

Leveraging LLM Reasoning Enhances Personalized Recommender Systems

Alicia Y. Tsai*^{†1,3}, Adam Kraft*³, Long Jin², Chenwei Cai²,
Anahita Hosseini³, Taibai Xu², Zemin Zhang², Lichan Hong³,
Ed H. Chi³, Xinyang Yi³

¹University of California, Berkeley ²Google ³Google DeepMind

Abstract

Recent advancements have showcased the potential of Large Language Models (LLMs) in executing reasoning tasks, particularly facilitated by Chain-of-Thought (CoT) prompting. While tasks like arithmetic reasoning involve clear, definitive answers and logical chains of thought, the application of LLM reasoning in recommendation systems (RecSys) presents a distinct challenge. RecSys tasks revolve around subjectivity and personalized preferences, an under-explored domain in utilizing LLMs' reasoning capabilities. Our study explores several aspects to better understand reasoning for RecSys and demonstrate how task quality improves by utilizing LLM reasoning in both zero-shot and fine-tuning settings. Additionally, we propose **RecSAVER** (Recommender Systems Automatic Verification and Evaluation of Reasoning) to automatically assess the quality of LLM reasoning responses without the requirement of curated gold references or human raters. We show that our framework aligns with real human judgment on the coherence and faithfulness of reasoning responses. Overall, our work shows that incorporating reasoning into RecSys can improve personalized tasks, paving the way for further advancements in recommender system methodologies.

1 Introduction

The rapid advancement of Large Language Models (LLMs) has ushered in a new era of transformative capabilities, demonstrating the potential across a spectrum of applications. Recent progress has showcased their ability in reasoning tasks. Particularly, the advent of Chain-of-Thought (CoT) [30] and zero-shot CoT prompting [19] has provided a pathway for these models to engage in multi-step reasoning. Tasks examined in these studies, ranging from arithmetic reasoning [1, 5, 16] to common-

sense question answering [14, 25], typically demand clear, definitive answers and logical chains of thought. In contrast, the landscape of recommender systems (RecSys), outlined in Figure 1, introduces a nuanced challenge, where reasoning extends beyond objective criteria to encompass subjectivity and personalized user preferences. This aspect remains an under-explored domain in leveraging the reasoning capabilities of LLMs. Prior works have utilized LLMs in recommender systems, employing techniques such as in-context learning and instruction tuning [12, 13, 18, 32]. However, a comprehensive understanding of how LLMs execute reasoning in the context of personalized preferences remains elusive. Our work fills this gap by investigating how the reasoning capabilities of LLMs can enhance personalized recommendations in zero-shot and fine-tuning, resulting in improved task performance.

In contrast to the prediction task performance, objectively assessing the quality of reasoning presents challenges in the absence of curated gold standard references or human raters. To surmount this, we propose **RecSAVER**. This framework provides an efficient means of assessing the quality of LLM outputs, contributing to our understanding of LLM reasoning dynamics in personalized recommendation scenarios. We gauge human assessments on *coherence*, *faithfulness*, and *insightfulness*. Our observations suggest that syntactic metrics such as BLEU and ROUGE, demonstrate suitability in evaluating the output faithfulness of LLMs. On the other hand, metrics like METEOR and BERTScore prove to be more adept at measuring coherence in the generated outputs. To the best of our knowledge, this is the first study that comprehensively examines the effects and quality of LLM reasoning for personalized RecSys tasks. In summary, our contributions are as follows:

1. We explore the utilization of LLMs for reasoning in personalized recommendations, showcasing notable task performance improvements in both zero-shot and fine-tuning scenarios.
2. We demonstrate the effectiveness of using larger models to generate reasoning data, enhancing the performance and reasoning abilities of smaller fine-tuned models.
3. We introduce **RecSAVER**, an automatic reasoning evaluation framework that doesn't require curated gold references. It offers meaningful insights

*Equal contribution. Contact: alicia-tsai@berkeley.edu, adamkraft@google.com.

[†]Work done while at Google DeepMind.

into LLMs’ reasoning capabilities, aligning with human judgment while providing cost and efficiency benefits.

2 Methodology

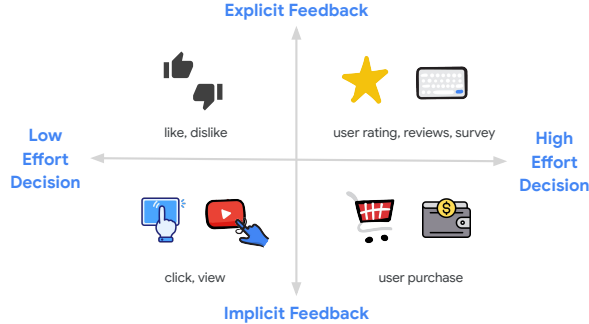


Figure 1: Landscape of recommender systems tasks, with user feedback extent on the vertical axis and decision-making effort on the horizontal. For example, a user clicking on websites requires low effort and does not provide much feedback about the user’s satisfaction. Conversely, a user rating and reviewing products requires more effort and provides better satisfaction signals.

2.1 Problem Setting

The landscape of RecSys tasks, segmented along two axes, is illustrated in Figure 1. In our experiments, we focus on the *user rating prediction* task, which involves a high degree of both user decision making effort and collected user feedback, making this task well suited for exploring the extent of LLM reasoning for RecSys.

Let \mathcal{R} represent a collection of user ratings and \mathcal{D} denote a collection of user-written reviews for a set of items \mathcal{I} belonging to a specific category (e.g., books), provided by users in \mathcal{U} . Each rating $\mathbf{r}_{u,i} \in \mathcal{R}$ is paired with a corresponding written review $\mathbf{d}_{u,i} \in \mathcal{D}$, reflecting the user’s $u \in \mathcal{U}$ overall satisfaction with item $i \in \mathcal{I}$. Each item i is associated with metadata \mathcal{M}_i , comprising details such as title, description, category, brand, and price. The user’s purchase history $\mathcal{H}_u = (h_{u,1}, h_{u,2}, \dots, h_{u,t})$ constitutes a chronologically ordered collection of past purchases. Each past purchase $h_{u,j} = (\mathcal{M}_j, \mathbf{r}_{u,j}, \mathbf{d}_{u,j})$ represents a triplet comprising the metadata for the purchased item j , the user’s rating for the purchase item, and the user’s review for that item. The primary objective of the rating prediction task is to forecast the unknown ratings for items that users have not yet reviewed. The objective can be formalized as follows:

$$\hat{\mathbf{r}}_{u,i} = \arg \max_k \mathbb{P}(\mathbf{r}_{u,i} = k \mid \mathcal{H}_u, \mathcal{M}_i), \quad (1)$$

where $i \notin \mathcal{H}_u$, $k \in \{1, 2, 3, 4, 5\}$.

In this equation, $\hat{\mathbf{r}}_{u,i}$ represents the predicted rating for user u on item i , chosen from the set of possible ratings 1 to 5. This prediction is based on the user’s purchase

Table 1: Abstract prompt template guiding our zero-shot approach, prompting the model to reason by leveraging user history and inferring preferences before making predictions.

<i>Preamble</i>	e.g. Here is information about a user and a new product ...
<i>User History</i>	$h_{u,1} = (\mathcal{M}_1, \mathbf{r}_{u,1}, \mathbf{d}_{u,1}), \dots, h_{u,t} = (\mathcal{M}_t, \mathbf{r}_{u,t}, \mathbf{d}_{u,t})$
<i>New Item</i>	\mathcal{M}_i , e.g. title, brand, category, ...
<i>Task Description</i>	e.g. Given the user’s past purchase history [...] how they will rate the new item? [...] After your reasoning, predict a numerical rating.

history \mathcal{H}_u and the item’s metadata \mathcal{M}_i , where the item i has not been previously reviewed by user u . Recent advancements in recommendation systems utilize LLMs to model the rating prediction task described by equation (1), denoted as $\hat{\mathbf{r}}_{u,i} = \text{LLM}(\mathcal{H}_u, \mathcal{M}_i)$.

2.2 Zero-shot Learning with Reasoning



Figure 2: We prompt the LLM to generate a reasoning output prior to outputting the final task prediction.

As shown in Figure 2, we employ zero-shot CoT prompting strategies [19] to guide the LLM in generating a reasoning response, denoted as $\hat{\mathbf{s}}_{u,i}$, alongside a rating prediction $\hat{\mathbf{r}}_{u,i}$ for a given user u and recommended item i :

$$\langle \hat{\mathbf{s}}_{u,i}, \hat{\mathbf{r}}_{u,i} \rangle = \text{LLM}(\mathcal{H}_u, \mathcal{M}_i). \quad (2)$$

Our prompt consists of four key elements: a *preamble*, the *user history* \mathcal{H}_u , the *new item metadata* \mathcal{M}_i , and a *task description*. The *preamble* provides context for the subsequent information and establishes the rating scale ranging from 1 to 5. Following the *preamble*, the *user history* \mathcal{H}_u is presented sequentially, detailing the user’s past interactions. A new item i is then introduced along with its metadata \mathcal{M}_i before the *task description*, prompting the model to make predictions. The *task description* also delineates the output requirements for the model responses. An abstract prompt template is illustrated in Table 1, while more prompt details are provided in Appendix A. Unlike traditional RecSys modeling techniques, our approach leverages natural language presentation for all information. This enables a more intuitive representation of rich content as natural language, as opposed to numerical IDs, enabling a more encompassing understanding of information [13].

2.3 Fine-tuning with Reasoning

Zero-shot learning with CoT prompting can be computationally intensive. Hence, fine-tuning with domain-specific datasets has emerged as a pragmatic strategy, especially when leveraging smaller pre-trained models [4]. Our interest lies in investigating whether training with reasoning outputs can further enhance task performance. Building on the prompting methods outlined in

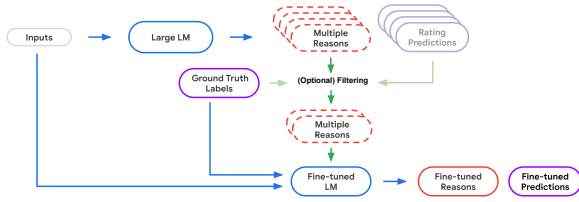


Figure 3: Fine-tuning a model with reasoning. We first collect multiple reasoning samples by prompting a Large LM. We then use the reasoning samples combined with the original rating ground truth labels to fine-tune a different (potentially smaller) LM. We can optionally filter the reasoning outputs by comparing the Large LM rating predictions with the ground truth ratings.

Section 2.2, we collect reasoning outputs generated by a larger language model to serve as training data for fine-tuning smaller models. For each user-item pair (u, i) with input $(\mathcal{H}_u, \mathcal{M}_i)$, we gather multiple reasoning responses and rating predictions by adjusting a decoding temperature parameter $T > 0$ during generation [17]:

$$\langle \hat{\mathbf{s}}_{u,i}^m, \hat{\mathbf{r}}_{u,i}^m \rangle = \text{LLM}(\mathcal{H}_u, \mathcal{M}_i), \quad (3)$$

where $m = 1, 2, \dots, M$ indexes the M candidate output pairs sampled from the decoder. This process yields a diverse set of reasoning paths, which is particularly advantageous for personalized recommendations, recognizing that the same rating can stem from various personal preferences and reasons. Optionally, we can use different methods to filter out reasoning responses $\hat{\mathbf{r}}_{u,i}^m$ that do not align with the ground truth $\mathbf{r}_{u,i}$. The fine-tuned model is then trained using the reasoning responses $\hat{\mathbf{s}}_{u,i}^m$ and the real ground truth rating label $\mathbf{r}_{u,i}$ as targets. The overall method is illustrated in Figure 3.

3 Rec-SAVER: Evaluation of Reasoning

In contrast to reasoning processes for solving mathematical problems or general question answering tasks, reasoning in RecSys rating prediction is highly subjective and personalized for individual users. Unlike in other domains where humans can provide reasoning steps and verify their validity, resulting in curated gold references, such references are challenging to obtain in RecSys due to the subjective nature of user preferences. To address this challenge, we propose a system called **Rec-SAVER: Recommender Systems Automatic Verification and Evaluation of Reasoning**. Rec-SAVER aims to automatically generate good reasoning references specifically tailored for RecSys tasks. These references can then be utilized to quantitatively evaluate the quality of reasoning responses generated by LLMs. Additionally, we conduct a human study to demonstrate the alignment of our method with real human judgment, thus providing validation for the effectiveness of Rec-SAVER in assessing reasoning quality in RecSys applications.

3.1 Reference Generation with Self-Verification

The core concept of Rec-SAVER involves a two-step process leveraging LLM-generated explanations and LLM self-verification. As illustrated in Figure 4, we present the LLM with a user-item pair $(\mathcal{H}_u, \mathcal{M}_i)$. Additionally the target user rating $\mathbf{r}_{u,i}$ is also provided as input. The model is instructed to provide a post hoc explanation, describing why the user assigned such a rating based on the given user history and new item information. We denote this post hoc explanation generated by LLM as $\hat{\mathbf{g}}_{u,i}$. Note that this is different from the aforementioned reasoning $\hat{\mathbf{s}}_{u,i}$, where the ground truth rating $\mathbf{r}_{u,i}$ is not included as an input. Similar to the approach in Section 2.3 where multiple responses are sampled, we generate N different explanations $\hat{\mathbf{g}}_{u,i}^n$ where $n = 1, 2, \dots, N$. These post hoc explanations are then passed onto a verification process.

To ensure the credibility and consistency of these LLM-generated explanations $\hat{\mathbf{g}}_{u,i}^n$, we implement a self-verification step atop the previous explanation generation process. The self-verification step involves making a second call to the same LLM, inputting the user-item information $(\mathcal{H}_u, \mathcal{M}_i)$, and the explanations $\hat{\mathbf{g}}_{u,i}^n$ generated from the previous call. The model is then tasked with making a rating prediction based on the user history, new item information, and the post hoc explanation, formally defined as $\tilde{\mathbf{r}}_{u,i}^n = \text{LLM}(\mathcal{H}_u, \mathcal{M}_i, \hat{\mathbf{g}}_{u,i}^n)$. However, in practice, we have observed that many explanations $\hat{\mathbf{g}}_{u,i}^n$ contain text snippets such as “the user gave a rating of 5 because ...”, which can lead to information leakage. To prevent $\hat{\mathbf{g}}_{u,i}^n$ from directly leaking the ground truth, we employ a simple post-processing step by removing sentences that mention “a rating of”, “stars”, or “scores” before performing the prediction. In future work we aim to improve upon this manual process to fully ensure the removal of information leakage.

We then validate whether the new rating $\tilde{\mathbf{r}}_{u,i}^n$ matches the original ground truth $\mathbf{r}_{u,i}$. Explanations $\hat{\mathbf{g}}_{u,i}^n$ that pass the self-verification step are retained as the final verified references, constituting a diverse pool of LLM-generated references $\hat{\mathcal{G}}$. This two-step process follows the intuition that good explanations based on the given information and the ground truth should enable the model to make a correct prediction. By validating the predictions based on the generated explanations against the ground truth ratings, we ensure that only high-quality explanations are retained to serve as the final references. These verified references then serve as proxies for the unknown set of gold references \mathcal{G} . Since the self-verification may result in different verified references per sample, we may have varying numbers of final references per sample. The full reference generation process is summarized in Algorithm 1.

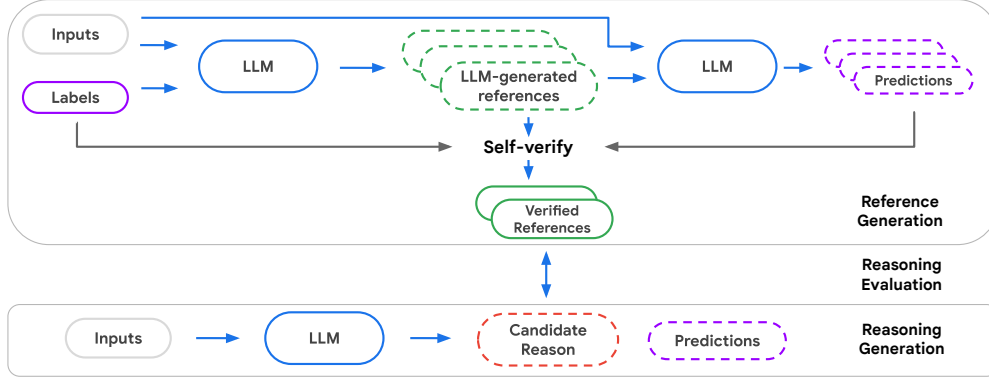


Figure 4: Overview of **Rec-SAVER** utilizing LLM-generated references and LLM self-verification. The first LLM call uses the ground truth rating labels as additional input to generate post hoc reasoning generated reference. We then do a subsequent LLM call passing in the generated reasoning reference and collect a new rating prediction. We keep only the predictions where the final rating prediction matches the ground truth rating label as our verified references. These verified references are then used to evaluate the reasoning outputs from other LLMs.

Algorithm 1 Reference generation with self-verification

```

1: Inputs:  $N$ 
2:  $\hat{\mathcal{G}} \leftarrow \emptyset$  ▷ verified references
3: for  $(\mathcal{H}_u, \mathcal{M}_i, \mathbf{r}_{u,i})$  in dataset do
4:   for  $n = 1 \dots N$  do
5:      $\hat{\mathbf{g}}_{u,i}^n \leftarrow \text{LLM}(\mathcal{H}_u, \mathcal{M}_i, \mathbf{r}_{u,i})$ 
6:      $\hat{\mathbf{g}}_{u,i}^n \leftarrow \text{post-process}(\hat{\mathbf{g}}_{u,i}^n)$ 
7:      $\tilde{\mathbf{r}}_{u,i}^n \leftarrow \text{LLM}(\mathcal{H}_u, \mathcal{M}_i, \hat{\mathbf{g}}_{u,i}^n)$ 
8:     if  $\tilde{\mathbf{r}}_{u,i}^n = \mathbf{r}_{u,i}$  then
9:        $\hat{\mathcal{G}} \leftarrow \hat{\mathcal{G}} \cup \{\hat{\mathbf{g}}_{u,i}^n\}$ 
10:    end if
11:  end for
12: end for

```

3.2 Human Judgment Alignment Study

To gauge the effectiveness of Rec-SAVER, we conduct a study to evaluate how our proposed method aligns with human judgment regarding the candidate reasons generated by LLMs. This study aims to provide insight into the reliability and validity of our proposed method. During the study, human raters are presented with sample input prompts and the reasoning outputs $\hat{\mathbf{s}}_{u,i}$ generated by LLMs. It is important to note that these reasoning outputs $\hat{\mathbf{s}}_{u,i}$ are produced using only user history \mathcal{H}_u and item metadata \mathcal{M}_i as inputs; no ground truth rating $\mathbf{r}_{u,i}$ is provided to the LLM. No ratings (ground truth or LLM predicted) are shown to the human raters. Human raters are asked to assess the reasoning outputs based on the following dimensions:

- **Coherence (5-point Likert):** Evaluate whether the generated reasoning makes sense and follows a clear and coherent logical flow that reflects the reasons behind the user’s preference.
- **Faithfulness (Binary):** Examine the presence of hallucination in the generated reasoning and whether it contains fabricated information.

- **Insightfulness (5-point Likert):** Assess the degree to which the generated reasoning delivers valuable, informative, interesting or delightful insights into the user’s preferences and purchasing patterns.

4 Experiments

4.1 Data Preparation and Tasks Setup

Experiments are conducted on the Amazon product review dataset*, which is a widely recognized benchmark in RecSys. This dataset offers comprehensive user feedback, including ratings and review text, as well as detailed product metadata such as descriptions, category information, price, and brand [22]. We focus our experiments on the rating prediction task, conducting evaluations in two distinct domains: BEAUTY and MOVIES/TV. To understand the extent that LLMs can understand user preferences, we filter examples where $4 \leq |\mathcal{H}_u| \leq 10$. The lower threshold ensures we have enough past purchases to see trends and patterns while the higher threshold prevents inputs from exceeding the LLM context window. The original label distribution is heavily skewed towards positive ratings, with a rating of 5 accounting for over 60 % of the data. We perform random subsampling to create a fully balanced dataset with an even label distribution, resulting in 4,000 training examples (800 per label) and 500 test examples (100 per label). The training split is used to test out different prompts for zero-shot learning and as training examples for the fine-tuning experiments. The user sets are mutually exclusive between training and test. This setup allows us to better test and understand the capabilities of LLM reasoning, while we acknowledge this may not reflect a full real world scenario.

In following sections, we report the rating task predic-

*https://cseweb.ucsd.edu/~jmcauley/datasets/amazon_v2/

tion metrics, including multi-class and binary metrics. Specifically, we compare the model’s output $\hat{r}_{u,i}$ against the ground truth rating $r_{u,i}$. The binary metrics were calculated using a cutoff threshold of $\hat{r}_{u,i} > 3$. The quality of the reasoning outputs will be analyzed in more detail later on in Section 5.

Table 2: Example task description prompt. The highlighted text represents portions removed when prompting the model to make predictions without the intermediate reasoning step.

Task Description
Given the user’s past purchase history [...] how they will rate the new item? [...] After your reasoning, predict a numerical rating.
=== Please follow the format below: ===
Reason
Write your reasoning explanation here.
Rating
Give a single numerical rating, e.g. 1

4.2 Zero-shot Learning Improves with Reasoning

In the following zero-shot experiments, we utilize the PaLM 2-M LLM [2], a highly capable model trained on a broad set of languages and tasks, including reasoning tasks. Initially, we investigate the impact of prompting the model to engage in reasoning prior to prediction (zero-shot chain-of-thought) compared to direct prediction (zero-shot). Table 2 shows the differences between the final task description of the input prompt for zero-shot CoT and direct prediction. Subsequently, Table 3 outlines the outcomes of different zero-shot ablation studies. We also compare against a naive baseline, where we use the historical rating average of the user’s history as a prediction for the future item. We observe a notable performance improvement across both product domains when the model is guided to output reasoning alongside the prediction (“Our Method (zero-shot CoT)” vs. “No Reasoning Outputs”). This suggests that personalized tasks are inherently difficult for LLMs to solve without further guidance such as engaging in an intermediate reasoning step.

Impact of explicit user feedback. A user’s past purchase, denoted as $h_t = (\mathcal{M}_j, \mathbf{r}_{u,j}, \mathbf{d}_{u,j}) \in \mathcal{H}_u$, encapsulates the user-item relations, providing explicit user feedback. To understand the helpfulness of this feedback, we conduct ablation studies. In the first case, we eliminate the written reviews $\mathbf{d}_{u,j}$ from the purchase history where $h_t = (\mathcal{M}_j, \mathbf{r}_{u,j})$. In the second case, we further eliminate the numerical ratings $\mathbf{r}_{u,j}$ from h_t , resulting in $h_t = (\mathcal{M}_j)$. The first case mirrors a common scenario in RecSys where only numerical rating information is available, while the second case simulates scenarios where only implicit feedback from a user, in the form of past purchases, is accessible. Table 3 presents the ablation results, highlighting a significant performance drop when the review text is excluded from h_t (“No Review”). The performance declines further when both reviews and ratings are excluded (“No

Review, No Rating”). When only the written review text is removed, the results are similar to direct predictions made without reasoning (“No Review” vs “No Reasoning Outputs”) and sometimes even worse than naive average baseline. This indicates that review text is essential for utilizing the reasoning capabilities of LLMs. Without user reviews, the model lacks detailed insights into past user interactions and can only rely on numerical ratings, resulting in performance similar to or worse than “No Reasoning Outputs” and the naive average baseline.

Furthermore, when both reviews and ratings are unavailable, the outcomes are akin to random guessing, as evidenced by the multi-class accuracy hovering around 0.2. This performance is strictly worse than the naive baseline and direct prediction without the reasoning outputs. Our ablation studies suggest that while the LLM can estimate some user preference information simply from the numerical user rating, it benefits even more when we have the full written user reviews and a guided reasoning step. Review texts help discern nuanced details about a user’s specific preferences. For instance, a user might rate the movie “*Top Gun*” as 5 stars for various reasons, such as their interest in airplanes, admiration for Tom Cruise, or a fondness for action movies in general. These preferences may be explicitly stated in the review text, enabling the LLM to make more informed decisions.

Effect of Pre-trained Knowledge. In addition to analyzing the impact of excluding explicit user feedback, we also investigate the removal of item descriptions. We found that while performance decreases in both domains when item descriptions are unavailable, the decline is less pronounced for MOVIES/TV compared to BEAUTY. This finding suggests that the LLM (PaLM 2-M) possesses a more extensive knowledge base in MOVIES/TV domain, enabling it to infer information about the recommended movies or TV shows even in the absence of descriptions.

One-shot learning. We compare one-shot learning against our zero-shot results (Table 3). Surprisingly, the one-shot results are significantly worse, yielding only slight improvements over random performance. We believe the inclusion of a single example considerably increases the input text length. This could potentially hinder the model’s ability to disentangle information from the provided example and the actual inputs $(\mathcal{M}_i, \mathcal{H}_u)$, resulting in poorer performance.

4.3 Fine-tuning with Reasoning Data

We utilize Flan-T5 [4] models as they are readily available to conduct fine-tuning experiments. Although these models are all trained on a variety of data and tasks, it should be noted that PaLM-2-M reports significantly higher quality than Flan-T5 on common benchmarks, including the Massive Multitask Language Understand-

Table 3: Comparisons and ablation studies on zero-shot learning with PaLM 2-M, investigating the role of reasoning outputs and input features. “Our Method” denotes zero-shot chain-of-thought with reasoning output, while “No Reasoning Outputs” refers to the rating prediction-only task. Additionally, we examine the impact of removing user reviews, user ratings, and item descriptions from the input to the LLM. The “Naive Baseline” uses the historical rating average of the user as the prediction. We also include results for one-shot learning.

Method		Binary Acc.	Binary F1	Multi. Acc.	Multi. MAE ↓	Multi. RMSE ↓	ROUGE-1 F1	METEOR	BLEU	BERT Score
BEAUTY	Naive Baseline (Avg.)	0.52	0.60	0.25	1.35	1.75	-	-	-	-
	Our Method (zero-shot CoT)	0.56	0.62	0.37	1.14	1.60	0.236	0.503	0.339	0.665
	- No Reasoning Outputs	0.49	0.57	0.23	1.35	1.70	-	-	-	-
	- No Review	0.48	0.57	0.21	1.35	1.69	0.237	0.507	0.337	0.667
	- No Review, No Rating	0.43	0.53	0.19	1.42	1.75	0.215	0.494	0.331	0.660
	- No Item Description	0.48	0.57	0.21	1.33	1.66	0.235	0.504	0.340	0.667
	One-shot	0.43	0.57	0.26	1.52	1.97	0.225	0.502	0.335	0.664
MOVIES/TV	Naive Baseline (Avg.)	0.59	0.63	0.30	1.21	1.63	-	-	-	-
	Our Method (zero-shot CoT)	0.62	0.66	0.40	1.06	1.53	0.194	0.465	0.296	0.647
	- No Reasoning Output	0.59	0.63	0.29	1.18	1.56	-	-	-	-
	- No Review	0.58	0.63	0.28	1.20	1.58	0.173	0.452	0.291	0.641
	- No Review, No Rating	0.43	0.54	0.20	1.42	1.75	0.150	0.434	0.283	0.633
	- No Item Description	0.54	0.62	0.28	1.22	1.60	0.183	0.460	0.296	0.647
	One-shot	0.47	0.59	0.23	1.32	1.68	0.182	0.452	0.276	0.641

Table 4: Fine-tuning results on Flan-T5, comparing different model sizes and fine-tuning with and without reasoning (predict numerical rating only). We also show the XL model results without any fine-tuning. Without fine-tuning, the model was unable to follow instructions with reasoning, and therefore we only show results without reasoning.

Model Size		Reasoning	Binary Acc.	Binary F1	Binary AUC	Multi. Acc.	Multi. AUC	Multi. MAE ↓	Multi. RMSE ↓	BLEU	ROUGE-1 F1	METEOR	BERT Score
BEAUTY	Small	✓	0.62	0.53	0.65	0.30	0.63	1.35	1.84	0.225	0.499	0.342	0.663
	Base	✓	0.59	0.47	0.66	0.27	0.64	1.37	1.83	0.239	0.507	0.344	0.667
	Large	✓	0.64	0.59	0.67	0.33	0.65	1.26	1.73	0.240	0.506	0.343	0.666
	XL	✓	0.67	0.61	0.78	0.30	0.69	1.24	1.68	0.241	0.510	0.339	0.667
	XL	✗	0.55	0.61	0.74	0.28	0.67	1.31	1.75	-	-	-	-
	XL (no fine-tuning)	✗	0.55	0.40	0.56	0.22	0.56	1.64	2.09	-	-	-	-
MOVIES/TV	Small	✓	0.60	0.55	0.70	0.33	0.66	1.23	1.71	0.137	0.423	0.272	0.627
	Base	✓	0.65	0.59	0.72	0.34	0.68	1.18	1.65	0.153	0.438	0.279	0.634
	Large	✓	0.64	0.58	0.72	0.32	0.67	1.23	1.70	0.165	0.448	0.286	0.639
	XL	✓	0.65	0.61	0.72	0.34	0.67	1.17	1.64	0.165	0.449	0.286	0.643
	XL	✗	0.62	0.57	0.70	0.33	0.66	1.27	1.75	-	-	-	-
	XL (no fine-tuning)	✗	0.61	0.43	0.61	0.23	0.61	1.56	2.00	-	-	-	-

Table 5: Comparison of fine-tuning Flan-T5 XL model with multiple reasoning paths per user-item pair and with different filtering methods. PaLM 2-M zero-shot (no fine-tuning) results are included for comparison.

Samples		Filter	Binary Acc.	Binary F1	Binary AUC	Multi. Acc.	Multi. AUC	Multi. MAE ↓	Multi. RMSE ↓	BLEU	ROUGE-1 F1	METEOR	BERT Score
BEAUTY	1	None	0.67	0.61	0.78	0.30	0.69	1.24	1.68	0.241	0.510	0.339	0.667
	8	None	0.68	0.64	0.79	0.31	0.70	1.25	1.71	0.248	0.509	0.333	0.671
	8	5-class	0.54	0.61	0.63	0.28	0.60	1.32	1.74	0.248	0.510	0.329	0.670
	8	Binary	0.53	0.59	0.64	0.29	0.60	1.40	1.88	0.246	0.508	0.335	0.669
	8	1-off	0.61	0.61	0.71	0.30	0.63	1.28	1.75	0.247	0.336	0.510	0.671
	PaLM 2-M Zero-shot			0.56	0.62	-	0.37	-	1.14	1.60	0.236	0.503	0.339
MOVIES/TV	1	None	0.65	0.61	0.72	0.34	0.67	1.17	1.64	0.165	0.449	0.286	0.643
	8	None	0.63	0.58	0.72	0.35	0.67	1.23	1.75	0.171	0.446	0.285	0.642
	8	5-class	0.59	0.63	0.69	0.32	0.63	1.17	1.61	0.176	0.449	0.291	0.642
	8	Binary	0.60	0.64	0.71	0.33	0.66	1.28	1.78	0.175	0.443	0.288	0.641
	8	1-off	0.62	0.63	0.74	0.36	0.67	1.16	1.64	0.180	0.451	0.291	0.643
	PaLM 2-M Zero-shot			0.62	0.66	-	0.40	-	1.06	1.53	0.194	0.465	0.296

ing (MMLU) [16]. Unless specified, we employ the Flan-T5 XL model (3B parameters) to output reasoning followed by a final rating prediction. We fine-tune the models for 100,000 steps with a batch size of 64, a dropout rate of 0.25, and a learning rate of 1e-4. We report the evaluation metrics on the test set at the end of training. In Table 4 we compare two fine-tuned XL models: one that outputs reasoning and rating, and an-

other that outputs rating only. For our first experiment, we use only 1 reasoning sample for each user-item pair, *i.e.* $m = 1$ for Eq. (3), and refrain from further filtering to maintain consistent training sample numbers for a fair comparison. We collect model probabilities by extracting and normalizing the logits corresponding to the 5 rating class tokens, allowing us to compute ROC-AUC metrics. The results from both domains con-

sistently indicate that fine-tuning models to engage in reasoning prior to making predictions leads to improved performance across all metrics.

In addition to comparing the XL models with reasoning outputs against those without, we fine-tune various Flan-T5 model sizes: Small (80M params), Base (250M), Large (780M), and XL (3B). All models are trained with one reasoning sample per example without any filtering. Table 4 illustrates a clear trend indicating that larger models perform better on the rating task. This observation aligns with the general understanding that larger models typically possess greater knowledge capacity, leading to enhanced performance on downstream tasks after fine-tuning. Moreover, larger models also tend to have better reasoning quality, which we discuss further in Section 5.2.

We also include results on the Flan-T5 XL model without any fine-tuning for reference, showing that fine-tuning is absolutely necessary to improve results for this model family. For this non fine-tuned model, we tried both with and without an additional reasoning output. When asked to output reasoning, the model was unable to follow instructions and did not output a final rating in almost all cases. Therefore, we only report no fine-tuning results without reasoning. Additionally, although we see drastic improvements for Flan-T5, our best Flan-T5 result still underperforms our best PaLM 2-M result without fine-tuning. This difference is likely attributable to the enhanced capabilities of PaLM 2-M relative to Flan-T5.

Training with multiple reasoning paths. In this experiment, we increase the number of reasoning samples m to 8 for each user-item pair, as defined in Eq. (3), providing diverse reasoning paths for each user-item relation. We also apply different filtering methods based on comparing the zero-shot LLM rating predictions $\hat{\mathbf{r}}_{u,i}^m$ to the ground truth ratings $\mathbf{r}_{u,i}$. Table 5 presents the results of training with multiple reasoning paths and with different filtering methods:

- **None:** No filter is applied. The model is trained with all reasoning samples $\hat{\mathbf{r}}_{u,i}^m$ per user-item pair.
- **5-class:** Reasoning samples are removed where $\hat{\mathbf{r}}_{u,i}^m \neq \mathbf{r}_{u,i}$ in terms of the 5-class rating.
- **Binary:** Reasoning samples are removed where the binary conversion of $(\hat{\mathbf{r}}_{u,i}^m > 3) \neq (\mathbf{r}_{u,i} > 3)$.
- **1-off:** Reasoning samples are removed where the absolute difference $|\hat{\mathbf{r}}_{u,i}^m - \mathbf{r}_{u,i}| > 1$.

In the BEAUTY domain, fine-tuning with 8 reasoning paths without any filtering slightly outperforms fine-tuning with only 1 reasoning path. Surprisingly, applying filtering methods significantly diminishes performance on the rating task. We hypothesize that filtering may remove a substantial portion of training samples, leading to poorer performance. This is particularly evident when the "5-class" filter, the most stringent filter, is applied. Conversely, in the MOVIES/TV domain,

the best results are achieved with the "1-off" filtering method. We attribute this to the LLM’s strong pre-trained knowledge in the MOVIES/TV domain, allowing it to tolerate the removal of examples where the reasoning does not align with the ground truth rating. In contrast, the BEAUTY domain may require more examples of user history and user-item relations for effective learning. Removing examples, even those with misaligned reasoning, may inadvertently reduce domain information, resulting in diminished performance. Further analysis of these effects will be conducted in future work.

5 Reasoning Evaluation

5.1 Human Judgment Alignment Analysis

As proposed in Sec. 3.2, we design a human judgement alignment study to evaluate the effectiveness of RECSAVER. We presented a total of 100 samples to human raters, with 50 examples from the BEAUTY domain and 50 examples from the MOVIES/TV domain. Each rating category was evenly represented. Each sample was annotated by 3 different annotators, resulting in a total of 300 annotated data points. Table 6 presents the inter-annotator agreement of weighted Cohen κ [6] and the average Pearson correlation (Avg. ρ) among the human annotated scores. The achieved statistical significance of the average correlation between 3 human annotators across all 3 measurements, as indicated by the p -values, signifies the level of consensus among annotators.

We evaluated four commonly used natural language generation (NLG) metrics: BLEU [23], ROUGE-1 [20], METEOR [3], and BERTScore [33]. BLEU and ROUGE-1 measure syntactic similarity by computing the exact n -gram overlap between the generated output and the reference texts. On the other hand, METEOR and BERTScore consider semantic similarity, providing a more comprehensive evaluation by incorporating contextual information. Table 7 reveals a consistently positive correlation between coherence and all NLG metrics, suggesting that our proposed evaluation method using LLM-generated references aligns well with coherence. However, insightfulness exhibits no correlation with BLEU and a low correlation with ROUGE-1 F1, while demonstrating a slightly positive correlation with METEOR and BERTScore. Unlike “coherence” and “faithfulness”, “insightfulness” is an exploratory metric aimed at understanding how LLM *surprise* human raters. It is anticipated that syntactic metrics such as BLEU and ROUGE-1 F1 may not correlate strongly with insightfulness. For instance, a response could be considered insightful even if it lacks significant n -gram overlap with the references. Although the semantic metrics METEOR and BERTScore show better correlation, they still do not align as closely with insightfulness as they do with coherence.

Table 6: Inter-annotator agreement (IAA) analysis on the human annotated scores.

	Mean	Cohen κ	Avg. ρ	p -value
Coherence	3.72	0.37	0.37	1e-10
Faithfulness	0.63	0.63	0.63	1e-12
Insightfulness	2.80	0.33	0.34	6e-4

Table 7: Correlation between coherence, insightfulness, and automatic NLG metrics. The annotated scores are averaged across the annotators for each sample.

	Coherence	Insightfulness
BLEU	0.36	0.02
ROUGE-1 F1	0.40	0.10
METEOR	0.40	0.25
BERTScore	0.36	0.20

Two-sample T-test. We conducted a two-sample t-test to compare faithful and unfaithful reasoning. Faithful reasoning refers to outputs without factual or logical errors, while unfaithful reasoning refers to outputs containing one or more such errors. Table 8 shows that when errors are present, the reasoning responses are less coherent and insightful. Additionally, the average scores for all automatic NLG metrics are higher for faithful reasoning compared to unfaithful reasoning. Syntactic metrics (BLEU and ROUGE-1 F1) exhibit a more pronounced difference than semantic metrics (METEOR and BERTScore), as indicated by the lower p -values. This discrepancy arises because small differences in n -grams can lead to factual errors. For example, changing “the user purchased 4 products” to “the user purchased 5 products” can render a sentence unfaithful. Such small discrepancies are not as easily detected by the semantic metrics.

Table 8: Two-sample t-test comparing the average of human annotated scores and NLG scores between faithful and unfaithful reasoning.

	Faithful	Unfaithful	p -value
Coherence	4.01	3.22	2e-8
Insightfulness	3.11	2.23	6e-9
BLEU	0.21	0.16	2e-3
ROUGE-1 F1	0.49	0.46	5e-3
METEOR	0.31	0.30	0.36
BERTScore	0.65	0.63	0.02

Effectiveness of Self-verification. To validate the effectiveness of the self-verification step in Rec-SAVER, we compared the metrics computed with and without the reference self-verification step. Table 9 shows that metrics computed from self-verified references show a stronger correlation with the coherence score compared to those without self-verification, indicating that self-verification contributes to increasing the credibility of LLM-generated references. Combining all results of this section, we observe a strong alignment between our proposed Rec-SAVER reasoning evaluation method and

human judgments regarding quality.

Table 9: Correlation between coherence and NLG metrics with and without using self-verified references.

Self-verification	Yes	No
BLEU	0.36	0.33
ROUGE-1 F1	0.40	0.35
METEOR	0.40	0.37
BERTScore	0.36	0.28

5.2 Analysis of Reasoning Quality

Having established Rec-SAVER as a method for generating references and evaluating reasoning, we can leverage it to further analyze model reasoning outputs. We display example reasoning outputs in Table 10. Our first investigation focuses on the question: “Are correct rating predictions generally associated with higher-quality reasons?” Table 11 demonstrates that reasoning metrics indeed improve for examples with correct predictions, both for zero-shot and fine-tuned models. Next, we analyze the reasoning quality across different methods discussed throughout the paper. Focusing on zero-shot models, we observe that in the MOVIES/TV domain, the NLG metrics decrease when we remove input information (Table 3). However, for BEAUTY, when we remove some input information like user reviews, some of the NLG metrics increase. This suggests that the LLM may encounter challenges synthesizing all of the information in BEAUTY but possesses a better overall knowledge of MOVIES/TV, allowing it to generate better reasons when provided with more information available.

Table 4 illustrates that beyond the rating prediction metrics, the NLG metrics also improve as we increase the fine-tuned model size. This suggests that in addition to producing better rating predictions, the models also generate better reasoning responses. In the MOVIES/TV domain, the “1-off” filtering method appears to yield the best rating metrics (Table 5), although a few other methods are comparable. However, when considering the NLG metrics, we observe more data showing that the “1-off” filtering method has an advantage over the other methods.

The comparison between the fine-tuned models and the zero-shot model for MOVIES/TV reveals that the zero-shot model outperforms the fine-tuned model in both rating and NLG metrics. This outcome suggests that the PaLM 2-M zero-shot model likely possesses superior pre-trained knowledge in this domain, which cannot be fully distilled into the Flan-T5 fine-tuned model. For example, the reasoning text used in the fine-tuning training data may include mentions of certain actors or directors in movies. However, this data might not cover a wide enough range of examples to provide the model a comprehensive understanding of the entire domain. Consequently, the final fine-tuned model may still exhibit information gaps in this domain compared

Table 10: Example reasons generated by PaLM 2-M in a zero-shot setting.

BEAUTY	The user has purchased 4 beauty products in the past, all of which are highly rated. The new product is also a beauty product, so it is likely that the user will be interested in it. The new product is a concealer, which is a type of makeup that is used to cover blemishes and imperfections. The user has not purchased any makeup products in the past, so it is possible that they are not familiar with concealers. However, the description of the new product states that it is “natural coverage” and “non-nonsense”, which suggests that it is a good choice for beginners. Overall it is likely that the user will rate the new product positively.
MOVIES/TV	The user has given high ratings to movies that are inspiring and thought-provoking, and/or a history lesson. The new video is a historical movie about a group of black flyers who broke the color barrier in the U.S. Air Force during World War II. It is likely that the user will find this movie inspiring and thought-provoking, and therefore rate it highly.

Table 11: Reasoning quality associated with correct and incorrect rating predictions for PaLM 2-M zero-shot and Flan-T5 XL fine-tuned (1 sample per example, no filtering) models.

	Model	Correct Prediction	BLEU	ROUGE-1 F1	METEOR	BERT Score
BEAUTY	PaLM 2-M	Yes	0.260	0.522	0.342	0.666
	PaLM 2-M	No	0.221	0.491	0.336	0.665
	Flan-T5 XL	Yes	0.254	0.515	0.342	0.667
	Flan-T5 XL	No	0.235	0.508	0.338	0.666
MOVIES/TV	PaLM 2-M	Yes	0.204	0.480	0.306	0.648
	PaLM 2-M	No	0.187	0.455	0.290	0.647
	Flan-T5 XL	Yes	0.177	0.457	0.292	0.644
	Flan-T5 XL	No	0.159	0.444	0.283	0.642

to the zero-shot model.

6 Related Work

LLM for RecSys. Recent advancements in the application of LLMs to RecSys have yielded diverse approaches. Typically, these approaches follow *pre-training*, *fine-tuning*, and *prompting* paradigms. In the context of recommendation tasks, fine-tuning LLMs is essential to acquire domain-specific knowledge. This fine-tuning process involves training the pre-trained model using task-specific recommendation datasets containing user-item interaction behaviors such as purchase, ratings, or click, and additional contextual information about users and items (*e.g.*, social relations or item descriptions) [10, 7, 21]. Beyond the pre-training and fine-tuning paradigm, prompting has emerged as a recent paradigm to tailor LLMs for specific downstream tasks, employing task-specific prompts [27, 9, 12], along with prompting techniques like in-context learning [11] and chain-of-thought [30]. More recently, there has been exploration into instruction tuning, a hybrid approach combining the pre-training and fine-tuning paradigm with prompting. This involves fine-tuning LLMs across multiple recommendation tasks using instruction-based prompts, enhancing the zero-shot performance of LLMs on previously unseen RecSys tasks [13, 32, 18].

LLM reasoning. Recent studies have suggested that the ability to reason may emerge in language models at a certain scale [5, 29]. These models, when provided with a few examples of “chain of thought”, which represent intermediate natural language reasoning steps, demonstrate the capability to generate explicit rationales before producing final answers [30]. Advances in this

direction include zero-shot CoT [19], where the model is prompted with the phrase “Let’s think step by step” to elicit reasoning without the inclusion of few-shot demonstrations. Various strategies have been proposed to enhance language model performance by prompting reasoning, such as multi-step reasoning [8, 34], tree-of-thoughts [31], iterative CoT prompting [26] and self-consistency [28]. Despite the impressive performance of LLMs on various reasoning tasks, the clarity of whether their predictions are based on true reasoning remains a challenge. This ambiguity arises because most existing evaluations focus on accuracy in end tasks rather than directly assessing the quality of the reasoning. Recent efforts have introduced metrics such as ROSCOE [15] and dataset such as PrOntoQA [24] for a more formal analysis of reasoning in LLMs. However, the application of these metrics and benchmarks to a broader range of tasks is still an area of limited depth.

7 Conclusion and Discussion

We explore reasoning in the context of personalized recommender systems, showing that adding reasoning steps can improve LLM task performance. It is important to have rich user context and explicit feedback in order for the LLMs to reason adequately. Having good pre-trained domain knowledge is also useful. **RecSAVER**, our proposed method for analyzing reasoning quality, aligns well with human judgment on the coherence of reasoning outputs and can be used to further evaluate model quality beyond numerical task results.

Limitations. In this work we started with rating predictions in the Amazon review dataset for two categories, BEAUTY and MOVIES/TV. However, the extent of recommender systems is vast. It is unclear to what extent our methods generalize more broadly to other categories such as music, video games, website articles, etc. Furthermore, more work is needed to explore these methods on other tasks, including candidate retrieval or ranking.

Now that we see evidence that reasoning is useful in RecSys, more work should be done to understand the extent and mechanisms behind this. Does the LLM actually reason in a manner that helps make a final decision similar to human thought? Or is there some other underlying procedure that yields these results, such as more overall computation or better attention? Future work

looking at different prompting strategies and reasoning plans could help uncover more details in this area.

Ethical Considerations. In this study, biases may exist for reasoning results for different users, including users that speak different languages or users with different genders. The dataset we use focuses on users that speak English. Also, users from different genders may interact differently with certain domains or with products in those domains, leading to skewed distributions in the data. Broader experiments are needed to understand these potential biases further in the context of reasoning for recommender systems.

8 Acknowledgements

We would like to thank Jianmo Ni, Nikhil Mehta, Ramkumar Rajendran, and Lakshmi Chakrapani for their helpful discussions and support.

References

- [1] Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. 2019. [MathQA: Towards interpretable math word problem solving with operation-based formalisms](#). 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)*, pages 2357–2367, Minneapolis, Minnesota. Association for Computational Linguistics.
- [2] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. 2023. Palm 2 technical report. *arXiv preprint arXiv:2305.10403*.
- [3] Satyanjeev Banerjee and Alon Lavie. 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- [4] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#).
- [5] Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- [6] Jacob Cohen. 1968. [Weighted kappa: nominal scale agreement with provision for scaled disagreement or partial credit](#). *Psychological Bulletin*, 70(4):213–220.
- [7] Zeyu Cui, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022. M6-rec: Generative pretrained language models are open-ended recommender systems. *arXiv preprint arXiv:2205.08084*.
- [8] Dheeru Dua, Shivanshu Gupta, Sameer Singh, and Matt Gardner. 2022. [Successive prompting for decomposing complex questions](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1251–1265, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- [9] Fernando Ferraretto, Thiago Laitz, Roberto Lotufo, and Rodrigo Nogueira. 2023. Exaranker: Synthetic explanations improve neural rankers. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2409–2414.
- [10] Luke Friedman, Sameer Ahuja, David Allen, Terry Tan, Hakim Sidahmed, Changbo Long, Jun Xie, Gabriel Schubiner, Ajay Patel, Harsh Lara, et al. 2023. Leveraging large language models in conversational recommender systems. *arXiv preprint arXiv:2305.07961*.
- [11] Tianyu Gao, Adam Fisch, and Danqi Chen. 2020. Making pre-trained language models better few-shot learners. *arXiv preprint arXiv:2012.15723*.
- [12] Yunfan Gao, Tao Sheng, Youlin Xiang, Yun Xiong, Haofen Wang, and Jiawei Zhang. 2023. Chat-rec: Towards interactive and explainable llms-augmented recommender system. *arXiv preprint arXiv:2303.14524*.
- [13] Shijie Geng, Shuchang Liu, Zuohui Fu, Yingqiang Ge, and Yongfeng Zhang. 2022. Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5). In *Proceedings of the 16th ACM Conference on Recommender Systems*, pages 299–315.
- [14] Mor Geva, Daniel Khoshdel, Elad Segal, Tushar Khot, Dan Roth, and Jonathan Berant. 2021. Did aristotle use a laptop? a question answering benchmark with implicit reasoning strategies. *Transactions of the Association for Computational Linguistics*, 9:346–361.
- [15] Olga Golovneva, Moya Chen, Spencer Poff, Martin Corredor, Luke Zettlemoyer, Maryam Fazel-Zarandi, and Asli Celikyilmaz. 2022. Roscoe: A suite of metrics for scoring step-by-step reasoning. *arXiv preprint arXiv:2212.07919*.
- [16] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the MATH dataset](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- [17] Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. In *International Conference on Learning Representations*.
- [18] Wang-Cheng Kang, Jianmo Ni, Nikhil Mehta, Maheswaran Sathiamoorthy, Lichan Hong, Ed Chi, and Derek Zhiyuan Cheng. 2023. Do llms understand user preferences? evaluating llms on user rating prediction. *arXiv preprint arXiv:2305.06474*.
- [19] Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:22199–22213.
- [20] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.

- [21] Junling Liu, Chao Liu, Peilin Zhou, Qichen Ye, Dading Chong, Kang Zhou, Yueqi Xie, Yuwei Cao, Shoujin Wang, Chenyu You, et al. 2023. Llmrec: Benchmarking large language models on recommendation task. *arXiv preprint arXiv:2308.12241*.
- [22] 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)*, pages 188–197.
- [23] 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*, pages 311–318.
- [24] Abulhair Saparov and He He. 2022. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. *arXiv preprint arXiv:2210.01240*.
- [25] Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. CommonsenseQA: A question answering challenge targeting commonsense knowledge. 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)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- [26] Boshi Wang, Xiang Deng, and Huan Sun. 2022. Iteratively prompt pre-trained language models for chain of thought. *arXiv preprint arXiv:2203.08383*.
- [27] Xiaolei Wang, Kun Zhou, Ji-Rong Wen, and Wayne Xin Zhao. 2022. Towards unified conversational recommender systems via knowledge-enhanced prompt learning. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1929–1937.
- [28] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*.
- [29] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022. Emergent abilities of large language models. *Transactions on Machine Learning Research*. Survey Certification.
- [30] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837.
- [31] Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik R Narasimhan, and Yuan Cao. 2023. React: Synergizing reasoning and acting in language models. In *The Eleventh International Conference on Learning Representations*.
- [32] Junjie Zhang, Ruobing Xie, Yupeng Hou, Wayne Xin Zhao, Leyu Lin, and Ji-Rong Wen. 2023. Recommendation as instruction following: A large language model empowered recommendation approach. *arXiv preprint arXiv:2305.07001*.
- [33] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- [34] Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei, Nathan Scales, Xuezhi Wang, Dale Schuurmans, Claire Cui, Olivier Bousquet, Quoc V Le, and Ed H. Chi. 2023. Least-to-most prompting enables complex reasoning in large language models. In *The Eleventh International Conference on Learning Representations*.

A Reasoning Generation Prompt

Table 12: Reasoning generation prompt used in our zero-shot and fine-tuned experiments.

Here is information about a user and a new {product / video (movies and tv)} being recommended to the user. For the user, we have the user’s past item information history and the user’s corresponding ratings. User ratings range from 1 to 5, where 1 is the lowest and 5 is the highest. For the new item being recommended, we have the item information.

```
### Past User History: ###
{Product / Video (Movies and TV)} Title: {title}
Brand: {brand}
Categories: {categories}
Description: {description}
Item Price: {price}
User Rating: {userRating}
User Review: {reviewText}
.
.
.
```

```
### New Item Information: ###
New {Product / Video (Movies and TV)}
{Product / Video (Movies and TV)} Title: {title}
Brand: {brand}
Categories: {categories}
Description: {description}
Item Price: {price}

#####
```

Given the user’s past {purchase / watch} history and the new item information, what information can you infer about the user’s preferences and how they will rate the new {product / video (movies and tv)}?

Your reasoning explanation should be based on any commonalities in the user history items and inferred user tastes or preferences.

After your reasoning, predict a numerical rating.

Please follow the format below:

```
### Reason ###
Write your reasoning explanation here. You can have line breaks.
```

```
### Rating ###
Give a single numerical rating, e.g. 1
```

B Additional Experimental Results

The weighted Cohen κ is calculated as followed:

$$\kappa = 1 - \frac{\sum_{i=1}^n \sum_{j=1}^n w_{ij} x_{ij}}{\sum_{i=1}^n \sum_{j=1}^n w_{ij} m_{ij}},$$

where $n = 5$ is the number of rating scale, and w , x and m are elements in the weight, observed, and expected matrices. Here, we use a quadratic weight where $w_{ij} = \frac{(i-j)^2}{(k-1)^2}$ to amplify the difference between scores.

Table 13 provides a comparison of human-rated reasoning outputs between incorrect and correct predictions, as depicted in Figure 5. The table shows the mean

scores for coherence, faithfulness, and insightfulness for reasoning outputs associated with incorrect and correct predictions. We observe that reasoning outputs corresponding to correct predictions receive higher scores across all three dimensions, indicating a higher reasoning quality when the prediction is correct.



Figure 5: Outputs reasons categorized based on the correctness of rating predictions.

Table 13: Comparison of human-rated reasoning outputs between incorrect and correct predictions. Higher scores indicate higher reasoning quality.

	Incorrect Prediction	Correct Prediction
Coherence	3.59	3.91
Faithfulness	0.61	0.67
Insightfulness	2.71	2.93
Types of Errors		
Incorrect Product Statistics	28%	21%
Incorrect Product Information	14%	16%
Arithmetic Errors	5%	3%
Others	2%	1%

Table 14: Example output generated by the fine-tuned FLAN-T5 XL model.

```
### Reason ###
The user has a history of watching action movies, especially those with a sci-fi or fantasy element. The new video is an action with a Batman theme, so it is likely to appeal to the user.

### Rating ###
5
```

Table 15: Comparing NLG metric statistics for a fine-tuned FLAN XL model and a zero-shot model in BEAUTY.

		ROUGE-1 F1	BLEU	METEOR	BERTScore
Fine-Tuned	Mean	0.509	0.248	0.333	0.671
	Min	0.256	0.028	0.163	0.524
	Max	0.804	0.771	0.575	0.816
	Range	0.548	0.743	0.412	0.292
	Std-Dev	0.087	0.131	0.073	0.046
Zero-Shot	Mean	0.506	0.245	0.332	0.665
	Min	0.088	0.010	0.110	0.480
	Max	0.852	0.740	0.772	0.838
	Range	0.765	0.730	0.662	0.359
	Std-Dev	0.090	0.132	0.076	0.047