G-Refer: Graph Retrieval-Augmented Large Language Model for Explainable Recommendation

Anonymous Author(s)

ABSTRACT

Explainable recommendation has demonstrated significant advantages in informing users about the logic behind recommendations, thereby increasing system transparency, effectiveness, and trustworthiness. To provide personalized and interpretable explanations, existing works often combine the generation capabilities of large language models (LLMs) with collaborative filtering (CF) information. CF information extracted from the user-item interaction graph captures the user behaviors and preferences, which is crucial for providing informative explanations. However, due to the complexity of graph structure, effectively extracting the CF information from graphs still remains a challenge. Moreover, existing methods often struggle with the integration of extracted CF information with LLMs due to its implicit representation and the modality gap between graph structures and natural language explanations. To address these challenges, we propose G-Refer, a framework using Graph Retrieval-augmented large language models (LLMs) for explainable recommendation. Specifically, we first employ a hybrid graph retrieval mechanism to retrieve explicit CF signals from both structural and semantic perspectives. The retrieved CF information is explicitly formulated as human-understandable text by the proposed graph translation and accounts for the explanations generated by LLMs. To bridge the modality gap, we introduce knowledge pruning and retrieval-augmented fine-tuning to enhance the ability of LLMs to process and utilize the retrieved CF information to generate explanations. Extensive experiments show that G-Refer achieves superior performance compared with existing methods in both explainability and stability. Codes and data are available at https://anonymous.4open.science/r/G-Refer.

CCS CONCEPTS

• Computing methodologies → Artificial intelligence; • Information systems \rightarrow Information retrieval.

KEYWORDS

Explainable Recommendation, Large Language Model, GraphRAG

ACM Reference Format:

Anonymous Author(s). 2018. G-Refer: Graph Retrieval-Augmented Large Language Model for Explainable Recommendation. In Proceedings of (Conference acronym 'XX). ACM, New York, NY, USA, 12 pages. https://doi.org/ 10.1145/nnnnnn.nnnnnn

58

1 INTRODUCTION

Recommendation systems (RS) play an essential role in our daily lives, as they affect how people navigate the vast amounts of products available in online services [21, 37, 56]. These systems have the ability to predict whether a user is interested in an item, e.g., by clicking on it or making a purchase. Recently, explainable recommendation [45, 55, 70, 71] has attracted more attention and demonstrated significant advantages in informing users about the logic behind their received recommendation results, thereby increasing system transparency, effectiveness, and trustworthiness. Specifically, explainable recommendation aims to generate humanunderstandable textual explanations for each user-item recommendation. This enables us to understand the underlying reasons behind each recommendation and develop accountable RS as a result [7].

Due to the impressive generative and reasoning capabilities of large language models (LLMs) [74], existing explainable recommendation methods [30, 31, 39] can now produce fluent and informative explanations in natural language based on user and item profiles, as well as their interactions. To provide more detailed and personalized explanations, recent works [39] have utilized user-item interaction graphs, which contain abundant collaborative filtering (CF) information from graph structure. The CF information [22] reveals the complex patterns between users and items, which can be used to generate more accurate and informative explanations. Nevertheless, the graph structure is inherently noisy and complex, making it challenging to effectively extract CF information for explanations.

In order to better model the CF information from graph structure in RS, graph neural networks (GNNs) have been widely adopted and demonstrated exceptional performance in recommendation tasks [2, 20, 57]. GNNs capture CF information by learning hidden representations of both users and items through an iterative process of feature transformation and information aggregation from neighboring nodes. While GNNs excel at capturing CF information for recommendations, simply incorporating GNNs with LLM-based RS such as XRec [39] still faces several challenges in producing satisfactory explanations for recommendations: (C1) Implicit CF signal. The CF signals are injected into LLMs by feeding user and item embeddings generated by GNNs, where the CF signals are represented as implicit node embeddings. Given the vagueness of the explainability of GNNs themselves [65], it is difficult to interpret the CF signals contained in these embeddings. Therefore, more explicit evidence constructed from the CF signals contained in graphs, especially favored from the view of human-understandable text, is eagerly needed. (C2) Modality gap. Since GNNs are primarily used to capture structural information, the rich semantics within user and item profiles are inevitably ignored. In addition, LLMs struggle to directly understand the structured CF signals captured by GNNs, because of their unstructured nature. This results in a significant modality gap between structured representations and

59

60

61

62

63

64

65

66

67

68

69

70

71

72

73

74

75

76

77

78

79

80

81

82

83

84

85

86

87

88

89

90

91

92

93

94

95

96

97

98

99

100

101

102

103

104

105

106

107

108

109

110

111

112

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.

Conference acronym 'XX, June 03-05, 2018, Woodstock, NY 55

^{© 2018} Copyright held by the owner/author(s). Publication rights licensed to ACM. 56 ACM ISBN 978-x-xxxx-xxxx-x/YY/MM 57

https://doi.org/10.1145/nnnnnnnnnnnn



Figure 1: Our proposed pipeline G-Refer facilitates explainable recommendation with three key components: (1) Hybrid Graph Retrieval employs multi-granularity retrievers to retrieve explicit CF signals and formulated as human-readable text by the Graph Translation; (2) Knowledge Pruning eliminates noise and improves training efficiency; and (3) Retrieval-augmented Fine-tuning instructs LLMs to leverage retrieved CF information in generating informative explanation.

natural language, making it difficult for LLMs to effectively leverage graph-derived CF information when generating explanations.

Presented Work. Motivated by these challenges, in this work, we propose G-Refer, a framework designed to use Graph Retrievalaugmented LLMs for explainable recommendation. We aim to address the limitations of existing methods and extract explicit, diverse, and semantically rich CF information to generate more accurate and personalized explanations. As shown in Figure 1, given a user-item recommendation to be explained, we leverage a hybrid graph retrieval mechanism that combines multi-granularity retrievers to extract explicit structural and semantic CF information from the graphs (to address C1). Specifically, we employ a path-level retriever to capture structural CF information by identifying k paths from graphs that account for recommendation. To utilize the rich semantics within nodes, we also employ a node-level retriever to retrieve semantic CF information from the graphs by finding the most relevant nodes to the recommendation. The retrieved CF information is then translated into human-understandable text with the graph translation module to facilitate LLMs in generating explanations. To bridge the modality gap and enhance the understanding of LLMs (to address C2), we adopt a lightweight retrieval-augmented fine-tuning (RAFT) approach to instruct LLMs in understanding the retrieved CF information and generating explanations. To further improve the training efficiency and reduce noise, we introduce a knowledge pruning technique to filter out training samples with less relevant CF information. After the training, LLMs exhibit a greater ability to leverage both the retrieved CF information and rich semantics in profiles to generate accurate and contextually relevant explanations for recommendations.

Our main contributions can be summarized as follows:

- **Comprehensive Analysis.** We identify the challenges in existing explainable recommendation works, as GNNs struggle to capture explicit and semantically rich CF information.
- Architecture Design. We propose G-Refer, a model leveraging hybrid graph retrieval to capture both structural and semantic CF signals from user-item interaction graphs. We also incorporate knowledge pruning and retrieval-augmented fine-tuning to enhance LLMs' ability to utilize retrieved knowledge.

• **Superior Performance.** Extensive experiments on public datasets demonstrate the effectiveness of our proposed G-Refer, surpassing a series of SOTA baselines by up to 8.67%.

2 PRELIMINARIES

2.1 Explainable Recommendation

The explainable recommendation aims to unveil the rationale behind the recommendation results and provide human-understandable textual explanations. In recommendation scenarios, the user-item interactions can be represented by a bipartite graph $\mathcal{G} = \{(u, i) | u \in \mathcal{U}, i \in I\}$, where \mathcal{U} and I are the sets of users and items, respectively. Each node in \mathcal{G} is associated with textual information from either the user profile set \mathcal{B} or the item profile set C. Specifically, each user u has a profile $b_u \in \mathcal{B}$ describing their preferences, and each item i has a profile $c_i \in C$ detailing attributes and descriptions.

Given a recommendation (u, i), which is a pair of user u and recommended item i, the goal of an explainable recommendation is to generate a clear textual explanation, which is formulated as:

$$explanation(u, i) = generator(u, i, \mathcal{G}, \mathcal{B}, C).$$
(1)

The graph models abundant structural CF information beneficial for explainable recommendations. Given the size of the entire graph, we commonly utilize the subgraph $\mathcal{G}(u, i) \subseteq \mathcal{G}$ as input [63, 67]. This subgraph represents the *L*-hop edge-centered ego-graph around the pair (u, i). It includes nodes that are at most *L* hops away from either *u* or *i*, effectively capturing the local connectivity structure for generating explanations.

2.2 Graph Retrieval-Augmented Generation

Retrieval-Augmented Generation (RAG) has achieved remarkable success in enhancing Large Language Models (LLMs) by retrieving external documents. However, existing RAGs fail to fully exploit the structural information in the graph to augment LLMs. The graph retrieval-augmented generation (GraphRAG), to address this issue, leverages structural information to enable more precise retrieval and facilitate context-aware generation [46]. Given a query q, GraphRAG aims to retrieve structural information from a graph \mathcal{G} and generate responses a conditioned by the retrieved information,

Anon

which can be formulated as:

$$G = G\text{-retriever}(q, \mathcal{G}),$$
(2)
$$P(a|q, \mathcal{G}) = P(a|q, G, \theta),$$
(3)

where
$$\theta$$
 is the parameters of the LLMs, and *G* is the structural information (e.g., nodes, paths, and subgraphs) retrieved by the graph retriever.

3 METHODOLOGY

Our proposed G-Refer is designed to excavate the CF information from the graph to enhance explainable recommendations. First, we leverage a hybrid graph retrieval mechanism that combines multigranularity retrievers to extract both structural and semantic CF information from the user-item interaction graph. In the second part, we employ knowledge pruning to filter out less relevant or redundant training samples, eliminating noise and improving training efficiency. Finally, a lightweight retrieval-augmented fine-tuning is conducted to bridge the modality gap, instructing LLMs to understand retrieved knowledge and generate accurate and contextually relevant explanations.

3.1 Hybrid Graph Retrieval

The user-item interaction graph encapsulates a wealth of knowledge, including user behaviors, preferences, and relationships among users and between items [56]. It can provide sufficient CF information from different perspectives, such as structural and semantic views, to explain the user behaviors, thereby enhancing the explainability of recommendations. For example, as illustrated in Figure 1, we can explain why user u_1 might be interested in item i_2 from both structural and semantic perspectives. From the structural perspective, the path $(u_1 \rightarrow i_1 \rightarrow u_2 \rightarrow i_2)$ illustrates the potentially shared preferences between user u_1 and u_2 , both of whom have watched i_1 , Iron Man, suggests that user u_1 may also enjoy i_2 , Doctor *Strange*, which user u_2 has watched, due to the overlapping Marvel fandom. In addition, from the semantic view, retrieving items i_2 and i3 could also lead to an explanation. Both items feature the same leading actor, Benedict Cumberbatch. This commonality suggests that user u_1 , having shown interest in related actors through their viewing habits, might find *i*₃, *Sherlock*, appealing as well.

According to [46], retrievers based on different levels of granularity in the graph (e.g., nodes, paths, and subgraphs) have unique strengths for addressing various aspects of retrieval scenarios. Inspired by recent works [16, 38], we adopt a hybrid graph retrieval strategy that employs both **path-level** and **node-level** retrievers to respectively retrieve structural and semantic CF information from the user-item interaction graph, which is composed of purchase histories with enriched text attributes. The CF information retrieved from both views is then translated into human-understandable text, facilitating the generation of explanations.

3.1.1Path-level Retriever. The path-level retriever aims to retrieve285the structural CF information from the graphs by identifying the286most k informative paths that connect users and items for inter-287pretable recommendations. Typically, recommendations on graphs288can be formalized as a link prediction (LP) problem, which predicts289the connections between user u and i based on the graph structure



Figure 2: The illustration of path-level retriever.

[33]. This can be formulated as

$$i = f_{\phi}(u, \mathcal{G}), \tag{4}$$

where f_{ϕ} denotes a link prediction model parameterized by ϕ . The f_{ϕ} captures complex patterns between users and items in the graph to provide the recommendation results.

Motivation. GNNs [58] have demonstrated exceptional performance in capturing graph structures and are widely used in recommender systems [20, 57]. While GNNs excel at extracting CF information for recommendations, they are often criticized for their lack of interpretability. GNNs struggle to provide clear explanations for their predictions [64], which significantly undermines transparency. To address this issue, inspired by recent GNN explanation methods [67], we propose a path-level retriever that generates explanations for GNN predictions and reveals the structural CF information captured by GNNs as interpretable paths, which is illustrated in Figure 2.

GNN Training. To capture the structural CF information, we first train a GNN to learn the user-item interactions in the graph for recommendations. We follow a standard two-step pipeline, which involves obtaining user and item representations with an *L*-hop GNN encoder and applying a prediction head to get a recommendation prediction. We use R-GCN [49] as the GNN encoder to learn node embeddings on user-item graphs, and an inner product is adopted as the prediction head. We also explore other GNN variants, i.e., LightGCN [20], in our experiments (Section 4.3). The forward-pass update of each user *u* in R-GCN is formalized as follows:

$$\mathbf{h}_{u}^{(l)} = \text{ReLU}\left(\sum_{i' \in \mathcal{N}_{u}} \frac{1}{|\mathcal{N}_{u}|} \mathbf{W}_{1}^{(l-1)} \mathbf{h}_{i'}^{(l-1)} + \mathbf{W}_{0}^{(l-1)} \mathbf{h}_{u}^{(l-1)}\right), \quad (5)$$

where $\mathbf{h}_{u}^{(l)}$ is the hidden state of user u in the l-th layer and \mathcal{N}_{u} denotes the neighbors of u. The $\mathbf{W}_{0}^{(l-1)}$ and $\mathbf{W}_{1}^{(l-1)}$ are learnable weight matrices at layer l-1. Likewise, the item representation $\mathbf{h}_{i}^{(l)}$ can be obtained in a similar manner.

The *m***-core Pruning.** To ensure concise and accurate explanations, we aim to extract paths that are short in length and avoid containing

high-degree nodes in the graph [67]. Given a prediction (u, i), the corresponding *L*-hop edge-centered ego-graph is initially extracted, which includes nodes that are at most *L* hops away from either *u* or *i*. A *m*-core pruning process is then applied to remove spurious neighbors and improve speed. Specifically, the *m*-core pruning is a recursive algorithm that eliminates nodes with degrees less than *m*, until only those with degrees $\geq m$ remain, forming the *m*-core.

Explanation Path Retrieval. We follow PaGE-Link [67] to perform mask learning on the learned GNN model to identify explanation paths from the user-item interaction graph via learning a mask over all edges, assigning high weights to important edges and low weights to others. Based on the mask, we apply the Dijkstra's shortest path algorithm [13] to retrieve explanation paths.

Formally, let \mathcal{E} denote all user-item interactions, $\tau(e)$ be the type of the edge e, and \mathcal{M} be learnable masks of all edge types. To optimize \mathcal{M} , the objective function \mathcal{L} that needs to be minimized is composed of two loss terms, which can be formulated as follows:

$$\mathcal{L}(\mathcal{M}) = \mathcal{L}_{pred}(\mathcal{M}) + \mathcal{L}_{path}(\mathcal{M}).$$
(6)

The first loss term \mathcal{L}_{pred} is to learn to select crucial edges for model prediction based on the perturbation-based explanation. Given a connected pair (*u*, *i*), the loss can be defined as follows:

$$\mathcal{L}_{pred}(\mathcal{M}) = -\log P(Y = 1 | \mathcal{G} = (\mathcal{U}, I, \mathcal{E} \odot \sigma(\mathcal{M})), (u, i)), \quad (7)$$

where $\sigma(\cdot)$ is the sigmoid function, \odot is the element-wise product, and *Y* is the original GNN prediction. Another loss term \mathcal{L}_{path} is to learn to select path-forming edges. We first use the following score function to assign a score to each candidate path *p* in the graph, prioritizing shorter paths and those without high-degree nodes:

$$Score(p) = \prod_{e \in p} \log \sigma(\mathcal{M}_{\tau(e)}) - \log(\deg(e)), \tag{8}$$

where $\mathcal{M}_{\tau(e)}$ denotes the mask corresponding to the type of the edge *e* and deg(*e*) is the degree of the target node of *e*. Using Dijkstra's shortest path algorithm, where path lengths are redefined according to the score function, a set of candidate edges considered concise and informative is selected, denoted by \mathcal{E}_{path} . Guided by \mathcal{E}_{path} , we define \mathcal{L}_{path} accordingly:

$$\mathcal{L}_{path}(\mathcal{M}) = -\sum_{r \in \mathcal{R}} (\sum_{e \in \mathcal{E}_{path}} \mathcal{M}_{\tau(e)} - \sum_{e \in \mathcal{E}, e \notin \mathcal{E}_{path}} \mathcal{M}_{\tau(e)}).$$
(9)

Upon convergence of mask learning, we can run the Dijkstra's shortest path to generate up to k explanation paths based on \mathcal{M} .

3.1.2 *Node-level Retriever.* The node-level retriever aims to retrieve the semantic CF information from the graphs by selecting the *k* most relevant user nodes and the *k* most relevant item nodes.

Motivation. Graphs are often known as noisy and incomplete, and the structural information alone may not be sufficient to provide accurate explanations. Additionally, graphs in recommendations typically contain rich semantics within user and item nodes de-rived from user interests and item properties. While the profiles provide some information, using only individual profiles for explanations may be inadequate, as it might not fully capture user topics or item features. Thus, retrieving semantics from graphs is crucial to uncovering latent connections and supplementing more comprehensive explanations for recommendations. Our node-level retriever incorporates a dense retrieval to calculate the semantic

similarity between users and items to retrieve the most relevant neighboring nodes for the current recommendation (u, i) to unveil the semantic CF knowledge for explanations.

Dense Retrieval. We employ a dual-encoder-based retriever architecture, which has demonstrated effectiveness across various retrieval tasks [42] and efficiency at the inference stage [24, 27]. Given a user-item pair (u, i), a text encoder maps the user profile b_u to an embedding $f(b_u) \in \mathbb{R}^d$, and the same encoder maps each item profile c_i to an embedding $f(c_i) \in \mathbb{R}^d$, where *d* is the hidden state dimension of the text encoder. The top-*k* relevant users are retrieved based on the user-user semantic similarities, which are computed via cosine similarity:

$$sim(u, u') = \frac{f(b_u) \cdot f(b_{u'})}{\|f(b_u)\| \|f(b_{u'})\|}, \ u' \in \mathcal{N}_i,$$
(10)

where N_i is the set of users connected to item *i*. We also retrieve top*k* items based on the item-item semantic similarities. The retriever is training-free and a pre-trained language model specifically utilized to generate high-quality sentence embeddings.

3.1.3 Graph Translation. After the hybrid retrieval, the complex nature of graph-type knowledge, particularly the paths retrieved, presents a challenge since it cannot be directly integrated with profiles of users and items for input into the LLMs [6, 53]. To address this, it is necessary to employ graph translation techniques that convert the graph-type knowledge into a format compatible with LLMs, enabling LLMs to effectively process and utilize structured information. Considering that instructions are typically presented in natural language, we follow [52, 62] to adopt a flatten-based method, which transforms retrieved paths and nodes into descriptive, easily comprehensible language. For each user-item pair (u, i) requiring an explanation, along with k retrieved paths, k user nodes, and k item nodes, we design the instruction prompt as follows:

Given the item title, item profile, and user profile, please explain why the user would enjoy this item. Item title: [Item title]. Item profile: [Item summary]. User profile: [User summary]. For the user-item pair, here are some related users and items. Users: [Top-k similar user profiles]. Items: [Top-k similar item profiles]. For the given user-item pair, here are several related paths connecting users and items through their interactions. [Top-k explanation paths] (each formatted as "<User profile> -> buys -> <Item profile> -> buys the profile> -> user profile> -> use

3.1.4 Discussion of Retrieval Granularity. In G-Refer, we consider node-level and path-level retrieval rather than subgraph-level based on the following observations. Firstly, both nodes and paths have a considerably smaller search space than subgraphs. As proven in [67], compared to the expected number of edge-induced subgraphs, the expected number of candidate nodes and paths grows strictly slower and becomes negligible. Therefore, such explanations exclude many less-meaningful subgraph candidates, making the explanation generation much more straightforward and accurate. Secondly, subgraphs contain complex high-order information, making it difficult to describe a retrieved subgraph with hundreds or even thousands of nodes in a way that is easily understandable, especially when all nodes have text attributes [36]. Consequently, we implement hybrid graph retrieval in the path and node levels.

G-Refer: Graph Retrieval-Augmented Large Language Model for Explainable Recommendation

Conference acronym 'XX, June 03-05, 2018, Woodstock, NY

3.2 Knowledge Pruning

Motivation. After graph retrieval, our goal is to enable LLMs to effectively generate explanations based on the graph CF information of varying granularities retrieved. However, it is noticed that for some user-item pairs, a sufficient explanation can be derived solely from their profiles, without the need for additional CF information. For instance, if a user's profile shows a preference for science fiction and the recommended item is a popular sci-fi film, then explaining this interaction might only require referencing these profile details. In such cases, the retrieved knowledge is not only redundant but employing such training samples could also weaken the LLM's ability to utilize CF information.

Re-ranking. To enhance LLMs' focus on CF information, we pro-pose knowledge pruning to filter out less relevant or redundant training samples [46], building on the inherent ability of LLMs to generate explanations by understanding user and item profiles. It also improves training efficiency through reduced data volume, particularly beneficial for large graphs as indicated in Table 1. To be specific, we use a re-ranking method that sorts all training samples based on their reliance on CF information for explanation, priori-tizing those that require more additional knowledge. We measure such reliance via the semantic similarity between the profiles and the ground truth explanations, which is defined as follows:

$$sim((u,i), \text{Explain}_{(u,i)}) = \frac{f(b_u \oplus c_i) \cdot f(\text{Explain}_{(u,i)})}{\|f(b_u \oplus c_i)\|\|f(\text{Explain}_{(u,i)})\|}, \quad (11)$$

where \oplus represents the concatenation and $\operatorname{Explain}_{(u,i)}$ is the ground truth explanation for (u, i). We introduce a pruning ratio t, which determines the portion of the total training data \mathcal{D} to be filtered after re-ranking. The top " $(1 - t) \cdot |\mathcal{D}|$ " training samples are preserved, ensuring that samples that are most likely to benefit from CF-based explanations are prioritized, thereby improving the overall efficiency and effectiveness of the training process.

3.3 Retrieval-Augmented Fine-Tuning

Motivation. To improve the LLM's capacity for generating bet-ter explanations using retrieved CF information, we fine-tune it with a parameter-efficient pre-training strategy, i.e., LoRA [23], on the pruned training set with in-context retrieval augmenta-tion. Integrating retrieval results during fine-tuning offers two key advantages. (1) It adapts the LLM to better utilize retrieved CF infor-mation to generate explanation, especially for the requirement of domain-specific knowledge it has never seen before. (2) Even state-of-the-art retrievers can falter and return inaccurate results [35]. By training the LLM to generate ground-truth responses even when irrelevant CF information is given, we enable the LLM to ignore misleading retrieval content and lean into its internal knowledge to reduce hallucination.

Fine-tuning. We employ retrieval-augmented fine-tuning (RAFT) with a conventional language modeling loss [69]. The model is trained to process input consisting of a user-item pair (u, i) with profiles b_u and c_i , its associated retrieved knowledge $\mathcal{K}_{(u,i)}$, and a corresponding prompt question Q to generate textual explanations E as outputs. The RAFT loss is computed based on the discrepancy between the model's predicted explanations and the ground truth

Table 1: Statistics of the experimental datasets.

Dataset	# Users	# Items	# Interactions	# Train (u-i)	# Test (u-i)
Amazon-books	15,349	15,247	360,839	95,841	3,000
Yelp	15,942	14,085	393,680	74,212	3,000
Google-reviews	22,582	16,557	411,840	94,663	3,000

explanations in the dataset. Formally, the loss is defined as follows:

$$\mathcal{L}_{\text{RAFT}} = -\sum_{(u,i)\in\mathcal{D}_{\text{prune}}} \log P(\text{Explain}_{(u,i)}|b_u, c_i, \mathcal{K}_{(u,i)}, Q; \theta),$$
(12)

where D_{prune} represents the pruned training set and θ is the parameters associated with the LoRA model.

4 EXPERIMENTS

We evaluate our model on real-world datasets to assess its performance in enhancing explainable recommendations. In particular, we aim to answer the following research questions: **Q1**: How effective is G-Refer compared with state-of-the-art models? **Q2**: How do the main components of our model impact the performance? **Q3**: What is the impact of hyper-parameters? **Q4**: How does G-Refer's training efficiency compare to state-of-the-art methods?

4.1 Experimental Protocols

4.1.1 Datasets. We evaluate G-Refer and leverage three prominent public datasets that offer distinct perspectives on user-item interactions, including **Amazon-books** [41], **Yelp** [39], and **Google-reviews** [28, 60]. Table 1 lists the statistics of three datasets. More details of datasets can be found in Appendix A.1.

4.1.2 Metrics. We follow XRec [39] to utilize a suite of metrics aimed at assessing the semantic explainability and stability of the generated explanations. Traditional n-gram-based metrics like BLEU [44] and ROUGE [34] are not adequate for this purpose due to their inability to fully capture semantic meaning. Specifically, we use GPT_{score} [54], $BERT_{score}$ [68], $BART_{score}$ [66], BLEURT [50], and USR [30] to measure the explainability. Notably, $BERT_{score}$ comprises the Recall score, which measures the completeness and quality of the retrieved CF information, allowing us to evaluate the performance of our graph retriever. More details of the used metrics can be found in Appendix A.2. To evaluate quality consistency, we also report the standard deviations of these metrics. Lower standard deviation values indicate more stable performance.

4.1.3 Baselines. We introduce five state-of-the-art baselines, including NRT [32], Att2Seq [14], PETER [30], PEPLER [31], and XRec [39]. These models are based on representative language models, such as GRU [10], LSTM [19], GPT, and LLaMA. More details of the compared baselines can be found in Appendix A.3.

4.1.4 *Implementation Details.* For path-level retrieval, we set the node embedding dimension to 128, the maximum retrieved path length to 5, and *m* is set to 2 for m-core pruning. We initialize the encoders in both the node-level retriever and the re-ranking mechanism with SentenceBERT [47]. The number of retrieved paths, as well as the number of retrieved nodes (including users and items), is set to 2, and the pruning ratio *t* is set to 70% across all datasets.

Anon

Table 2: Overall comparison in terms of Explainability and Stability. Superscripts "P", "R", and "F1" denote Precision, Recall, and F1-Score, respectively. The subscript "std" indicates the standard deviation of each metric. Bold indicates the best results, while underlined denotes the second-best. "7B" and "8B" denote LLaMA 2-7B and LLaMA 3-8B, respectively.

Models	Explainability ↑						Stability ↓						
wiouels	GPT _{score}	BERT ^P _{score}	BERT ^R _{score}	BERT ^{F1} score	BART _{score}	BLEURT	USR	GPT _{std}	BERT ^P std	BERT ^R std	BERT ^{F1} std	BART _{std}	BLEURTstd
					Amazo	on-books							
NRT	75.63	0.3444	0.3440	0.3443	-3.9806	-0.4073	0.5413	12.82	0.1804	0.1035	0.1321	0.5101	0.3104
Att2Seq	76.08	0.3746	0.3624	0.3687	-3.9440	-0.3302	0.7757	12.56	0.1691	0.1051	0.1275	0.5080	0.2990
PETER	77.65	0.4279	0.3799	0.4043	-3.8968	-0.2937	0.8480	11.21	0.1334	0.1035	0.1098	0.5144	0.2667
PEPLER	78.77	0.3506	0.3569	0.3543	-3.9142	-0.2950	0.9563	11.38	0.1105	0.0935	0.0893	0.5064	0.2195
XRec	82.57	0.4193	0.4038	0.4122	-3.8035	-0.1061	1.0000	9.60	0.0836	0.0920	0.0800	0.4832	0.1780
G-Refer (7B)	82.70	0.4076	0.4476	0.4282	-2.4772	-0.1246	1.0000	9.04	0.0937	0.0845	0.0820	0.2769	0.1893
G-Refer (8B)	82.82	0.4073	0.4494 (+4.56%)	0.4289 (+1.67%)	-2.4699	-0.1203	1.0000	8.95	0.0945	0.0855	0.0825	0.2791	0.1912
					J	/elp							
NRT	61.94	0.0795	0.2225	0.1495	-4.6142	-0.7913	0.2677	16.81	0.2293	0.1134	0.1581	0.5612	0.2728
Att2Seq	63.91	0.2099	0.2658	0.2379	-4.5316	-0.6707	0.7583	15.62	0.1583	0.1074	0.1147	0.5616	0.2470
PETER	67.00	0.2102	0.2983	0.2513	-4.4100	-0.5816	0.8750	15.57	0.3315	0.1298	0.2230	0.5800	0.3555
PEPLER	67.54	0.2920	0.3183	0.3052	-4.4563	-0.3354	0.9143	14.18	0.1476	0.1044	0.1050	0.5777	0.2524
XRec	74.53	0.3946	0.3506	0.3730	-4.3911	-0.2287	1.0000	11.45	0.0969	0.1048	0.0852	0.5770	0.2322
G-Refer (7B)	74.91	0.3573	0.4264	0.3922	-2.8716	-0.1451	1.0000	10.88	0.1050	0.0952	0.0862	0.3571	0.2197
G-Refer (8B)	75.16	0.3629	0.4373 (+8.67%)	0.4003 (+2.73%)	-2.8717	-0.1335	<u>1.0000</u>	10.76	0.1068	0.0995	0.0885	0.3534	0.2182
					Google	e-reviews							
NRT	58.27	0.3509	0.3495	0.3496	-4.2915	-0.4838	0.2533	19.16	0.2176	0.1267	0.1571	0.6620	0.3118
Att2Seq	61.31	0.3619	0.3653	0.3636	-4.2627	-0.4671	0.5070	17.47	0.1855	0.1247	0.1403	0.6663	0.3198
PETER	65.16	0.3892	0.3905	0.3881	-4.1527	-0.3375	0.4757	17.00	0.2819	0.1356	0.2005	0.6701	0.3272
PEPLER	61.58	0.3373	0.3711	0.3546	-4.1744	-0.2892	0.8660	17.17	0.1134	0.1161	0.0999	0.6752	0.2484
XRec	69.12	0.4546	0.4069	0.4311	-4.1647	-0.2437	0.9993	14.24	0.0972	0.1163	0.0938	0.6591	0.2452
G-Refer (7B)	71.47	0.4253	0.4873	0.4566	-2.6013	-0.1561	1.0000	13.46	0.1184	0.0872	0.0921	0.4004	0.2415
G-Refer (8B)	71.73	0.4245	0.4935 (+7.48%)	0.4592 (+2.81%)	-2.6117	-0.1517	1.0000	13.23	0.1175	0.0920	0.0916	0.3921	0.2511



Figure 3: Human evaluation comparing XRec and G-Refer.

For retrieval-augmented fine-tuning, we use the models from the open-source LLaMA family. Specifically, for a fair comparison, we utilize the same model, i.e., LLaMA 2-7B adopted by baselines [39]. We also report results based on the advanced LLaMA 3-8B model. We set the learning rate, epochs, and max length as 2e-5, 2, and 2048 for RAFT, which is trained on 8 NVIDIA A100 GPUs. The total batch sizes are set to 32 and 16, respectively, for the 7B and 8B models. The rank of the LoRA adapter is set to 8. For inference, we set the temperature as 0 and the maximum output tokens as 256, ensuring a stable and reasonable generation. In addition, we employ the GPT-3.5-turbo model for computing the GPTScore metric.

4.2 Model Performance (RQ1)

Overall Performance. We first compare the quality of gener-4.2.1 ated explanations against baselines across three datasets, the results are summarized in Table 2. It is observed that G-Refer demonstrates superior performance in both explainability and stability, outper-forming baselines across semantic evaluators such as GPT, BERT, and BART. Compared to the most powerful baseline, XRec [39],

which uses GNNs to implicitly capture CF information, our model shows improvements in $\mathrm{BERT}^{\mathrm{F1}}_{\mathrm{score}}$ across three datasets, with increases of 1.67%, 2.73%, and 2.81% respectively. This indicates that the explicit CF information retrieved by our hybrid graph retriever is more accurate and better utilizes both structures and semantics of user-item interaction graphs compared to the implicit CF information captured by XRec. In addition, graph translation enables the adapter-free RAFT, allowing LLMs to better comprehend the structural inputs and leverage their powerful capabilities to generate human-understandable explanations. Notably, G-Refer significantly outperforms all baselines in BERT^R_{score}, with increases of 4.56%, 8.67%, and 7.48% across the three datasets, while slightly decreasing in BERT^P_{score}. The retrieved CF information enables generated explanations to include more key information (e.g., user topic modeling and interaction history). Although introducing retrieved knowledge inevitably introduces some noise, reducing the precision slightly, we consider this trade-off beneficial, as the completeness of explanations is crucial for user understanding compared to accurate but less informative expressions. Finally, we notice that G-Refer's performance improvement on Amazon is relatively modest compared to other datasets. This can be attributed to the sparsity of its user-item interaction graph, particularly in the test set where the average node degree is only 2.76. In such cases, even with retrieval, the effectiveness of CF information from the graph is limited.

4.2.2 Human Evaluation. The ultimate goal of model explanation is to aid human understanding and decision-making. Human evaluation is thus the best way to evaluate the effectiveness of an explainer, which has been widely used in previous works [18, 48]. We conduct a human evaluation by randomly selecting 20 user-item

Table 3: Ablation study for G-Refer, with	the best results
highlighted red and the worst in blue for th	e component.

698

699

700

701

702

703

704

705

706

707

708

709

710

711

713

714

715

716

718

719

720

721

722

723

724

725

726

727

728

729

730

731

732

733

734

735

736

737

738

739

740

741

742

743

744

746

747

748

749

750

751

754

Datasets	Ye	elp	Google-reviews		
Ablations	$BERT^{F1}\uparrow$	$ $ BERT _{std} \downarrow	$BERT^{F1}\uparrow$	BERT _{std} ↓	
Variants on graph	ı retriever:				
w/o path-level	0.3927	0.0868	0.4560	0.0924	
w/o node-level	0.3966	0.0894	0.4544	0.0922	
w/o GraphRAG	0.3880	0.0896	0.4468	0.0940	
Variants on link p	prediction mod	lel:			
w/ LightGCN	0.3941	0.0870	0.4589	0.0922	
w/ R-GCN	0.4003	0.0885	0.4592	0.0916	
Variants on LLMs with different scales:					
w/ Qwen-0.5B	0.3201	0.6530	0.4129	0.2171	
w/ Qwen-1.5B	0.3557	0.3940	0.4451	0.0940	
w/ Qwen-3B	0.3994	0.0861	0.4602	0.0903	
w/ Qwen-7B	0.3991	0.0851	0.4582	0.0914	
Knowledge pruni	ng v.s. full trai	ning set:			
w/o pruning	0.4002	0.0892	0.4605	0.0909	
G-Refer	0.4003	0.0885	0.4592	0.0916	

pairs from the test set of each dataset and generating explanations for each sample using XRec [39] and G-Refer. We designed a survey with single-choice questions and distributed it to five senior researchers, asking them to select the best explanations. As shown in Figure 3, explanations generated by G-Refer were consistently favored across all datasets, especially for Yelp and Google-reviews, where our explanations were chosen in over 80% of cases.

4.2.3 Case Study. To demonstrate the effectiveness of G-Refer and show how retrieved CF signals benefit the generated explanations, we provide several cases and give analysis in Appendix B.

Ablation Study (RQ2) 4.3

To investigate how each component affects the model performance, we conduct the following ablation studies in Table 3.

4.3.1 The Effects of Hybrid Graph Retriever. From Table 3 we can observe that both path-level and node-level retrievers contribute to the final results, with their combination yielding the best performance. We observe that semantic information is more crucial for Yelp, while structural knowledge offers greater advantages for Google-reviews. This suggests that the importance of retrievers may differ based on dataset characteristics and highlights the complementarity of the two retrievers.

4.3.2 The Effects of Various GNNs. The GNN encoder in our pathretriever captures structural CF information and can be replaced with any advanced GNN-based recommender model. We compare LightGCN [20], which is one of the representative GNN models for 745 recommendation, with R-GCN [49] in the ablation study. The results show that both of the GNNs achieve competitive performance. This demonstrates their effectiveness in capturing structural information for recommendation and the flexibility of our path-level retriever, which could adapt to different graphs and recommendation scenarios using different GNN architectures.

4.3.3 The Effects of Various LLMs with Different Scales. We in-752 753 vestigate the impact of LLM scales on G-Refer's performance by



Figure 4: Performance of different retrieved number k.

comparing the advanced Owen 2.5 family [61] of different scales. Firstly, we observe a clear performance scaling from 0.5B to 3B models. Interestingly, the 3B model achieves comparable results to the 7B model, indicating that with the introduction of retrieved knowledge, even a relatively smaller LLM can achieve impressive results through RAFT. This also demonstrates the necessity of finetuning with retrieved knowledge, which can significantly enhance the model's ability to generate high-quality explanations.

4.3.4 The Effects of Knowledge Pruning. Training with the full dataset increases the number of samples by several times but does not improve performance proportionally, and may even lead to slight degradation (e.g., on Yelp). This is because some user-items are self-explanatory, and additional CF information may introduce noise. Thus, knowledge pruning is essential for removing these samples from training and improving model performance.

4.4 Hyperparameter Study (RQ3)

We explore the variation of G-Refer's performance with respect to two hyperparameters: k (the number of retrieved elements, including both paths and nodes) and t (the knowledge pruning ratio).

4.4.1 Retrieved Number k. Figure 4 illustrates the impact of k (ranging from 1 to 5) on BERT_{score} precision and BERT_{score} recall (abbr. precision and recall in the following). The best precision is achieved at k=2, after which it decreases, showing that more retrieved paths and nodes can introduce noise and deteriorate the performance. Conversely, a low k (i.e., k=1) fails to provide sufficient CF information, leading to lower results. We also notice that recall is not sensitive to the increased k, suggesting that a modest k can effectively retrieve essential information for generating explanations.

4.4.2 Pruning Ratio t. In Table 4, we report G-Refer's explainability and stability on three datasets with *t* ranging from 90% to 50%. In terms of explainability, both Amazon-books and Yelp perform best at t = 70%, while Google-reviews at t = 50%. This suggests that a low t value, which corresponds to retaining more training samples, does not always lead to improved model performance, as irrelevant samples can introduce noise. However, it generally enhances model stability, indicating that fine-tuning with more data can enhance the model's ability to generate consistent explanations.

4.5 Efficiency Analysis (RQ4)

In Figure 5, we analyze the efficiency of our model, including the time taken to train one epoch, performance results, and the number

833

834

835

836

837

838

839

840

841

842

843

844

845

870

Table 4: Performance of different knowledge pruning ratio t.

Ratio	Ex	plainabili	ity↑		Stability	\downarrow
itutio	BERT ^{F1}	BART	BLEURT	BERT _{std}	BART _{std}	BLEURT _{std}
			Amazon-	books		
t=90%	0.4211	-2.5045	-0.1388	0.0817	0.2811	0.1905
t=80%	0.4253	-2.5051	-0.1339	0.0821	0.2803	0.1913
t=70%	0.4289	-2.4799	-0.1203	0.0825	0.2791	0.1912
t=60%	0.4279	-2.4743	-0.1263	0.0830	0.2799	0.1910
t=50%	0.4255	-2.4900	-0.1326	0.0819	0.2744	0.1893
			Yelj	þ		
t=90%	0.3974	-2.9125	-0.1455	0.0872	0.3724	0.2240
t=80%	0.3961	-2.8893	-0.1458	0.0866	0.3591	0.2207
t=70%	0.4003	-2.8717	-0.1335	0.0885	0.3534	0.2182
t=60%	0.3927	-2.8859	-0.1461	0.0856	0.3559	0.2178
t=50%	0.3967	-2.8934	-0.1461	0.0856	0.3518	0.2164
			Google-re	eviews		
t=90%	0.4462	-2.6987	-0.1798	0.0919	0.4042	0.2437
t=80%	0.4518	-2.6202	-0.1645	0.0927	0.3912	0.2508
t=70%	0.4592	-2.6117	-0.1517	0.0916	0.3921	0.2511
t=60%	0.4583	-2.6169	-0.1465	0.0918	0.3908	0.2477
t=50%	0.4606	-2.6073	-0.1459	0.0907	0.3852	0.2498

of parameters that require tuning for G-Refer, XRec [39], and fullset training. For a fair comparison, all experiments were conducted on a single GPU with a batch size of 1. Though requiring more learnable parameters, G-Refer achieves faster training and better performance compared to XRec. This can be attributed to (1) our knowledge pruning which significantly reduces training data; (2) XRec's speed limitation due to adapter insertion. In addition, full-set training reduces efficiency without performance gains, indicating the necessity of knowledge pruning.

RELATED WORKS 5

5.1 **Explainable Recommendation**

Explainable recommendation [45, 70, 71] has demonstrated signifi-846 847 cant advantages in informing users about the logic behind recom-848 mendations, thereby increasing system transparency, effectiveness, and trustworthiness. Early works focus on generating explanations 849 with predefined templates [29] or extracting logic reasoning rules 850 from recommendation models [5, 51, 75]. To provide more detailed 851 and personalized explanations, recent works have explored gener-852 ating explanations from graph structure, which contains abundant 853 collaborative information (CF) for providing explanations [1, 15, 59]. 854 855 For example, PGPR [59] proposes a reinforcement learning-based method to find a path in the graph to explain the recommendation. 856 857 However, the graph structure is often complex and hard for users to 858 understand. To address this issue, some works have explored generating explanations in natural language [4, 8, 9, 11, 14, 30-32, 32, 39]. 859 With the advance of large language models (LLMs), like ChatGPT 860 861 [43], they can generate more fluent and informative explanations based on user's and item's profile [31]. Recently, researchers have 862 tried to combine the advantages of graphs to enhance the explana-863 tion generated by LLMs. XRec [39] adopts a graph neural network 864 (GNN) to model the graph structure and generate embeddings. 865 Then, the embeddings are fed into LLMs to generate explanations. 866 This approach allows LLMs to produce more informative explana-867 868 tions by considering CF information within the graph structure. 869 However, XRec represents CF information as hidden embeddings,



Figure 5: Efficiency Analysis of G-Refer.

leaving it unclear what specific CF information is considered when generating explanations. This ambiguity makes it difficult to verify the explanations generated by LLMs, which suffer from severe hallucination issues [72] and further diminishes the trustworthiness of the explainable recommendation.

5.2 Graph Retrieval-Augmented Generation

Although retrieval-augmented generation (RAG) has been widely used to enhance the LLMs by incorporating external knowledge [17], it has shown limitations in considering the graph structure. To address this issue, graph retrieval-augmented generation (GraphRAG) has been proposed to effectively incorporate graph information into the LLM generation process [46]. The GraphRAG typically involves two processes: graph retrieval and graph-augmented generation. In graph retrieval, we need to design a retrieval model to retrieve relevant graph information for the given query, which can be some simple non-parametric retrievers like L-hop neighbors. To better consider the graph structure, GNN-based retrievers have been proposed. For example, GNN-RAG [40] first encodes the graph, assigns a score to each entity, and retrieves entities relevant to the query. Instead of GNNs, RoG [38] proposes an LLM-based graph retriever by planning, retrieving, and reasoning on graphs. In the graph-augmented generation, the retrieved graph information is used to enhance the generation process of LLMs. Due to the unstructured nature of the LLMs, some works introduce an adapter layer to bridge the modality gap. GraphLLM [3] first encodes the graph as embeddings and then feeds them into LLMs with a graph adapter layer. However, the hidden embeddings used by these methods are difficult to interpret and understand. To improve the interpretability, recent research adopts some graph description language to translate the graph structure into natural language, such as edge list, node list, and syntax tree [25, 73]. In this way, the graph structure can be understood by both LLMs and humans. Thus, we adopt a simple prompt to describe the graph structure for the explainable recommendation.

6 CONCLUSION

We systematically analyze the limitations of existing explainable recommendation methods and introduce G-Refer, a novel model to address these challenges. The key design of our model involves leveraging a hybrid graph retrieval to extract explicit CF signals, employing knowledge pruning to filter out less relevant samples, and utilizing retrieval-augmented fine-tuning to integrate retrieved knowledge into explanation generation. Comprehensive experiments validate the effectiveness of our model. G-Refer reveals great potential for graph retrieval-augmented generation in recommendation scenarios. Future work includes exploring fully training-free retrievers and investigating the transferability to other tasks.

Anon

G-Refer: Graph Retrieval-Augmented Large Language Model for Explainable Recommendation

Conference acronym 'XX, June 03-05, 2018, Woodstock, NY

987

988

989

990

991

992

993

994

995

996

997

998

999

1000

1001

1002

1003

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

1018

1019

1020

1021

1022

1023

1024

1025

1026

1027

1028

1029

1030

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043 1044

929 **REFERENCES**

930

931

932

933

934

935

936

937

938

939

940

941

942

943

944

945

946

947

948

949

950

951

952

953

954

955

956

957

958

959

960

961

962

963

964

965

966

967

968

969

970

971

972

973

974

975

976

977

978

979

980

981

982

983

986

- Giacomo Balloccu, Ludovico Boratto, Christian Cancedda, Gianni Fenu, and Mirko Marras. 2023. Faithful Path Language Modelling for Explainable Recommendation over Knowledge Graph. arXiv preprint arXiv:2310.16452 (2023).
- [2] Yixin Cao, Xiang Wang, Xiangnan He, Zikun Hu, and Tat-Seng Chua. 2019. Unifying knowledge graph learning and recommendation: Towards a better understanding of user preferences. In *The world wide web conference*. 151–161.
- [3] Ziwei Chai, Tianjie Zhang, Liang Wu, Kaiqiao Han, Xiaohai Hu, Xuanwen Huang, and Yang Yang. 2023. Graphllm: Boosting graph reasoning ability of large language model. arXiv preprint arXiv:2310.05845 (2023).
- [4] Hanxiong Chen, Xu Chen, Shaoyun Shi, and Yongfeng Zhang. 2021. Generate natural language explanations for recommendation. arXiv preprint arXiv:2101.03392 (2021).
- [5] Hanxiong Chen, Shaoyun Shi, Yunqi Li, and Yongfeng Zhang. 2021. Neural collaborative reasoning. In Proceedings of the Web Conference 2021. 1516–1527.
- [6] Nuo Chen, Yuhan Li, Jianheng Tang, and Jia Li. 2024. GraphWiz: An Instruction-Following Language Model for Graph Computational Problems. In Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. 353–364.
- [7] Xu Chen, Yongfeng Zhang, and Ji-Rong Wen. 2022. Measuring" why" in recommender systems: A comprehensive survey on the evaluation of explainable recommendation. arXiv preprint arXiv:2202.06466 (2022).
- [8] Zhongxia Chen, Xiting Wang, Xing Xie, Tong Wu, Guoqing Bu, Yining Wang, and Enhong Chen. 2019. Co-attentive multi-task learning for explainable recommendation.. In *IJCAI*, Vol. 2019. 2137–2143.
- [9] Hao Cheng, Shuo Wang, Wensheng Lu, Wei Zhang, Mingyang Zhou, Kezhong Lu, and Hao Liao. 2023. Explainable recommendation with personalized review retrieval and aspect learning. arXiv preprint arXiv:2306.12657 (2023).
- [10] Kyunghyun Cho. 2014. Learning phrase representations using RNN encoderdecoder for statistical machine translation. arXiv preprint arXiv:1406.1078 (2014).
- [11] Felipe Costa, Sixun Ouyang, Peter Dolog, and Aonghus Lawlor. 2018. Automatic generation of natural language explanations. In Companion Proceedings of the 23rd International Conference on Intelligent User Interfaces. 1–2.
- [12] Jacob Devlin. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018).
- [13] Edsger W Dijkstra. 2022. A note on two problems in connexion with graphs. In Edsger Wybe Dijkstra: his life, work, and legacy. 287-290.
- [14] Li Dong, Shaohan Huang, Furu Wei, Mirella Lapata, Ming Zhou, and Ke Xu. 2017. Learning to generate product reviews from attributes. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers. 623–632.
- [15] Zuohui Fu, Yikun Xian, Ruoyuan Gao, Jieyu Zhao, Qiaoying Huang, Yingqiang Ge, Shuyuan Xu, Shijie Geng, Chirag Shah, Yongfeng Zhang, et al. 2020. Fairnessaware explainable recommendation over knowledge graphs. In Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval. 69–78.
- [16] Yifu Gao, Linbo Qiao, Zhigang Kan, Zhihua Wen, Yongquan He, and Dongsheng Li. 2024. Two-stage Generative Question Answering on Temporal Knowledge Graph Using Large Language Models. arXiv preprint arXiv:2402.16568 (2024).
- [17] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. arXiv preprint arXiv:2312.10997 (2023).
- [18] Azin Ghazimatin, Oana Balalau, Rishiraj Saha Roy, and Gerhard Weikum. 2020. Prince: Provider-side interpretability with counterfactual explanations in recommender systems. In Proceedings of the 13th International Conference on Web Search and Data Mining. 196–204.
- [19] Alex Graves and Alex Graves. 2012. Long short-term memory. Supervised sequence labelling with recurrent neural networks (2012), 37–45.
- [20] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgen: 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.
- [21] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In Proceedings of the 26th international conference on world wide web. 173–182.
- [22] Jonathan L Herlocker, Joseph A Konstan, and John Riedl. 2000. Explaining collaborative filtering recommendations. In Proceedings of the 2000 ACM conference on Computer supported cooperative work. 241–250.
- [23] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2021. LoRA: Low-Rank Adaptation of Large Language Models. In International Conference on Learning Representations.
- [24] Gautier Izacard, Mathilde Caron, Lucas Hosseini, Sebastian Riedel, Piotr Bojanowski, Armand Joulin, and Edouard Grave. 2022. Unsupervised Dense Information Retrieval with Contrastive Learning. *Transactions on Machine Learning Research* (2022).
- [25] Jiho Kim, Yeonsu Kwon, Yohan Jo, and Edward Choi. 2023. Kg-gpt: A general framework for reasoning on knowledge graphs using large language models.

arXiv preprint arXiv:2310.11220 (2023).

- [26] M Lewis. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. arXiv preprint arXiv:1910.13461 (2019).
- [27] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. Advances in Neural Information Processing Systems 33 (2020), 9459–9474.
- [28] Jiacheng Li, Jingbo Shang, and Julian McAuley. 2022. Uctopic: Unsupervised contrastive learning for phrase representations and topic mining. arXiv preprint arXiv:2202.13469 (2022).
- [29] Lei Li, Li Chen, and Ruihai Dong. 2021. Caesar: context-aware explanation based on supervised attention for service recommendations. *Journal of Intelligent Information Systems* 57, 1 (2021), 147–170.
- [30] Lei Li, Yongfeng Zhang, and Li Chen. 2021. Personalized Transformer for Explainable Recommendation. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers). 4947–4957.
- [31] Lei Li, Yongfeng Zhang, and Li Chen. 2023. Personalized prompt learning for explainable recommendation. ACM Transactions on Information Systems 41, 4 (2023), 1–26.
- [32] Piji Li, Zihao Wang, Zhaochun Ren, Lidong Bing, and Wai Lam. 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.
- [33] Xin Li and Hsinchun Chen. 2013. Recommendation as link prediction in bipartite graphs: A graph kernel-based machine learning approach. *Decision Support Systems* 54, 2 (2013), 880–890.
- [34] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In Text summarization branches out. 74–81.
- [35] Xi Victoria Lin, Xilun Chen, Mingda Chen, Weijia Shi, Maria Lomeli, Richard James, Pedro Rodriguez, Jacob Kahn, Gergely Szilvasy, Mike Lewis, et al. 2024. RA-DIT: Retrieval-Augmented Dual Instruction Tuning. In *The Twelfth International Conference on Learning Representations*.
- [36] Hao Liu, Jiarui Feng, Lecheng Kong, Ningyue Liang, Dacheng Tao, Yixin Chen, and Muhan Zhang. 2024. One For All: Towards Training One Graph Model For All Classification Tasks. In The Twelfth International Conference on Learning Representations.
- [37] Jie Lu, Dianshuang Wu, Mingsong Mao, Wei Wang, and Guangquan Zhang. 2015. Recommender system application developments: a survey. *Decision support* systems 74 (2015), 12–32.
- [38] Linhao Luo, Yuan-Fang Li, Reza Haf, and Shirui Pan. [n.d.]. Reasoning on Graphs: Faithful and Interpretable Large Language Model Reasoning. In *The Twelfth International Conference on Learning Representations*.
- [39] Qiyao Ma, Xubin Ren, and Chao Huang. 2024. XRec: Large Language Models for Explainable Recommendation. arXiv preprint arXiv:2406.02377 (2024).
- [40] Costas Mavromatis and George Karypis. 2024. GNN-RAG: Graph Neural Retrieval for Large Language Model Reasoning. arXiv preprint arXiv:2405.20139 (2024).
- [41] Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP). 188–197.
- [42] Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernandez Abrego, Ji Ma, Vincent Zhao, Yi Luan, Keith Hall, Ming-Wei Chang, et al. 2022. Large Dual Encoders Are Generalizable Retrievers. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing. 9844–9855.
- [43] OpenAI. 2022. Introducing ChatGPT. https://openai.com/index/chatgpt/
- [44] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 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.
- [45] Georgina Peake and Jun Wang. 2018. Explanation mining: Post hoc interpretability of latent factor models for recommendation systems. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. 2060–2069.
- [46] Boci Peng, Yun Zhu, Yongchao Liu, Xiaohe Bo, Haizhou Shi, Chuntao Hong, Yan Zhang, and Siliang Tang. 2024. Graph Retrieval-Augmented Generation: A Survey. arXiv preprint arXiv:2408.08921 (2024).
- [47] Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. arXiv preprint arXiv:1908.10084 (2019).
- [48] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. "Why should i trust you?" Explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. 1135–1144.
- [49] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. 2018. Modeling relational data with graph convolutional networks. In *The semantic web: 15th international conference, ESWC 2018, Heraklion, Crete, Greece, June 3–7, 2018, proceedings 15.* Springer, 593–607.

1046

1047

1048

1049

1050

1051

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

1080

1081

1082

1083

1084

1085

1086

1087

1088

1089

1090

1091

1092

1093

1094

1095

1096

1102

- [50] Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. BLEURT: Learning Robust Metrics for Text Generation. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics. 7881–7892.
- [51] Shaoyun Shi, Hanxiong Chen, Weizhi Ma, Jiaxin Mao, Min Zhang, and Yongfeng Zhang. 2020. Neural logic reasoning. In Proceedings of the 29th ACM International Conference on Information & Knowledge Management. 1365–1374.
- [52] Yanchao Tan, Zihao Zhou, Hang Lv, Weiming Liu, and Carl Yang. 2024. Walklm: A uniform language model fine-tuning framework for attributed graph embedding. Advances in Neural Information Processing Systems 36 (2024).
- [53] Heng Wang, Shangbin Feng, Tianxing He, Zhaoxuan Tan, Xiaochuang Han, and Yulia Tsvetkov. 2024. Can language models solve graph problems in natural language? Advances in Neural Information Processing Systems 36 (2024).
 - [54] Jiaan Wang, Yunlong Liang, Fandong Meng, Zengkui Sun, Haoxiang Shi, Zhixu Li, Jinan Xu, Jianfeng Qu, and Jie Zhou. 2023. Is ChatGPT a Good NLG Evaluator? A Preliminary Study. In Proceedings of the 4th New Frontiers in Summarization Workshop. 1–11.
 - [55] Xiang Wang, Xiangnan He, Fuli Feng, Liqiang Nie, and Tat-Seng Chua. 2018. Tem: Tree-enhanced embedding model for explainable recommendation. In Proceedings of the 2018 world wide web conference. 1543–1552.
 - [56] Likang Wu, Zhi Zheng, Zhaopeng Qiu, Hao Wang, Hongchao Gu, Tingjia Shen, Chuan Qin, Chen Zhu, Hengshu Zhu, Qi Liu, et al. 2024. A survey on large language models for recommendation. World Wide Web 27, 5 (2024), 60.
 - [57] Shiwen Wu, Fei Sun, Wentao Zhang, Xu Xie, and Bin Cui. 2022. Graph neural networks in recommender systems: a survey. *Comput. Surveys* 55, 5 (2022), 1–37.
 - [58] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24.
- [59] Yikun Xian, Zuohui Fu, Shan Muthukrishnan, Gerard De Melo, and Yongfeng Zhang. 2019. Reinforcement knowledge graph reasoning for explainable recommendation. In *Proceedings of the 42nd international ACM SIGIR conference on research and development in information retrieval.* 285–294.
 - [60] An Yan, Zhankui He, Jiacheng Li, Tianyang Zhang, and Julian McAuley. 2023. Personalized showcases: Generating multi-modal explanations for recommendations. In Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval. 2251–2255.
 - [61] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024. Qwen2 technical report. arXiv preprint arXiv:2407.10671 (2024).
 - [62] Ruosong Ye, Caiqi Zhang, Runhui Wang, Shuyuan Xu, Yongfeng Zhang, et al. 2023. Natural language is all a graph needs. arXiv preprint arXiv:2308.07134 4, 5 (2023), 7.
 - [63] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining. 974–983.
 - [64] Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. 2019. Gnnexplainer: Generating explanations for graph neural networks. Advances in neural information processing systems 32 (2019).
 - [65] Hao Yuan, Haiyang Yu, Shurui Gui, and Shuiwang Ji. 2022. Explainability in graph neural networks: A taxonomic survey. *IEEE transactions on pattern analysis* and machine intelligence 45, 5 (2022), 5782–5799.
 - [66] Weizhe Yuan, Graham Neubig, and Pengfei Liu. 2021. Bartscore: Evaluating generated text as text generation. Advances in Neural Information Processing Systems 34 (2021), 27263–27277.
 - [67] Shichang Zhang, Jiani Zhang, Xiang Song, Soji Adeshina, Da Zheng, Christos Faloutsos, and Yizhou Sun. 2023. PaGE-Link: Path-based graph neural network explanation for heterogeneous link prediction. In Proceedings of the ACM Web Conference 2023. 3784–3793.
 - [68] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. arXiv preprint arXiv:1904.09675 (2019).
 - [69] Tianjun Zhang, Shishir G Patil, Naman Jain, Sheng Shen, Matei Zaharia, Ion Stoica, and Joseph E Gonzalez. 2024. Raft: Adapting language model to domain specific rag. arXiv preprint arXiv:2403.10131 (2024).
 - [70] 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.
 - [71] Yongfeng Zhang, Guokun Lai, Min Zhang, Yi Zhang, Yiqun Liu, and Shaoping Ma. 2014. Explicit factor models for explainable recommendation based on phrase-level sentiment analysis. In Proceedings of the 37th international ACM SIGIR conference on Research & development in information retrieval. 83–92.
- [72] Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu, Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang, Yulong Chen, et al. 2023. Siren's song in the AI ocean: a survey on hallucination in large language models. arXiv preprint arXiv:2309.01219 (2023).
- [73] Jianan Zhao, Le Zhuo, Yikang Shen, Meng Qu, Kai Liu, Michael Bronstein, Zhaocheng Zhu, and Jian Tang. 2023. Graphtext: Graph reasoning in text space.
 arXiv preprint arXiv:2310.01089 (2023).

1103

1104

1105

1106

1107

1108

1109

1110

1111

1112

1113

1114

1115

1116

1117

1118

1119

1120

1121

1123

1124

1125

1126

1127

1129

1130

1131

1132

1133

1134

1136

1137

1138

1139

1140

1141

- [74] 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).
- [75] Yaxin Zhu, Yikun Xian, Zuohui Fu, Gerard De Melo, and Yongfeng Zhang. 2021. Faithfully explainable recommendation via neural logic reasoning. arXiv preprint arXiv:2104.07869 (2021).

1154 1155 1156

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

1275

1276

A MORE EXPERIMENT DETAILS

A.1 Details of Datasets

1161

1162

1163

1164

1165

1166

1167

1176

1177

1185

1186

1187

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

- Amazon-books [41] is a subset of the Amazon review dataset, specifically focused on book recommendations. This dataset encompasses user-book interactions, including both numerical ratings and textual reviews submitted by users after their purchases.
- Yelp [39] is a widely used dataset derived from Yelp¹, where local businesses such as restaurants and bars are considered items. It provides rich information about local businesses, including user reviews and ratings across various categories.
- Google-reviews [28, 60] comprises restaurant reviews collected from Google Local². This dataset incorporates both business metadata and user feedback, offering a broad perspective on dining establishments globally.

A.2 Details of Metrics

- IT78 GPT_{score} [54] leverages large language models to evaluate text quality, providing a context-aware assessment.
- **BERT**_{score} [68] computes the similarity between reference and generated texts using contextual embeddings from BERT [12]. Given a reference sentence $x = \langle x_1, x_2, ..., x_n \rangle$ and a generated sentence $\hat{x} = \langle \hat{x}_1, \hat{x}_2, ..., \hat{x}_m \rangle$, A sequence of word embeddings are first generated using BERT:

$$BERT(\langle x_1, x_2, ..., x_n \rangle) = \langle \mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n \rangle$$
$$BERT(\langle \hat{x}_1, \hat{x}_2, ..., \hat{x}_m \rangle) = \langle \hat{\mathbf{x}}_1, \hat{\mathbf{x}}_2, ..., \hat{\mathbf{x}}_m \rangle$$
(13)

The similarity between two individual embeddings $(\mathbf{x}_i, \hat{\mathbf{x}}_j)$ is measured using cosine similarity, which simply reduces to $\mathbf{x}_i^{\top} \hat{\mathbf{x}}_j$ since both embeddings are pre-normalized. With these definitions, the Precision, Recall, and F1-score are calculated as follows:

cosine similarity

$$BERT_{score}^{P} = \frac{1}{|\hat{x}|} \sum_{\hat{x}_{j} \in \hat{x}} \max_{x_{i} \in x} \quad \mathbf{x}_{i}^{\top} \hat{x}_{j}$$
(14)

greedy matching

cosine similarity

$$BERT_{score}^{R} = \frac{1}{|x|} \sum_{\mathbf{x}_{i} \in \mathbf{x}} \max_{\hat{\mathbf{x}}_{j} \in \hat{\mathbf{x}}} \mathbf{x}_{i}^{\top} \hat{\mathbf{x}}_{j}$$
(15)

greedy matching

$$BERT_{score}^{F1} = 2 \times \frac{BERT_{score}^{P} \times BERT_{score}^{R}}{BERT_{score}^{P} + BERT_{score}^{R}}$$
(16)

- BART_{score} [66] conceptualizes the evaluation as a text generation task, assigning scores based on the probability of regenerating reference texts using the BART model [26].
- BLEURT [50] employs a novel language model pre-trained with synthetic data to assess the similarity between the generated and reference texts.
 - USR [30] assesses the uniqueness of generated explanations by calculating the ratio of unique sentences to total sentences.

¹https://www.yelp.com/dataset/

1217 ²https://www.google.com/maps

1218

1213

1214

1215

1216

A.3 Details of Baselines

- NRT [32] employs multi-task learning to predict ratings and generate tips for recommendations simultaneously based on user and item IDs. The generation component is a GRU [10].
- Att2Seq [14] implements an attention-based attribute-to-sequence model that generates reviews by leveraging attribute information. The generation component is a two-layer LSTM [19].
- **PETER** [30] is a personalized Transformer model that maps user and item IDs to generated explanations. It bridges IDs and words through a "context prediction" task. PETER is used instead of PETER+ due to the absence of word features in the datasets.
- **PEPLER** [31] proposes sequential tuning and recommendation as regularization strategies to bridge the gap between prompts (incorporating user and item ID vectors) and the pre-trained transformer model for generating explanations.
- **XRec** [39] utilizes the encoded user/item embeddings from GNNs as implicit collaborative signals, which are then integrated into each layer of LLMs, enabling the generation of explanations.

B CASE STUDY

We present two cases in Table 5 and Table 6 to demonstrate the effectiveness of G-Refer and illustrate how retrieved CF signals benefit the generated explanations. All retrieved knowledge is presented as human-readable text after graph translation for better user understanding. We also provide the ground truth explanations and explanations generated by XRec [39] for comparison.

From Table 5, we can observe that the user profile contains no explicit drinking interests. However, node-level CF signals reveal that other users who have similar interests to the current user show a preference for drinking. This allows G-Refer to infer that the current user might appreciate this restaurant specializing in "crafted beers". In contrast, XRec's explanation is notably generic, which relies on common phrases like "delicious food" and "great service", missing the key recommendation reasons, thus leading to a low Recall score. Table 6 provides another case that leverages path-level retrieved CF information for more comprehensive explanations. While the user profile only indicates general preferences for comfort food and efficient services, path-level CF signals uncover broader connections between the current user and other users/items. Users along the retrieved path have an interest in "frozen yogurt", allowing us to infer that the current user might be interested in this restaurant due to its frozen yogurt offerings. In contrast, XRec incorrectly explains that the user might choose this restaurant for its Thai food, failing to capture the user's true interests. These cases show that node-level and path-level CF signals are complementary. While not equally effective in every situation, gathered CF signals can contribute to generating more accurate explanations in recommendation scenarios. The better explanations demonstrate G-Refer has the ability to effectively retrieve, understand, and utilize such knowledge.

Table 5: A case from Yelp, which shows how G-Refer primarily leverages node-level retrieved CF information to generate better explanations.

	Inputs
User Profile	The user is likely to enjoy Mexican restaurants with a focus on burritos and tacos, as well as baseball games at scenic locations and coffee shops for cozy work or socializing.
tem Profile	Beer enthusiasts looking for a laid-back atmosphere with a variety of craft beers and casual dining options, including dog owners seeking a pet-friendly environment, would enjoy this business.
Node-level CF	For the user-item pair, here are some related users and items: Users: The user is likely to enjoy businesses with unique decor traditional family recipes, craft beers, outdoor dining, gourmet wood-fired entrees, live music, and waterfront locations., The user is likely to enjoy fine dining experiences, laid-back accommodations with resort amenities, American pub fare with craft beer, and relaxed atmospheres with delicious comfort food and craft beer selections. Items: Coffee enthusiasts looking for a cozy atmosphere with delicious pastries and brunch options would enjoy Lulu Carpenter's., Food enthusiasts who appreciate classic Italian cuisine and casual dining experiences would enjoy Joe's Pizza & Subs.
	Explanations
Ground Truth	The user would enjoy the business for its awesome selection of beers and ciders, great outdoor spaces, and the option to order bar food from Carpos, even though it predominantly focuses on beer.
<u>KRec</u>	The user would enjoy the business because of its delicious food, great service, and cozy atmosphere, making it a great place to relax and enjoy a good meal.
<u>G-Refer</u> (ours)	The user would enjoy the business because of the great selection of beers, friendly staff, and the opportunity to try new and unique brews like the hazy IPA.

Table 6: A case from Google-review, which shows how G-Refer primarily leverages path-level retrieved CF information to generate better explanations.

User Profile This user is likely to enjoy businesses that offer comfort food with unique twists, efficient waxing services with minimal pain, upscale dining with stunning views and impressive service, and flavorful Asian cuisine with fast service in Philadelphia. Item Profile Fans of fozen yogurt with a wide variety of flavors and toppings, in a vibrant and welcoming setting in Philadelphia, would enjoy Berry Sweet Frozen Yogurt. Ideal for those who appreciate a fun atmosphere, quality treats, and convenient location on South Street. Path-level CF For the given user-item pair, here are several related paths connecting users and items through their interactions: 1. User (Profile: This user is likely to enjoy businesses that offer comfort food with unique twists, efficient waxing services with minimal pain, upscale dining with stunning views and impressive service, and flavorful Asian cuisine with fast service in Philadelphia.) -> buys -> Item (Profile: Users who appreciate high-quality skincare products, attentive customer service, and are willing to splurge for effective beauty items would enjoy Keihl's in Philadelphia,) -> bought by -> User (Profile: This user is likely to enjoy businesses offering high-quality od and attentive service in cozy settings, such as restaurants like Kanella, Village Whiskey, and Cheu Noodle Bar. They appreciate flavorful meals, knowledgeable staff, and unique dining experiences.) -> buys -> Item (Profile: Users who appreciate and high-quality vegan options, a personalized dining experience with small restaurant owners, unique and experimental dishes with varied influences, and a great wine and cocktail selection would enjoy Meritage in Philadelphia,) -> bought by -> User (Profile: The user is likely to enjoy upscale dining with Latin American and Mexican cuisine, diverse frozen yogurt flavors and toppings, in a vibrant and welcoming setti		Inputs
Item Profile Fans of frozen yogurt with a wide variety of flavors and toppings, in a vibrant and welcoming setting in Philadelphia, would enjoy Berry Sweet Frozen Yogurt. Ideal for those who appreciate a fun atmosphere, quality treats, and convenient location on South Street. Path-level CF For the given user-item pair, here are several related paths connecting users and items through their interactions: 1. User (Profile: This user is likely to enjoy businesses that offer comfort food with unique twists, efficient waxing services with minimal pain, upscale dining with stunning views and impressive service, and flavorful Asian cuisine with fast service in Philadelphia.) -> buys -> Item (Profile: Users who appreciate high-quality skincare products, attentive customer service, and are willing to splurge for effective beauty items would enjoy Kiehl's in Philadelphia.) -> bought by -> User (Profile: This user is likely to enjoy businesses offering high-quality food and attentive service in cozy settings, such as restaurants like Kanella, Village Whiskey, and Cheu Noodle Bar. They appreciate flavorful meals, knowledgeable staff, and unique dining experiences.) -> buys -> Item (Profile: Users who appreciate creative and high-quality vegan options, a personalized dining experience with small restaurant owners, unique and experimental dishes with varied influences, and a great wine and cocktail selection would enjoy Meritage in Philadelphia.) -> buyst by -> User (Profile: The user is likely to enjoy upscale dining with Latin American and Mexican cuisine diverse frozen yogurt flavors about service and chicken quality.) -> buys -> Item (Profile: Fans of frozen yogurt, setting and toppings, in a vibrant and welcoming setting in Philadelphia, would enjoy Berry Sweet Frozen Yogurt. Ideal for those who appreciate a fun atmosphere, quality treats, and convenient location on South Street.)	User Profile	This user is likely to enjoy businesses that offer comfort food with unique twists, efficient waxing services with minimal pain, upscale dining with stunning views and impressive service, and flavorful Asian cuisine with fast service in Philadelphia.
Path-level CF For the given user-item pair, here are several related paths connecting users and items through their interactions: 1. User (Profile: This user is likely to enjoy businesses that offer comfort food with unique twists, efficient waxing services with minimal pain upscale dining with stunning views and impressive service, and flavorful Asian cuisine with fast service in Philadelphia.) -> buys -> Item (Profile: Users who appreciate high-quality skincare products, attentive customer service, and are willing to splurge for effective beauty items would enjoy Kiehl's in Philadelphia.) -> bought by -> User (Profile: This user is likely to enjoy businesses offering high-quality food and attentive service in cozy settings, such as restaurants like Kanella, Village Whiskey, and Cheu Noodle Bar. They appreciate flavorful meals, knowledgeable staff, and unique dining experiences.) -> buys -> Item (Profile: Users who appreciate creative and high-quality vegan options, a personalized dining experiences with small restaurant owners, unique and experimental dishes with varied influences, and a great wine and cocktail selection would enjoy Meritage in Philadelphia.) -> buys -> User (Profile: The user is likely to enjoy upscale dining with Latin American and Mexican cuisine, diverse frozen yogurt flavors with unique bubble tea options, good food in a low-key ambiance, and authentic spicy Indian cuisine despite some reservations about service and chicken quality.) -> buys -> Item (Profile: Fans of frozen yogurt with a wide variety of flavors and toppings, in a vibrant and welcoming setting in Philadelphia, would enjoy Berry Sweet Frozen Yogurt. Ideal for those who appreciate a function on South Street. Cround Truth The user would enjoy Berry Sweet because of its delicious frozen yogurt, friendly staff, variety of toppings, and fun atmosphere that makes it a great place to hang out with friends on South Street.	Item Profile	Fans of frozen yogurt with a wide variety of flavors and toppings, in a vibrant and welcoming setting in Philadelphia, would enjoy Berry Sweet Frozen Yogurt. Ideal for those who appreciate a fun atmosphere, quality treats, and convenient location on South Street.
Explanations Ground Truth The user would enjoy Berry Sweet because of its delicious frozen yogurt, friendly staff, variety of toppings, and fun atmosphere that makes it a great place to hang out with friends on South Street. XRec You would enjoy the business because it offers delicious, authentic Thai food with a variety of dishes to choose from, including spicy and non-spicy options, as well as a cozy atmosphere with friendly service. G-Refer (ours) The user would enjoy Berry Sweet because of the wide variety of flavors and the opportunity to mix and match them, making it a fun for yogurt lovers seeking a customizable and fun experience.	Path-level CF	For the given user-item pair, here are several related paths connecting users and items through their interactions: 1. User (Profile: This user is likely to enjoy businesses that offer comfort food with unique twists, efficient waxing services with minimal pain, upscale dining with stunning views and impressive service, and flavorful Asian cuisine with fast service in Philadelphia.) -> buys -> Item (Profile: Users who appreciate high-quality skincare products, attentive customer service, and are willing to splurge for effective beauty items would enjoy Kiehl's in Philadelphia.) -> bought by -> User (Profile: This user is likely to enjoy businesses offering high-quality food and attentive service in cozy settings, such as restaurants like Kanella, Village Whiskey, and Cheu Noodle Bar. They appreciate flavorful meals, knowledgeable staff, and unique dining experiences.) -> buys -> Item (Profile: Users who appreciate creative and high-quality vegan options, a personalized dining experience with small restaurant owners, unique and experimental dishes with varied influences, and a great wine and cocktail selection would enjoy Meritage in Philadelphia.) -> bought by -> User (Profile: The user is likely to enjoy upscale dining with Latin American and Mexican cuisine, diverse frozen yogurt flavors with unique bubble tea options, good food in a low-key ambiance, and authentic spicy Indian cuisine despite some reservations about service and chicken quality.) -> buys -> Item (Profile: Fans of frozen yogurt with a wide variety of flavors and toppings, in a vibrant and welcoming setting in Philadelphia, would enjoy Berry Sweet Frozen Yogurt. Ideal for those who appreciate a fun atmosphere, quality treats, and convenient location on South Street.)
Ground Truth The user would enjoy Berry Sweet because of its delicious frozen yogurt, friendly staff, variety of toppings, and fun atmosphere that makes it a great place to hang out with friends on South Street. XRec You would enjoy the business because it offers delicious, authentic Thai food with a variety of dishes to choose from, including spicy and non-spicy options, as well as a cozy atmosphere with friendly service. G-Refer (ours) The user would enjoy Berry Sweet because of the wide variety of flavors and the opportunity to mix and match them, making it a fun for yogurt lovers seeking a customizable and fun experience.		Explanations
XRec You would enjoy the business because it offers delicious, authentic Thai food with a variety of dishes to choose from, including spicy and non-spicy options, as well as a cozy atmosphere with friendly service. G-Refer (ours) The user would enjoy Berry Sweet because of the wide variety of flavors and the opportunity to mix and match them, making it a fun for yogurt lovers seeking a customizable and fun experience.	Ground Truth	The user would enjoy Berry Sweet because of its delicious frozen yogurt, friendly staff, variety of toppings, and fun atmosphere that makes it a great place to hang out with friends on South Street.
<u>G-Refer</u> (ours) The user would enjoy Berry Sweet because of the wide variety of flavors and the opportunity to mix and match them, making it a fun for yogurt lovers seeking a customizable and fun experience.	XRec	You would enjoy the business because it offers delicious, authentic Thai food with a variety of dishes to choose from, including spicy and non-spicy options, as well as a cozy atmosphere with friendly service.
	<u>G-Refer</u> (ours)	The user would enjoy Berry Sweet because of the wide variety of flavors and the opportunity to mix and match them, making it a fun for yogurt lovers seeking a customizable and fun experience.
		12