, 14(1): 1–101.

Zhao, S.; Wei, W.; Zou, D.; and Mao, X. 2022. Multi-view intent disentangle graph networks for bundle recommendation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 4379–4387.

Zheng, B.; Hou, Y.; Lu, H.; Chen, Y.; Zhao, W. X.; Chen, M.; and Wen, J.-R. 2024. Adapting large language models by integrating collaborative semantics for recommendation. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 1435–1448. IEEE.

Zheng, L.; Li, C.; Lu, C.-T.; Zhang, J.; and Yu, P. S. 2019. Deep distribution network: Addressing the data sparsity issue for top-n recommendation. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1081–1084.