

## Towards Explainable Temporal User Profiling with LLMs

Sabouri, Milad; Mansoury, Masoud; Lin, Kun; Mobasher, Bamshad

**DOI**

[10.1145/3708319.3733655](https://doi.org/10.1145/3708319.3733655)

**Publication date**

2025

**Document Version**

Final published version

**Published in**

UMAP Adjunct '25: Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization

**Citation (APA)**

Sabouri, M., Mansoury, M., Lin, K., & Mobasher, B. (2025). Towards Explainable Temporal User Profiling with LLMs. In *UMAP Adjunct '25: Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization* (pp. 219-227). ACM. <https://doi.org/10.1145/3708319.3733655>

**Important note**

To cite this publication, please use the final published version (if applicable). Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

**Green Open Access added to [TU Delft Institutional Repository](#)  
as part of the Taverne amendment.**

More information about this copyright law amendment  
can be found at <https://www.openaccess.nl>.

Otherwise as indicated in the copyright section:  
the publisher is the copyright holder of this work and the  
author uses the Dutch legislation to make this work public.



# Towards Explainable Temporal User Profiling with LLMs

Milad Sabouri  
DePaul University  
Chicago, IL, USA  
msabouri@depaul.edu

Kun Lin  
DePaul University  
Chicago, IL, USA  
klin13@depaul.edu

Masoud Mansoury  
Delft University of Technology  
Delft, Netherlands  
m.mansoury@tudelft.nl

Bamshad Mobasher  
DePaul University  
Chicago, IL, USA  
mobasher@cs.depaul.edu

## Abstract

Accurately modeling user preferences is vital not only for improving recommendation performance but also for enhancing transparency in recommender systems. Conventional user-profiling methods—such as averaging item embeddings—often overlook the evolving, nuanced nature of user interests, particularly the interplay between short-term and long-term preferences. In this work, we leverage large language models (LLMs) to generate natural language summaries of users’ interaction histories, distinguishing recent behaviors from more persistent tendencies. Our framework not only models temporal user preferences but also produces natural language profiles that can be used to explain recommendations in an interpretable manner. These textual profiles are encoded via a pre-trained model, and an attention mechanism dynamically fuses the short-term and long-term embeddings into a comprehensive user representation. Beyond boosting recommendation accuracy over multiple baselines, our approach naturally supports explainability: the interpretable text summaries and attention weights can be exposed to end users, offering insights into why specific items are suggested. Experiments on real-world datasets underscore both the performance gains and the promise of generating clearer, more transparent justifications for content-based recommendations.

## CCS Concepts

• Information systems → Recommender systems.

## Keywords

Explainable User Modeling, Content-Based Recommendations, Large Language Models

## ACM Reference Format:

Milad Sabouri, Masoud Mansoury, Kun Lin, and Bamshad Mobasher. 2025. Towards Explainable Temporal User Profiling with LLMs. In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization (UMAP Adjunct '25)*, June 16–19, 2025, New York City, NY, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3708319.3733655>

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).

*UMAP Adjunct '25*, New York City, NY, USA

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-1399-6/25/06  
<https://doi.org/10.1145/3708319.3733655>

## 1 Introduction

Explainability has emerged as a critical requirement for recommender systems, enabling end users to understand—and hence trust—the reasons behind particular suggestions [23]. In content-based settings [10], where item metadata and textual descriptions drive user–item matching, many approaches continue to focus narrowly on maximizing predictive accuracy. They often overlook interpretability, commonly representing user preferences by simply averaging the embeddings of all items a user has interacted with [14]. While this “centric” approach is straightforward, it neglects the evolving nature of user interests and offers limited insight into *why* a recommendation is made.

To address these limitations, we propose a method that not only models user preferences more accurately but also inherently supports explanation. Our framework leverages large language models (LLMs) [24] to transform each user’s interaction history into two textual “profiles”: one highlighting short-term behaviors and another capturing stable, long-term tastes. These natural-language profiles are then encoded via a pre-trained BERT [5] model, resulting in two high-dimensional embeddings. An attention mechanism [22] dynamically fuses these embeddings into a single representation that drives recommendations. Crucially, because the original short-term and long-term summaries are human-readable, they serve as transparent rationales—an essential step toward explainability. We validate our approach on real-world datasets across two distinct domains. In particular, these domains differ in their average user-profile size, enabling us to investigate whether an LLM-driven fusion of short-term and long-term preferences is effective in contexts with more dynamic or stable user behaviors. The textual summaries generated for each user not only support interpretability for end users but also can serve as a valuable tool for system designers, offering insight into the model’s representation of user preferences and the rationale behind specific recommendations.

In addition, the attention mechanism applied over short-term and long-term embeddings exposes the relative influence of recent versus persistent user interests in shaping each recommendation. This dual-level transparency, combining human-readable summaries with interpretable model dynamics, can address practical demands in numerous settings where stakeholders increasingly require model outputs to be interpretable.

In this paper, we explore whether our LLM-based approach can improve ranking performance across domains characterized by different average profile sizes. We also show how the framework’s

textual summaries and attention mechanism can contribute to interpretability, clarifying for end users whether a recommendation is driven more by recent consumption patterns or by stable, enduring interests. Moreover, we investigate the relative impact of each architectural component in achieving strong results.

Our contributions are threefold. First, we propose a user-profiling strategy that integrates temporal signals into semantically rich, text-based representations, naturally suited for explanation. Second, we incorporate an attention mechanism that flexibly adapts to each user’s evolving interests, providing a clear, interpretable view of how these interests guide recommendations. Finally, we present an empirical evaluation on datasets reflecting contrasting profile sizes and preference dynamics, confirming that our approach not only boosts recommendation accuracy but also furnishes valuable explanatory insights. By unifying temporal user modeling with inherent interpretability, this work sets the stage for more transparent, trustworthy content-based recommender systems.

## 2 Related Works

Explainable recommendations [23] have increasingly become a focal point in recommender systems research, with recent works emphasizing the need to balance transparency and accuracy. For example, [4] proposes a theoretical framework for integrating explanations into recommendation pipelines, illustrating how even simple explanatory interfaces can boost user trust. Similarly, [2] interpretability by summarizing content and applying a linear attention mechanism, showing that clear textual justifications can coexist with competitive predictive performance. The trade-off between transparency and accuracy is examined in [25], which leverages user reviews to produce more interpretable suggestions. Meanwhile, user studies such as [3] highlight how varying the depth of explanations can significantly impact user satisfaction, indicating that the detail level should be adaptive rather than fixed.

Recent research has also begun exploring large language models (LLMs) as a foundation for explainability. [11] demonstrates that LLMs can generate high-quality, personalized explanations that help users better evaluate recommended items. Building on such capabilities, XRec[12] uses a lightweight collaborative adaptor to help LLMs interpret user-item interactions, thus producing thorough, human-readable explanations.

Meanwhile, Knowledge graph-based approaches [7] have leveraged side information to enhance transparency. Studies like [17] integrates knowledge graphs with LLM-generated reviews to highlight the reasoning paths behind recommendations, enabling interpretable item discovery even in cross-selling scenarios. Likewise, [18] introduces a powerful model-intrinsic method that extends KGAT (Knowledge Graph Attention Network) [21]. Specifically, they leverage large volumes of structured side information—including users’ demographics, item attributes, and multiple one-to-many relations—and compress it via latent-class modeling. Doing so reduces computational overhead and preserves interpretability by enabling direct visualization of which edges or neighbors drive the recommendation. This approach is particularly relevant to real-world e-commerce scenarios, in which companies store extensive multi-factor data. Their improvements highlight how “soft edges” can be used to refine attention weights and

speed up training, thereby removing a key practical bottleneck of knowledge-graph-based explainable models.

Our work shares similarities with [18], as both emphasize explainability in recommendation systems enhanced by advanced neural components. However, our methodology diverges in the nature of input data utilized: whereas [18] relies on side information derived from knowledge graphs, we leverage textual representations derived from user interaction histories. Specifically, we prompt large language models (LLMs) to summarize short-term versus long-term user behaviors, generating coherent and interpretable narratives suitable for direct user consumption. This integration of natural language within the recommendation pipeline effectively connects semantic insights from LLMs with temporal user behavior dynamics. Furthermore, our model explicitly differentiates between recent interests and persistent preferences, merging these aspects via an attention mechanism to produce intrinsically interpretable user profiles.

Despite substantial advances [2, 4, 18, 23] in generating robust and comprehensible explanations, existing methods predominantly focus either on static user profiles or strictly sequential behavioral patterns. Consequently, they often fail to account for the significant differences between transient interests and stable, long-term user preferences. While certain recommendation methods [9, 20, 26] incorporate temporal dynamics, they generally do not utilize LLM-generated narratives that clarify the temporal evolution of recommendations to end users. In contrast, our method explicitly partitions user interaction histories into short-term and long-term textual summaries, subsequently encoding and integrating these summaries through an attention mechanism to support recommendation generation. Importantly, the original textual summaries remain directly accessible as built-in explanations of user interests, transparently illustrating the interplay between recent and enduring preferences without requiring additional annotations or extensive post-processing. Thus, our approach bridges the gap between advanced LLM-driven explainability and nuanced temporal modeling of user behavior.

## 3 Explainable LLM-Driven Temporal User Profiling

Let  $\mathcal{U} = \{u_1, u_2, \dots, u_{|\mathcal{U}|}\}$  denote the set of users, and  $\mathcal{I} = \{i_1, i_2, \dots, i_{|\mathcal{I}|}\}$  denote the set of items. Each user  $u \in \mathcal{U}$  interacts with a subset of items  $\mathcal{I}_u \subseteq \mathcal{I}$ , where each interaction is associated with a timestamp  $t_{u,i}$ . Let  $\mathcal{H}_u = \{(i, t_{u,i}) : i \in \mathcal{I}_u\}$  represent the interaction history of user  $u$ , sorted in ascending order of timestamps.

The objective is to learn a *user embedding*  $\mathbf{e}_u \in \mathbb{R}^d$  and an *item embedding*  $\mathbf{e}_i \in \mathbb{R}^d$  for each user  $u$  and item  $i$ , such that the likelihood of user  $u$  interacting with item  $i$  can be predicted accurately. Specifically, we aim to estimate a function  $f : \mathbb{R}^d \times \mathbb{R}^d \rightarrow [0, 1]$  that predicts the interaction probability  $\hat{y}_{u,i}$  as:

$$\hat{y}_{u,i} = f(\mathbf{e}_u, \mathbf{e}_i), \quad (1)$$

where  $\hat{y}_{u,i}$  denotes the predicted likelihood of user  $u$  interacting with item  $i$ , and  $f$  is implemented as a multi-layer perceptron (MLP).

The proposed architecture for explainable LLM-driven temporal user profiling, illustrated in Figure 1, leverages temporal information and semantic representations to generate personalized *and*

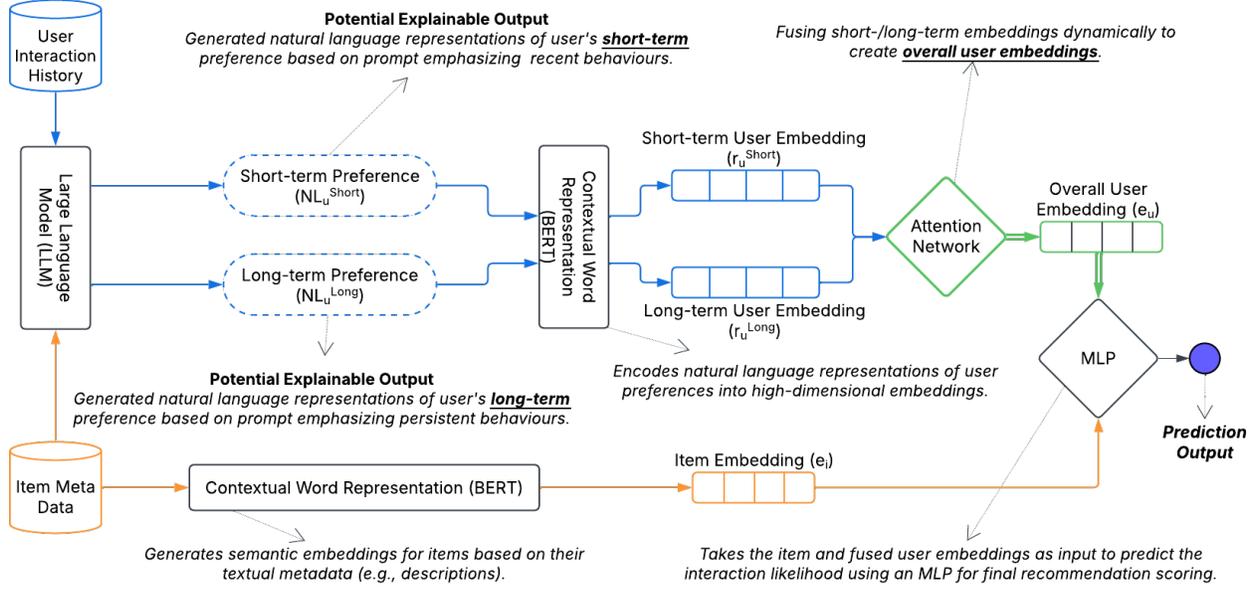


Figure 1: Proposed Architecture for LLM-Driven Temporal User Profiling

interpretable user embeddings. The model consists of three key components: (i) Temporal User Profile Creation, (ii) Embedding Representation, and (iii) Recommendation Generation.

Beyond capturing user dynamics, each component also contributes to transparency. Our framework generates human-readable text and user-specific attention weights, both of which can be exposed to end users or analysts as intrinsic explanations of why each recommendation is made.

### 3.1 User Profile Creation

To construct a temporal-aware user profile, we generate two distinct **natural language representations** of user preferences: *short-term preferences* and *long-term preferences*. Let  $\mathcal{H}_u = \{(i, t_{u,i}) : i \in \mathcal{I}_u\}$  denote the interaction history of user  $u$ , where  $t_{u,i}$  represents the timestamp of interaction with item  $i$ . The interaction history is sorted in ascending order of timestamps to maintain chronological order. We utilize a Large Language Model (LLM) to process the entire user interaction history twice, using distinct prompts to generate separate representations for short-term and long-term preferences:

*Pass 1: Short-Term Profile Generation.* The complete interaction history  $\mathcal{H}_u$  is provided as input to the LLM along with a task-specific prompt that instructs the model to extract the user’s *short-term interests*, placing greater emphasis on the most recent interactions while still considering temporal context. The LLM generates a natural language representation of short-term preferences, denoted as:

$$NL_u^{\text{short}} = \text{LLM}(\mathcal{H}_u, \text{Prompt}^{\text{short}}). \quad (2)$$

*Pass 2: Long-Term Profile Generation.* The same interaction history  $\mathcal{H}_u$  is passed to the LLM again, but this time with a different prompt that instructs the model to generate a *long-term preference profile* by considering the user’s entire history, capturing overarching interests and persistent behavioral patterns. The resulting representation is:

$$NL_u^{\text{long}} = \text{LLM}(\mathcal{H}_u, \text{Prompt}^{\text{long}}). \quad (3)$$

By explicitly utilizing distinct prompts while keeping the full interaction history intact, this approach enables the LLM to contextually differentiate between recent and persistent user preferences, addressing the limitations of traditional centric-based approaches that fail to incorporate temporal user dynamics.

Because these profiles are output as coherent text, they can be exposed directly to end users or system designers as an explanatory rationale. For example, the short-term profile might read:

*“The user recently interacted with several action-thriller titles and has shown interest in fast-paced content.”*

Such natural language descriptions *explicitly* reveal which recent patterns are driving the recommendation. Meanwhile, the long-term profile might include broader statements about the user’s historical favorites or consistent behavioral patterns, giving a transparent view of persistent interests.

### 3.2 Embedding Representation

The generated natural language representations are encoded into high-dimensional embeddings using a pre-trained BERT model [5], resulting in short-term user embedding and long-term user embedding, as follows:

$$\mathbf{r}_u^{\text{short}} = \text{BERT}(\text{NL}_u^{\text{short}}), \quad (4)$$

$$\mathbf{r}_u^{\text{long}} = \text{BERT}(\text{NL}_u^{\text{long}}). \quad (5)$$

Here,  $\mathbf{r}_u^{\text{short}} \in \mathbb{R}^d$  and  $\mathbf{r}_u^{\text{long}} \in \mathbb{R}^d$  are  $d$ -dimensional embeddings representing the short-term and long-term preferences of user  $u$ , respectively. Equations 4 and 5 describe the encoding process in which BERT converts each natural language profile into a dense vector representation. These embeddings capture both the semantic richness of user preferences and contextual nuances from interaction histories, serving as the basis for constructing the final user profile. Similarly, item descriptions are processed by BERT to generate item embeddings, ensuring that user and item representations lie in the same embedding space.

To construct the overall user embedding, an attention mechanism [22] is applied to dynamically fuse the short-term and long-term user embeddings. The attention mechanism computes personalized importance weights for short-term and long-term preferences, enabling the model to adaptively emphasize recent interactions for users with rapidly changing interests or highlight stable, overarching tendencies for users with persistent behaviors. This dynamic fusion mechanism is a key differentiator from the centric approach, where user profiles are created by simply averaging item embeddings, thereby neglecting temporal nuances.

Let  $\alpha_u^{\text{short}}$  and  $\alpha_u^{\text{long}}$  denote the attention weights assigned to the short-term and long-term preference embeddings, respectively. The attention weights are computed as:

$$\alpha_u^{\text{short}} = \frac{\exp(\mathbf{W}_a \mathbf{r}_u^{\text{short}})}{\exp(\mathbf{W}_a \mathbf{r}_u^{\text{short}}) + \exp(\mathbf{W}_a \mathbf{r}_u^{\text{long}})}, \quad (6)$$

$$\alpha_u^{\text{long}} = 1 - \alpha_u^{\text{short}}, \quad (7)$$

where  $\mathbf{W}_a \in \mathbb{R}^{1 \times d}$  is a learnable parameter vector, and  $\exp(\cdot)$  denotes the exponential function. Equation 6 computes the attention weight for short-term preferences, while Equation 7 ensures that the sum of these attention weights equals 1.

The overall user embedding  $\mathbf{e}_u$  is then obtained as a weighted sum of the short-term and long-term embeddings:

$$\mathbf{e}_u = \alpha_u^{\text{short}} \cdot \mathbf{r}_u^{\text{short}} + \alpha_u^{\text{long}} \cdot \mathbf{r}_u^{\text{long}}, \quad (8)$$

where  $\mathbf{e}_u \in \mathbb{R}^d$  represents the final user embedding. As shown in Equation 8, the attention mechanism enables the model to dynamically balance the influence of short-term and long-term preferences based on user-specific interaction patterns.

The learned attention weights  $\alpha_u^{\text{short}}$  and  $\alpha_u^{\text{long}}$  can be surfaced in a real system to communicate the *relative importance* of the user’s recent vs. historical activities. For instance, if  $\alpha_u^{\text{short}}$  is very high for a user, it indicates that the model believes the user’s latest interactions dominate their current recommendations—a straightforward, interpretable explanation for why certain items appear more prominently.

### 3.3 Recommendation Generation

The final user embedding  $\mathbf{e}_u$ , obtained from the attention mechanism, is concatenated with the corresponding item embedding  $\mathbf{e}_i$  and passed through a multi-layer perceptron (MLP) [15] to predict the interaction probability of user  $u$ :

$$\hat{y}_{u,i} = \text{MLP}([\mathbf{e}_u; \mathbf{e}_i]), \quad (9)$$

where  $[\mathbf{e}_u; \mathbf{e}_i] \in \mathbb{R}^{2d}$  denotes the concatenation of the user embedding  $\mathbf{e}_u$  and the item embedding  $\mathbf{e}_i$ , and  $\hat{y}_{u,i} \in [0, 1]$  represents the predicted probability that user  $u$  will interact with item  $i$ .

As shown in Equation 9, the MLP acts as a scoring function that learns non-linear interactions between user and item embeddings. The final output is a probability indicating the likelihood of interaction. The MLP consists of multiple fully connected layers with ReLU activation functions, followed by a sigmoid output layer to ensure that the output lies in the range  $[0, 1]$ . The model is trained using a binary cross-entropy loss, where the target label  $y_{u,i} \in \{0, 1\}$  indicates whether user  $u$  interacted with item  $i$ :

$$\mathcal{L} = -\frac{1}{|\mathcal{D}|} \sum_{(u,i,y_{u,i}) \in \mathcal{D}} (y_{u,i} \log \hat{y}_{u,i} + (1 - y_{u,i}) \log(1 - \hat{y}_{u,i})), \quad (10)$$

where  $\mathcal{D}$  represents the set of all user–item interaction pairs in the training dataset. Equation 10 penalizes the difference between the predicted probability  $\hat{y}_{u,i}$  and the true label  $y_{u,i}$ , driving the model to learn accurate predictions.

### 3.4 Model Highlights and Strengths

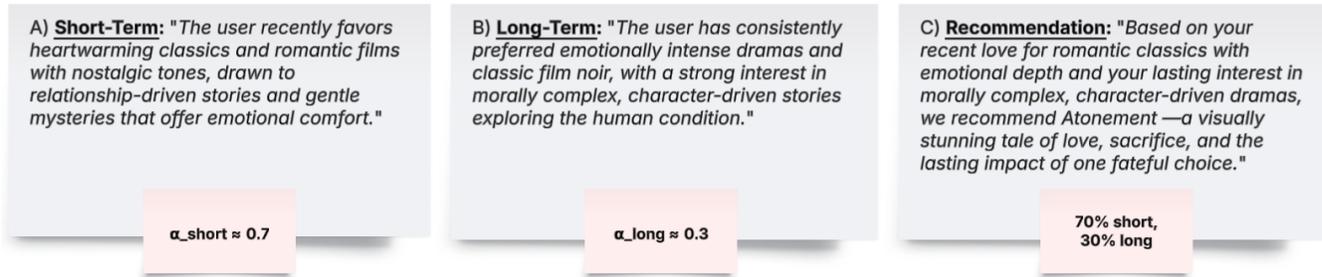
Compared to a simpler “centric” baseline that averages item embeddings without distinguishing short-term from long-term preferences, the present architecture offers two key benefits:

- (1) *Temporal Awareness*: By explicitly modeling short-term and long-term preferences and fusing them via an attention mechanism, the model better captures *how* user interests shift over time, leading to improved performance.
- (2) *Intrinsic Explainability*: Each user’s short-term and long-term text profiles can be exposed directly, while the attention weights inform whether recent or persistent behaviors drive the recommendation. This provides a transparent rationale far more meaningful than a single averaged vector of past items.

In essence, the synergy of LLM-based text generation and attention-based fusion yields highly *semantic* user embeddings that not only boost recommendation metrics (e.g., Recall@K, NDCG@K) but also enable end users and system developers to *see* precisely which aspects of a user’s recent or long-running behaviors are influencing each recommendation.

## 4 Experiments & Explainability Demonstration

In this research, our primary focus is to propose an LLM-based method for creating explainable user profiles, specifically for content-based recommendation systems. The goal of these experiments is to evaluate the effectiveness of our approach in incorporating temporal dynamics, improving recommendation quality, and generating more interpretable user profiles compared to baseline models. Additionally, we showcase how our framework’s outputs can be used to



**Figure 2: A conceptual illustration of our framework’s potential explainability. (A) and (B) depict the short-term vs. long-term textual profiles our current approach generates and encodes. (C) shows a hypothetical extension wherein the final recommendation is explicitly justified by both sets of preferences—an aspect we plan to explore in future work.**

offer transparent justifications for recommended items. The code, data, prompts, and experimental results are available on GitHub<sup>1</sup>.

#### 4.1 Datasets

We evaluated our proposed method on two public datasets drawn from Amazon Product Reviews [8], specifically the Movies&TV and Video Games categories. We deliberately selected these two domains to observe our approach’s effectiveness under contrasting user-profile characteristics:

*Movies&TV:* Users in this domain often have larger and more diverse interaction histories, which are subject to frequent changes in taste (e.g., seasonal shifts, trend-driven watching habits). This makes Movies&TV a natural testbed for short-term vs. long-term preference modeling.

*Video Games:* By contrast, many gamers have fewer overall interactions and more stable preferences—often remaining loyal to specific genres or franchises. This dataset thus enables us to examine whether our method’s LLM-driven, temporally aware profile construction retains benefits when user histories are relatively sparse and less variable.

To ensure high-quality item representations, we filtered the metadata to keep only items with mostly English text and descriptions of at least 500 characters, providing enough context for the LLM to model user preferences. Both datasets provide interaction logs with timestamps, allowing us to incorporate temporal behavior into recommendation. Additionally, each item’s textual metadata (titles, descriptions) was leveraged to create BERT-based item embeddings, used by our method and all baselines for a fair comparison. Table 1 presents the high-level dataset details—total users, total items, and overall interaction counts. Meanwhile, Table 2 provides user-profile metrics such as the mean, median, and standard deviation of user interactions. Taken together, these two tables reveal the notable difference between Movies&TV—where users typically have larger, more varied histories—and Video Games—where user profiles tend to be smaller and more stable.

#### 4.2 Baselines

Because our work centers on user profiling for content-based recommendation, we primarily compare our proposed LLM-driven strategy against user-profiling methods commonly employed in

such systems. These baselines differ in how they model user histories, incorporate (or exclude) temporal factors, and utilize textual or collaborative information. By including diverse approaches, we thoroughly test the incremental benefit of (1) large language models and (2) explicit short-term vs. long-term preference profiling.

**4.2.1 Centric Approach (Centric).** As a main baseline, we use a *centric* approach wherein user profiles are created by averaging all item embeddings the user has interacted with (no temporal segmentation). Other components, such as item embeddings and the final MLP architecture, remain identical to our proposed method to ensure a fair comparison. This reflects the common practice [6] of deriving user vectors by straightforward aggregation, offering a clear benchmark against which we can measure the gains from explicitly modeling recent vs. persistent behavior.

**4.2.2 Temporal Fusion without LLM-Based Representations (Temp-Fusion).** This baseline represents a time-aware content-based approach that forgoes LLM-generated textual summaries. It splits a user’s interaction history into short-term and long-term segments (with dataset-specific cutoffs), encodes each segment via BERT embeddings, and applies an attention mechanism to fuse them—similar to our framework’s temporal fusion. However, rather than generating natural language profiles to be encoded by BERT, it simply aggregates the raw item embeddings from the user’s recorded interactions. By contrasting this design against our LLM-based strategy, we isolate how much semantically rich text-based representation impacts performance and explainability.

**4.2.3 Popularity-Based Recommendation (Popularity).** This non-personalized baseline ranks items solely based on their overall popularity, measured by the frequency of interactions aggregated across all users. As a control condition, this baseline represents a minimal standard of recommendation performance that does not involve personalization or user-item modeling. Its inclusion allows us to contextualize the effectiveness of our proposed method by contrasting against a naïve strategy that ignores both user preference and content signals. Any substantial improvement over this baseline indicates the necessity of modeling beyond mere popularity trends.

**4.2.4 Matrix Factorization (MF).** Matrix Factorization is a widely used collaborative filtering method that learns latent representations for users and items based only on observed interaction patterns. As a control condition, MF offers a direct point of comparison

<sup>1</sup><https://github.com/milsab/TETUP>

**Table 1: Datasets Statistics**

Dataset	No. Unique Users	No. Unique Items	No. Interactions
Movies&TV	10,000	14,420	202,583
Games	10,371	3,790	83,842

**Table 2: User Profiles Statistics**

Dataset	Mean	Median	Mode	Std Dev
Movies&TV	11.79	9.00	6	9.80
Games	4.55	3.00	3	3.97

for evaluating our content-based approach. It captures user preferences through interaction data but lacks any understanding of item content or metadata. By including MF, we isolate the impact of incorporating content features into the recommendation process. This baseline helps clarify whether observed improvements are due to the content-aware nature of our model, rather than the presence of collaborative patterns alone—particularly in settings where user-item interactions are sparse or noisy.

Collectively, these baselines illuminate the contributions of our approach’s core components: *temporal modeling*, *LLM-based semantic representation*, and *attention-based fusion*. By systematically comparing results, we show that our LLM-driven model, which yields both accurate predictions and transparent textual rationales, consistently outperforms less sophisticated user-profiling methods.

### 4.3 Evaluation Methodology

Because we aim to improve *content-based user profiling*, we focus on how effectively each method learns and applies user embeddings for ranking items. Thus, the centric approach—which simply averages item embeddings over a user’s history—serves as a fundamental baseline. Unlike sequential recommenders [9, 19, 20] that evaluate next-item prediction via metrics like Hit Rate or MRR, we adopt a ranking-oriented protocol. Specifically, we compute Recall@K and NDCG@K by asking each model to rank relevant (held-out) items amidst a larger candidate set. This directly tests *how well the user profiles capture user intent*, which is central to content-based personalization rather than pure sequence prediction.

To incorporate temporal awareness in the training process, we sort each user’s interactions chronologically and split them into: 1) *Training set*: Earliest 60% of interactions, 2) *Validation set*: The subsequent 20% of interactions, and 3) *Test set*: Most recent 20% of interactions. This time-wise splitting strategy ensures that user preferences evolve naturally from training to test, allowing us to evaluate the models’ ability to predict future interactions based on historical behavior.

### 4.4 Experimental Setup

All experiments were implemented in PyTorch [13] and trained using the Adam optimizer. Early stopping (patience=5) halted training if no validation improvement was seen for 5 epochs, preventing overfitting. Each model trained up to 100 epochs, with a batch size tuned among {512, 1024, 2048}; we present results at 2048 for best overall performance.

To encode the textual user summaries we generate for short-term and long-term profiles, we used the all-MiniLM-L6-v2 variant of SBERT [16], which yields 384-dimensional embeddings. Our method extracts short-term and long-term natural language representations of each user’s interaction history via the GPT-4o-mini model [1], known for high-quality text generation. The short-term, long-term, and item embeddings all have dimension 384.

A multi-layer perceptron (hidden dimension=128, dropout=0.2) was used to predict user-item interaction probabilities. We included 5 negative samples per positive interaction. The learning rate was fixed at 1e-3, which provided stable and reliable convergence.

The experiments were executed on an environment using four NVIDIA A100 GPUs, each with 80GB of memory, enabling efficient parallelization and acceleration of the training process. This setup, along with robust optimization techniques and careful hyperparameter tuning, ensured reliable and reproducible results across all experiments.

### 4.5 Experimental Results & Discussion

We now evaluate the effectiveness of our LLM-based approach on two datasets: Movies&TV and Video Games. Table 3 details our method’s ranking performance for Movies&TV and Video Games, while Table 2 provides a closer look at user profiles characteristics. Two observations emerge:

#### 1) Effectiveness in Larger, Frequent-Shifting Profiles:

In the Movies&TV domain, each user’s profile tends to be bigger (mean of 11.79 interactions, with a high standard deviation of 9.80), implying diverse and more rapidly changing preferences. Table 3 shows our approach significantly improves Recall@10 (by 17%) and NDCG@10 (by 14%) over Centric, with similarly strong gains for Recall@20 and NDCG@20 (12% and 9% increases, respectively). These results suggest that when users accumulate larger histories and exhibit more varied consumption patterns, short-term vs. long-term fusion—enriched by LLM-based text—delivers notable accuracy gains. Among the baselines, Temp-Fusion comes closest but lacks the semantic depth of our natural language profiles, reinforcing the advantage of explicit textual summaries in capturing nuanced, evolving tastes.

#### 2) Mixed Results in Smaller, Steadier Profiles:

By contrast, Video Games users’ profiles are typically shorter (mean of 4.55 interactions, median of 3.0) with a lower standard deviation (3.97), indicating more stable and less frequent content shifts. Under such conditions, our approach still achieves the highest Recall@20, affirming that even minimal LLM-derived text can help with broader item ranking. However, Temp-Fusion narrowly edges out our method on Recall@10, NDCG@10, and NDCG@20. We attribute this partly to:

*Sparse Histories*: With only a few recorded interactions per user, splitting into short-term vs. long-term segments can dilute the available signals.

**Table 3: Performance Comparison of our proposed method and baselines on Movies&TV and Video Games datasets for  $K = 10$  and  $K = 20$ . (asterisk indicates that the improvement of the proposed method over the baseline method (Centric) is statistically significant ( $p < 0.05$ ))**

Method	Movies&TV				Video Games			
	Recall@10	NDCG@10	Recall@20	NDCG@20	Recall@10	NDCG@10	Recall@20	NDCG@20
(a) Centric	0.0113	0.0191	0.0199	0.0269	0.0645	0.0532	0.0932	0.0649
(b) Popularity	0.0082	0.0145	0.0133	0.0191	0.0397	0.0324	0.0706	0.0453
(c) MF	0.0048	0.0085	0.0087	0.0124	0.0457	0.0370	0.0754	0.0491
(d) Temp-Fusion	<u>0.0118</u>	<u>0.0201</u>	<u>0.0207</u>	<u>0.0276</u>	<b>0.0693</b>	<b>0.0589</b>	<u>0.0982</u>	<b>0.0712</b>
(e) LLM-Based (Proposed)	<b>0.0132*</b>	<b>0.0217*</b>	<b>0.0223*</b>	<b>0.0293*</b>	<u>0.0665*</u>	<u>0.0547*</u>	<b>0.1021*</b>	<u>0.0683*</u>
<b>Gain of Proposed vs. (a)</b>	<b>17%</b>	<b>14%</b>	<b>12%</b>	<b>9%</b>	<b>3%</b>	<b>3%</b>	<b>10%</b>	<b>5%</b>

*Relatively Fixed Interests:* Many gamers stay loyal to specific franchises or genres, reducing the benefit of dynamic preference modeling.

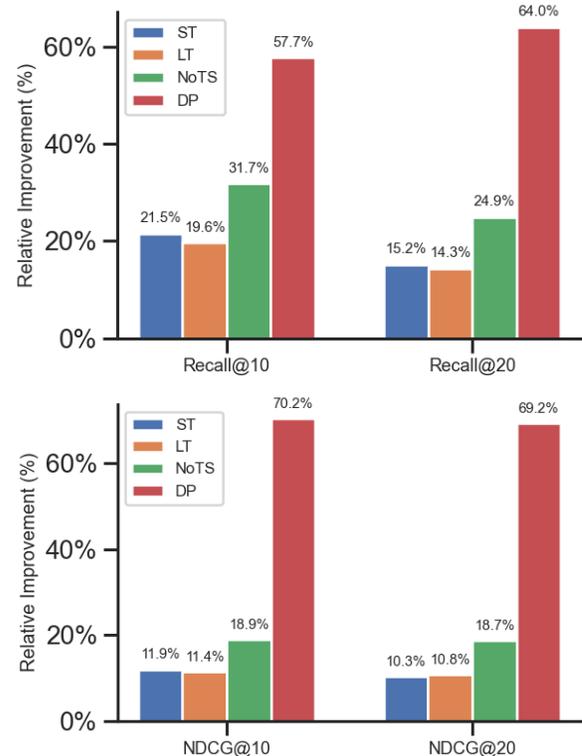
Overall, these findings highlight how dataset properties, particularly user-profile size and preference variability, impact the efficacy of our short-term–long-term LLM-driven approach. For domains like Movies&TV, where user preferences shift more frequently and historical data is abundant, our method excels. In domains with smaller, steadier profiles, its edge is narrower, yet it still maintains the best performance in some metrics (Recall@20) and preserves an explainability advantage through textual profiles.

With regard to explainability, Figure 2 demonstrates how our framework’s short-term and long-term textual profiles contribute to explainable user modeling, while also suggesting a future extension for full user-facing recommendations. Concretely, we achieve *explainability by construction* in two ways:

1) *Textual Short-Term and Long-Term Summaries:* Rather than producing only abstract embeddings, our model outputs natural language describing the user’s recent (Figure 2-A) versus historical (Figure 2-B) preferences. For instance, the short-term text might emphasize “heartwarming classics” or “romantic films with nostalgic tones,” while the long-term text highlights “emotionally intense dramas” or “morally complex, character-driven stories.” Because these summaries are easily comprehensible, users (or developers) can see exactly why certain items fit their current or enduring interests.

2) *Attention-Driven Transparency:* Each recommendation is governed by attention weights  $\alpha_u^{\text{short}}$  and  $\alpha_u^{\text{long}}$ . In a future scenario (sketched in Figure 2-C), if short-term interests carry 70% of the weight, the recommendation might highlight the user’s most recent preferences. Conversely, a higher long-term weight could favor timeless favorites. This numeric breakdown is straightforward to communicate (e.g., “70% recent tastes, 30% long-standing preferences”), which would further reinforce user trust by clarifying why one item outperforms another.

While our current experiments focus on the textual profiles depicted in parts (A) and (B), these same signals lay the foundation for part (C): a user-facing explanation that explicitly merges short-term and long-term rationales. By unifying textual summaries with attention proportions, the final recommendation could be explained in a way that goes beyond typical opaque embeddings. Although a formal user study remains an opportunity for future work, our

**Figure 3: Relative improvement of our complete approach over four ablation variants on the Movies&TV dataset. Bars show percentage gains in Recall@K and NDCG@K.**

framework already offers transparent user profiling, which many baselines lack.

## 5 Ablation Study

We conducted an ablation study to isolate key architectural and design choices, allowing us to assess the contribution of individual components in our proposed method. The experiments include variations that use only short-term or long-term preferences, general preferences without temporal distinction, and an alternative scoring mechanism (dot product) in place of the MLP. Since the results for the Games dataset exhibited a similar pattern, we focus on the Movies dataset for clarity. Tables 4 and 5 show the numerical

**Table 4: Recall@K Comparison for Ablation Variants (Movies&TV Dataset)**

Experiment	Recall@10	Recall@20
Short-Term Only (ST)	0.0109	0.0193
Long-Term Only (LT)	0.0110	0.0195
General Preferences (No TS)	0.0100	0.0178
Dot-Product Scoring (DP)	0.0084	0.0136
LLM-Based (Proposed)	<b>0.0132</b>	<b>0.0223</b>

**Table 5: NDCG@K Comparison for Ablation Variants (Movies&TV Dataset)**

Experiment	NDCG@10	NDCG@20
Short-Term Only (ST)	0.0194	0.0266
Long-Term Only (LT)	0.0195	0.0265
General Preferences (No TS)	0.0183	0.0247
Dot-Product Scoring (DP)	0.0128	0.0173
LLM-Based (Proposed)	<b>0.0217</b>	<b>0.0293</b>

outcomes (Recall@K, NDCG@K), while Figure 3 presents a bar chart illustrating how these variants compare. To structure our discussion effectively, we analyze each variant below.

### 5.1 Short-Term Preferences Only (ST)

This variant exclusively focuses on short-term user interests, relying on  $NL_u^{\text{short}}$  as the profile representation. We encode these textual summaries via BERT, then combine the resulting user embeddings with item embeddings in an MLP for final score computation. The goal is to isolate how well rapidly changing behaviors alone can drive recommendations, ignoring broader historical context. As shown in Tables 4 and 5, the ST variant underperforms compared to Long-Term Only (LT). This suggests that while short-term signals are helpful, they alone cannot produce a robust user profile. The gap is most pronounced at larger  $K$  values. In the bar chart (Figure 3), ST’s relative improvement remains smaller than that of the full model, reinforcing the need to account for enduring user tendencies along with recent shifts.

### 5.2 Long-Term Preferences Only (LT)

Here, we rely on  $NL_u^{\text{long}}$  exclusively—reflecting general, persistent user preferences—again encoded via BERT and combined with item embeddings in an MLP. This setup measures how effectively stable user interests can drive recommendations in the absence of short-term signals. From Tables 4 and 5, LT yields respectable performance, highlighting that many users show relatively consistent tastes. However, LT still lags behind our full method, as visible in Figure 3, indicating the advantage of incorporating both recent and historical behaviors rather than relying on long-term context alone.

### 5.3 General Preferences without Temporal Distinction (NoTS)

Unlike ST and LT, the “NoTS” setting collapses all user interactions into a single textual summary, dropping any notion of *when* an item was consumed. By discarding short-term vs. long-term segmentation, we test whether explicit temporal partitioning is vital for robust user modeling.

Tables 4 and 5 show that NoTS outperforms the dot-product variant but remains behind ST and LT, signifying that ignoring time-based differentiation dilutes important signals. In Figure 3, NoTS consistently trails the short-term and long-term variants across the evaluated metrics. This confirms that temporal segmentation improves personalization, as recent behaviors and historical interests each contain unique information about user preferences.

### 5.4 Substitute MLP with Dot Product (DP)

This variant preserves short-term and long-term user embeddings but replaces the MLP scoring function with a simple dot product. We thus examine whether a non-linear component is necessary to fully exploit the embeddings’ expressive capacity. The dot-product variant (DP) underperforms all others in Tables 4 and 5. Figure 3 further illustrates its shortcomings. Linear scoring captures fewer complex correlations, confirming that a non-linear MLP significantly reinforces ranking accuracy.

### 5.5 Proposed LLM-Based Method

Our complete approach combines short-term and long-term textual profiles, each encoded by BERT, then fused by attention before being fed into an MLP. This design accounts for both immediate user context (short-term) and stable interests (long-term), while the MLP models non-linear user–item interactions. Across all ablation variants, our proposed method dominates in Tables 4 and 5. Figure 3 shows especially large gains when integrating all components—temporal separation, LLM-generated text, attention-based fusion, and an MLP. Collectively, these results confirm that each element is crucial. Short-term and long-term segmentation balances recent context with persistent tastes, LLM-based text enriches user representations, attention weights adapt to each user’s dynamic interest, and non-linear scoring provides deeper interaction modeling. Hence, the ablation study demonstrates that integrating *all* these aspects is essential for maximizing both accuracy and interpretability, validating our design choice to treat short-term vs. long-term signals as separately encoded textual profiles, fuse them adaptively, and employ a learnable MLP for final predictions.

## 6 Conclusion

We presented a content-based recommendation framework that integrates temporal user modeling and LLM-driven textual profiling into an inherently explainable pipeline. Rather than averaging user interactions, our approach explicitly separates short-term and long-term preferences, encodes them via BERT, and fuses the resulting embeddings with an attention mechanism. Experiments across Movies&TV (larger, frequently shifting profiles) and Video Games (smaller, more stable profiles) confirm that time-aware, text-based user representations yield up to 17% higher Recall@10 and 14% higher NDCG@10 compared to a standard centric baseline.

Our method already offers built-in explainability by exposing short-term and long-term textual profiles and attention weights, clarifying whether recommendations stem from recent spikes or deeper, enduring interests. For future work, we plan to extend this design to generate fully user-facing explanations—explicitly weaving together both textual profiles and their attention ratios to explain *why* a particular item is recommended.

## References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [2] Jun Ai, Haolin Li, Zhan Su, and Fengyu Zhao. 2025. An explainable recommendation algorithm based on content summarization and linear attention. *Neuro-computing* 630 (2025), 129692.
- [3] Mohamed Amine Chatti, Mouadh Guesmi, Laura Vorgerd, Thao Ngo, Shoeb Joarder, Qurat Ul Ain, and Arham Muslim. 2022. Is more always better? The effects of personal characteristics and level of detail on the perception of explanations in a recommender system. In *Proceedings of the 30th ACM Conference on User Modeling, Adaptation and Personalization*. 254–264.
- [4] Janneth Chicaiza and Priscila Valdiviezo-Diaz. 2023. Explainable recommender systems: from theory to practice. In *Intelligent Sustainable Systems: Selected Papers of WorldS4 2022, Volume 2*. Springer, 449–459.
- [5] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Jill Burstein, Christy Doran, and Thamar Solorio (Eds.). Association for Computational Linguistics, Minneapolis, Minnesota, 4171–4186. <https://doi.org/10.18653/v1/N19-1423>
- [6] Ziwei Fan, Zhiwei Liu, Shelby Heinecke, Jianguo Zhang, Huan Wang, Caiming Xiong, and Philip S Yu. 2023. Zero-shot item-based recommendation via multi-task product knowledge graph pre-training. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 483–493.
- [7] Qingyu Guo, Fuzhen Zhuang, Chuan Qin, Hengshu Zhu, Xing Xie, Hui Xiong, and Qing He. 2020. A survey on knowledge graph-based recommender systems. *IEEE Transactions on Knowledge and Data Engineering* 34, 8 (2020), 3549–3568.
- [8] 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).
- [9] Wang-Cheng Kang and Julian McAuley. 2018. Self-attentive sequential recommendation. In *2018 IEEE international conference on data mining (ICDM)*. IEEE, 197–206.
- [10] Pasquale Lops, Marco De Gemmis, and Giovanni Semeraro. 2011. Content-based recommender systems: State of the art and trends. *Recommender systems handbook* (2011), 73–105.
- [11] Sebastian Lubos, Thi Ngoc Trang Tran, Alexander Felfernig, Seda Polat Erdeniz, and Viet-Man Le. 2024. Llm-generated explanations for recommender systems. In *Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization*. 276–285.
- [12] Qiyao Ma, Xubin Ren, and Chao Huang. 2024. Xrec: Large language models for explainable recommendation. *arXiv preprint arXiv:2406.02377* (2024).
- [13] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).
- [14] Marco Polignano, Cataldo Musto, Marco de Gemmis, Pasquale Lops, and Giovanni Semeraro. 2021. Together is better: Hybrid recommendations combining graph embeddings and contextualized word representations. In *Proceedings of the 15th ACM conference on recommender systems*. 187–198.
- [15] Marius-Constantin Popescu, Valentina E Balas, Liliana Perescu-Popescu, and Nikos Mastorakis. 2009. Multilayer perceptron and neural networks. *WSEAS Transactions on Circuits and Systems* 8, 7 (2009), 579–588.
- [16] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. <https://arxiv.org/abs/1908.10084>
- [17] Guangsi Shi, Xiaofeng Deng, Linhao Luo, Lijuan Xia, Lei Bao, Bei Ye, Fei Du, Shirui Pan, and Yuxiao Li. 2024. Llm-powered explanations: Unraveling recommendations through subgraph reasoning. *arXiv preprint arXiv:2406.15859* (2024).
- [18] Ryotaro Shimizu, Megumi Matsutani, and Masayuki Goto. 2022. An explainable recommendation framework based on an improved knowledge graph attention network with massive volumes of side information. *Knowledge-Based Systems* 239 (2022), 107970.
- [19] 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.
- [20] Yong Kiam Tan, Xinxing Xu, and Yong Liu. 2016. Improved recurrent neural networks for session-based recommendations. In *Proceedings of the 1st workshop on deep learning for recommender systems*. 17–22.
- [21] Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. 2019. Kgat: Knowledge graph attention network for recommendation. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*. 950–958.
- [22] A Waswani, N Shazeer, N Parmar, J Uszkoreit, L Jones, A Gomez, L Kaiser, and I Polosukhin. 2017. Attention is all you need. In *NIPS*.
- [23] Yongfeng Zhang, Xu Chen, et al. 2020. Explainable recommendation: A survey and new perspectives. *Foundations and Trends® in Information Retrieval* 14, 1 (2020), 1–101.
- [24] Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223* (2023).
- [25] Yao Zhou, Haonan Wang, Jingrui He, and Haixun Wang. 2025. Review-Based Explainable Recommendations: A Transparency Perspective. *ACM Transactions on Recommender Systems* (2025).
- [26] Yu Zhu, Hao Li, Yikang Liao, Beidou Wang, Ziyu Guan, Haifeng Liu, and Deng Cai. 2017. What to do next: Modeling user behaviors by time-LSTM.. In *IJCAI*, Vol. 17. 3602–3608.