DaLimeLlama: A Persuasive XAI Framework For Learning-to-Rank Product Recommenders

Sebastian Lubos[1], Seda Polat Erdeniz[1], Thi Ngoc Trang Tran[1], Alexander Felfernig[1], Ilhan Adıyaman[2], Ata Gür[2] and Tevfik Ince[2]

[1]Institute of Software Technology, Graz University of Technology, Inffeldgasse 16b, 8010 Graz, Austria
[2]Frizbit, Carrer del Doctor Trueta 183, 08005 Barcelona, Spain

Abstract
As e-commerce platforms evolve, the need for effective recommender systems becomes increasingly crucial. However, traditional recommendation models often lack transparency and interpretability, hindering user trust and satisfaction. In this paper, we present a novel approach to address this challenge by employing the *Local Interpretable Model-agnostic Explanations (LIME)* library in combination with the *Large Language Model (LLM) Llama* to generate persuasive explanations for recommendation outputs produced by a *learning-to-rank* algorithm. Our method aims to enhance the understandability and persuasiveness of recommendation explanations, thereby improving user engagement and satisfaction. We tested our approach on a real-world e-commerce dataset and conducted comprehensive experiments to assess its effectiveness. The results, including a preliminary user study, demonstrate that integrating LIME and Llama improves the understandability, transparency, and persuasiveness of recommendation explanations compared to XAI baseline methods. Our research contributes to advancing persuasive and explainable AI in e-commerce by offering a practical solution.

Keywords
Persuasive XAI, Large Language Models, Recommender Systems, Explanations, Learning-to-rank

1. Introduction

Learning-to-rank (LTR) product recommenders [1, 2] find applications across various domains where the ranking of items significantly influences user satisfaction and business performance. In e-commerce platforms, they enhance the relevance of search results by adapting to user preferences. Similarly, in content recommendation systems, LTR improves the ranking of suggested content, thereby enhancing user engagement and retention. Job portals benefit from LTR by matching job seekers with relevant listings based on individual skills and preferences. Travel websites and financial services utilize LTR to personalize recommendations, ensuring optimal ranking of options like hotels, flights, loans, or insurance policies. Additionally, in online advertising, LTR optimizes ad ranking, maximizing performance and revenue by displaying ads that are most relevant to users.

However, standard methods to explain recommendations are not applicable to LTR recommenders. Since LTR recommenders are usually employing classification machine learning models [3] such as *Random Forest* [4] or *XGB* [5], *explainable artificial intelligence (XAI)* methods such as LIME or SHAP [6] are suitable to explain the results of LTR. However, these explanations only highlight the most important features of a specific recommendation result. While this can be useful for domain experts, it is often unintuitive and ineffective in e-commerce, where users are unfamiliar with the underlying machine learning models and their features. In these cases, users require understandable natural language explanations. Moreover, users need to be convinced of the recommended items, for example, by using sales strategies, including persuasion phrases.

Persuasive XAI [7, 8, 9, 10] holds significant importance in e-commerce due to its potential to enhance user experience, increase sales, and build trust between consumers and AI-driven systems. State-of-the-art methods in persuasive XAI [11, 12, 13, 14, 15] encompass a range of sophisticated approaches to improve user engagement and influence behavior. Personalization techniques lie at the forefront, tailoring explanations and persuasive strategies to individual user characteristics and decision-making styles. These methods often leverage insights from behavioral economics and psychology, aligning explanations with common cognitive biases to increase their effectiveness [16]. Interactive interfaces play a significant role, allowing users to explore recommendations and explanations in-depth, thereby increasing understanding and trust. Emotional engagement is another key aspect, with methods incorporating storytelling, emotive language, and visual metaphors to evoke strong emo-

The 1st Workshop on Risks, Opportunities, and Evaluation of Generative Models in Recommender Systems (ROEGEN@RecSys 2024), October 2024, Bari, Italy.
slubos@ist.tugraz.at (S. Lubos); sedapolat@gmail.com (S. P. Erdeniz); ttrang@ist.tugraz.at (T. N. T. Tran); alexander.felfernig@ist.tugraz.at (A. Felfernig)
0000-0002-5024-3786 (S. Lubos); 0009-0006-5106-3986 (S. P. Erdeniz); 0000-0002-3550-8352 (T. N. T. Tran); 0000-0003-0108-3146 (A. Felfernig)

tional responses from users. Additionally, explainable recommender systems integrate XAI techniques to generate transparent and persuasive recommendations, enhancing user engagement and conversion rates. Ethical considerations are pervasive, with the development of frameworks ensuring that persuasive XAI methods adhere to principles of transparency, fairness, and user autonomy [12]. Robust evaluation metrics, encompassing user engagement, trust, satisfaction, and behavioral outcomes, are essential for assessing the effectiveness and impact of persuasive XAI interventions across various domains, including e-commerce, healthcare, finance, and education. As research in persuasive XAI evolves, these state-of-the-art methods foster more transparent, effective, and ethically sound AI-driven persuasion strategies.

A promising way to advance this research is the integration of *Large Language Model (LLM)* [17] to improve explainability. Models like *GPT-4* [18] and *Llama 3* [19, 20], known for their impressive natural language generation capabilities, could help address some challenges in XAI for recommenders. LLMs are advanced natural language processing models that respond to specific input *prompts* to perform various tasks. Those prompts include instructions guiding the model to generate the desired output [21, 22]. Different possibilities to enhance recommendations using LLMs have already been indicated in various studies [23, 24], including approaches to explain recommendations using their extensive encoded knowledge and natural language abilities [25, 26]. Based on these promising findings, we expect LLMs to support LTR product recommenders by providing users with understandable and persuasive explanations.

In this paper, we address the research question, "*How can we integrate persuasive explainable AI with learning-to-rank recommender systems?*" To this end, we propose *DaLimeLlama* as a novel method combining XAI and LLMs to provide understandable explanations for recommended products. We explain our proposed approach using a working example and discuss the results from a real-world e-commerce dataset and a preliminary user study to explore how the approach can enhance the explainability of LTR recommenders[1].

2. Proposed Method

To employ persuasive XAI in e-commerce, we propose a novel method by combining the powers of XAI and LLMs, specifically using the *LIME* library [27] and the *Llama 3* LLM [19, 20]. *LIME (Local Interpretable Model-agnostic Explanations)* [27] is a widely used XAI technique used to provide interpretable explanations for individual predictions made by machine learning models. LIME offers insights into why a particular model output was generated by approximating the model's behavior locally around the prediction of interest. *Llama* [19] is a collection of powerful open-source LLMs with different model sizes, reaching state-of-the-art performance in competitive benchmarks[2]. In the scope of this paper, we used the recent *Llama 3* model version [20] with 70 billion parameters, fine-tuned to complete chat responses[3].

Our proposed *DaLimeLlama* framework consists of three stages: *classification stage*, *explanation stage*, and *persuasion stage*. In the following, we explain our suggested approach using a small recommendation scenario to illustrate its functionality. In this scenario, we used real-world commercial transactions involving cosmetic products to generate recommendations and explanations. Detailed information about the dataset and the methods applied are provided subsequently. The implementation of the approach using the dataset can be accessed in our GitHub repository[4].

Step 0: Data Preparation The dataset for our experiments was obtained from an e-commerce company selling high-quality cosmetics. It contains two days of session data, including users' purchased products and a product table with product features. In total, the dataset includes 78 sessions and 850 products. A subset of product features (see Figure 1) was used to compute recommendations and enrich the session data (see Figure 2).

Step 1: Classification Stage The first stage applies a machine learning model to solve the recommendation as a classification problem. We used *Learning-to-rank (LTR)* [28] for this purpose using the *XGBRanker* implementation of the *XGBoost* library[5]. LTR is a machine learning algorithm trained to rank a list of products based on their relevance to users' preferences or needs. In the context of e-commerce, LTR for product recommendations involves predicting the order in which products should be presented to a user, aiming to maximize user satisfaction, engagement, and conversion rates.

Step 2: Explanation Stage The second stage uses LIME to generate local explanations for a given session input and recommended product. *Local explanations* provide detailed insights into the model's decision-making process for individual predictions. Specifically, the explanations highlight the degree of influence product and session features have on the predicted score of the recommended item, helping to understand why the model made a particular recommendation in that specific instance. Table 2 shows example LIME explanations.

[1] A short video presentation is available on YouTube: https://youtu.be/7IVD59x1Jvo

[2] https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard

[3] We used the hosted model version *meta/meta-llama-3-70b-instruct* on *Replicate*

[4] https://github.com/AIG-ist-tugraz/DaLimeLlama

[5] https://xgboost.readthedocs.io/en/stable/tutorials/learning_to_rank.html

product_id	total purchases of this product	product price	product is on sale	product is a best seller	discount on product price	primary product category	product subcategory	product category	product collection	product by gender
01BC500K20	0	30.0	false	false	0.0	Body	not available	Foaming Bath	Karité	female
01BK010K14	23	9.0	false	true	0.0	Body	not available	Shower Gel	Almendra	female
15LM300L21	11	29.0	false	false	0.0	Hands	not available	Hand cream	Lavanda	female

Figure 1: Subset of the product details in the dataset, showing the product features considered for the recommendation.

session_id	preference of sale products	preference of best seller products	average price per product in purchase	average discount per product in purchase	preferred primary product category	preferred product subcategory	preferred product category	preferred product collection	preference of female products	preference of male products
5792f3be…	false	false	47.0	0.0	Perfume	not available	Women's perfumes	Forgotten Flowers	true	false
4c21880b…	true	false	55.0	0.13	Hands	not available	Hand cream	Les Classiques	true	true
f6d66abf…	false	true	21.67	0.0	Body	not available	Shower Gel	Karité	true	false

Figure 2: Subset of the enriched session details, showing the session features considered for the recommendation.

Step 3: Persuasion Stage The third stage forms an important part of our approach by making the LIME output more understandable using the extensive background knowledge of the LLM to transform it into a natural language output. The required prompt includes instructions on the task and expected output and instructs the LLM to use a persuasive explanation style. Furthermore, it includes the specific task description for a product recommendation and the LIME output. The prompt template is shown in Listing 1. Examples of generated explanations using this prompt are shown in Table 2.

3. Discussion

In Table 2, we provide three recommended products with explanations generated by LIME and our method, *DaLimeLlama*. For initial insights into the user perception of DaLimeLlama explanations, we conducted a small online user study with 8 participants using Google Forms. Using a within-subject design, each participant received product recommendations, including product descriptions and two explanations (LIME baseline and DaLimeLlama). After reading, participants were asked to rate[6] the understandability of those explanations, whether they helped understand why the product was recommended (transparency), and if they were convincing (persuasiveness). The mean ratings (see Table 1) were consistently higher for DaLimeLlama in all three categories.

As observed in the second column of Table 2, LIME provides the most important features and their influence direction specific to this recommendation. For example, the first LIME explanation indicates that this product is recommended, among other aspects, because it was purchased 13 times, suggesting its relative popularity. However, this information is only provided with the label of the most important feature "*total purchases of this product*" and its respective value "*13*". For standard users, this explanation is not intuitively interpretable, and even domain experts might not easily benefit from this kind of explanation. For instance, if LIME is used in the healthcare domain to explain the suggested recommendation of a prognosis prediction classifier model to clinicians, it is inefficient to use this format, e.g., by showing a laboratory result value "*Hemoglobin = 24 g/dl*", to explain the results.

[6] using a 5-point Likert scale

	LIME	DaLimeLlama
Persuasiveness	1.8	4.1
Transparency	2.2	4.0
Understandability	2.2	4.4

Table 1
Mean ratings of explanation characteristics for baseline LIME and DaLimeLlama explanations.

In the last column of Table 2, we demonstrate the explanations of our proposed method for the same recommendation. DaLimeLlama provides a natural language explanation of the most important features of the recommendation result as identified by LIME. For instance, for the first recommendation, DaLimeLlama describes the most important feature calculated by LIME ("*total purchase of the product = 13*") as: "*Our recommendation is based on the fact that many customers have purchased this product before, with a total of 13 purchases.*" For an e-commerce user, this sentence is more intuitively understandable than the feature and value itself, as also indicated by our preliminary user study.

Moreover, DaLimeLlama uses persuasive language to motivate users to accept recommendations. The prompt template (see Listing 1) instructs the LLM to be persuasive without specifying which principles to follow. Taking Cialdini's persuasion principles [29] as the basis, we see that the examples primarily follow the *consistency* principle, highlighting the agreement between previous preferences and the recommended product. This is illustrated in specific phrases, such as "*We also know that you tend to prefer products within a certain price range and this soap falls comfortably within that range*", "*you don't seem to prioritize best-seller products, which is great because this soap is a hidden gem*", or "*Since you've purchased a similar product before, we know you appreciate high-quality body care*". Additionally, at least in the first example, the *social proof* principle is used as well by emphasizing the purchases of other customers, i.e., "*Our recommendation is based on the fact that many customers have purchased this product before.*" The results of the user study suggest that the explanations are perceived as persuasive. However, as different persuasion principles might work

```
# System instructions
You are an explanation generator for recommended beauty products.
Recommendations are computed based on individual user preferences.
Your task is to explain why a recommended product is a good fit for a given user.
You will receive the LIME output and should use it exclusively to create your explanation.
LIME is a technique ...
You must only use the LIME information for your explanation.
The explanation should be persuasive and address the user directly.
Formatting instructions: ...
# Request
The following product was recommended: '{product}'.
The most important features of the product and user preferences for this recommendation are: {features}
The LIME output is the following: {lime_result}
Generate an explanation why the recommended product is a good fit for the user, given the LIME output.
```

Listing 1: Prompt template for DaLimeLlama explanations, including the system instructions and the actual request. Variables are written in curly brackets "{...}". Formatting instructions were shortened, and the complete prompt template is available in our GitHub Repository: https://github.com/AIG-ist-tugraz/DaLimeLlama.

better depending on the application and domain [30, 10], future research should explore instructing the LLM to incorporate specific principles and evaluate their impact on user response.

To generate understandable explanations with the LLM, features must be clearly labeled for accurate interpretation. For more complex domains or features, the LLM prompt could be extended with a "dictionary" that explicitly includes the interpretation and background of these features. This approach allows DaLimeLlama explanations to highlight user preferences considered in recommendations, enhancing system transparency by clarifying why a product was recommended and persuading users to purchase it.

In summary, the presented examples emphasize that DaLimeLlama improves the understandability and readability of LIME explanations by incorporating additional information from encoded LLM knowledge. While custom templates could be used to format LIME explanations, using an LLM simplifies this process and reduces manual effort, especially when handling numerous features or different contexts. Although we demonstrated this approach in a single domain using Llama 3, it can be adapted to other scenarios and LLMs by modifying the prompt accordingly. Overall, this approach can improve the explainability of LTR product recommenders.

4. Conclusion

Providing persuasive explanations in recommender systems is challenging. In this paper, we proposed a novel method combining the capabilities of XAI and LLM methods to improve the explainability of LTR product recommenders. We integrated the XAI method *LIME* with the *Llama 3* LLM to demonstrate our approach using a real-world e-commerce dataset. In our experiments, we observed that these *DaLimeLlama* explanations have a great potential to increase user acceptance compared to using LIME explanations directly. To investigate the usefulness of this approach in a real-life setting, we plan to evaluate the framework by applying *A/B testing* on various scenarios, including e-commerce and healthcare applications. The initial results presented in this paper encourage further research in this direction. An extension of the presented framework, including additional XAI methods, strategies to cope with LLM-related hallucination, and support of recommendation approaches beyond LTR, is left for future work.

References

[1] A. Karatzoglou, L. Baltrunas, Y. Shi, Learning to rank for recommender systems, in: Proceedings of the 7th ACM Conference on Recommender Systems, 2013, pp. 493–494.

[2] T. Chen, Xgboost: extreme gradient boosting, R package version 0.4-2 1 (2015).

[3] J. D. Novaković, A. Veljović, S. S. Ilić, Ž. Papić, M. Tomović, Evaluation of classification models in machine learning, Theory and Applications of Mathematics & Computer Science 7 (2017) 39.

[4] A. Parmar, R. Katariya, V. Patel, A review on random forest: An ensemble classifier, in: International conference on intelligent data communication technologies and internet of things (ICICI) 2018, Springer, 2019, pp. 758–763.

[5] Z. E. Aydin, Z. K. Ozturk, Performance analysis of xgboost classifier with missing data, Manchester Journal of Artificial Intelligence and Applied Sciences (MJAIAS) 2 (2021) 2021.

<table>
<tr><th>Recommendation</th><th>LIME Explanation</th><th>DaLimeLlama Explanation</th></tr>
<tr><td>
Shea Butter Lavender Liquid Soap
Hands / Hand Soap
Relax with the iconic soothing scent of lavender from Provence with this liquid hand and body soap. Formulated without palm oil, our liquid soaps are made from vegetable oil according to the traditional cooking recipe in a cauldron of the master soap makers of Provence. Enriched with shea butter known for its nourishing properties, this liquid soap cleanses the skin leaving it soft and delicately perfumed with the iconic Provencal scent of Lavender.
Brand: L'Occitane en Provence
Collection: Karité
Size: 500 ml

24€
</td><td>
<table>
<tr><th>Feature</th><th>Value</th></tr>
<tr><td>total purchases of this product</td><td>13.00</td></tr>
<tr><td>preferred product category=Eau de toilette</td><td>True</td></tr>
<tr><td>preference of best seller products</td><td>0.00</td></tr>
<tr><td>average price per product in purchase</td><td>27.50</td></tr>
<tr><td>product subcategory=not available</td><td>True</td></tr>
<tr><td>product category=Hand Soap</td><td>True</td></tr>
<tr><td>product is a best seller</td><td>0.00</td></tr>
<tr><td>preferred product collection=Les Classiques</td><td>True</td></tr>
<tr><td>product price</td><td>24.00</td></tr>
</table>
</td><td>
We think you'll love the Shea Butter Lavender Liquid Soap! Our recommendation is based on the fact that many customers have purchased this product before, with a total of 13 purchases. We also know that you tend to prefer products within a certain price range, and this soap falls comfortably within that range, with an average price per product in your purchases being €27.50. Additionally, you don't seem to prioritize best-seller products, which is great because this soap is a hidden gem! What's more, we've avoided recommending products from the Eau de toilette category, which isn't your preferred choice. Lastly, we've made sure to suggest a product with a clear subcategory which is important to you.
</td></tr>
<tr><td>
Ultra Rich Body Cream
Body / Moisturizing Body Cream & Milk
With a high concentration of shea butter (25%), this cream is a universal care for (very) dry and sensitive skin. It helps nourish the skin and soothes it, while helping to reinforce its protective barrier. Its rich texture is easy to apply and absorbs quickly without leaving a greasy film, and from the first use, the skin is moisturized for 72 hours, providing softness and comfort. Specially developed formula for sensitive skin
Brand: L'Occitane en Provence
Collection: Karité
Size: 200 ml
Bestseller

39€
</td><td>
<table>
<tr><th>Feature</th><th>Value</th></tr>
<tr><td>total purchases of this product</td><td>1.00</td></tr>
<tr><td>preferred product category=Eau de toilette</td><td>True</td></tr>
<tr><td>average price per product in purchase</td><td>27.50</td></tr>
<tr><td>preference of best seller products</td><td>0.00</td></tr>
<tr><td>product is a best seller</td><td>1.00</td></tr>
<tr><td>preferred product collection=Les Classiques</td><td>True</td></tr>
<tr><td>product subcategory=not available</td><td>True</td></tr>
<tr><td>product category=Moisturizing Body Cream & Milk</td><td>True</td></tr>
</table>
</td><td>
We think you'll love the Ultra Rich Body Cream! Since you've purchased a similar product before, we know you appreciate high-quality body care. Although you haven't specifically looked for body creams in the past, we believe this product will still meet your expectations. You tend to opt for products with a slightly higher price point, and at €39, this cream fits your budget. What's more, you don't seem to prioritize best-seller products, which is great because this one happens to be a popular choice. Overall, we're confident that the Ultra Rich Body Cream will become a new favorite of yours!
</td></tr>
<tr><td>
Almond Shower Oil
Body / Shower gel
This almond shower oil transforms, on contact with water, into a delicate milky veil that gently cleanses the body. Enriched with sweet almond oil, it respects the skin's hydrolipidic film to preserve its softness and comfort, keeping it hydrated throughout the day. Easy to rinse, this body cleanser leaves a subtle and delicious scent of fresh almonds for irresistible skin! L'Occitane almond shower oil is an ideal alternative to the usual shower gels and is suitable for the most delicate skins.
Brand: L'Occitane en Provence
Collection: Karité
Size: 200 ml
Bestseller

20€
</td><td>
<table>
<tr><th>Feature</th><th>Value</th></tr>
<tr><td>total purchases of this product</td><td>3.00</td></tr>
<tr><td>preferred product category=Eau de toilette</td><td>True</td></tr>
<tr><td>average price per product in purchase</td><td>27.50</td></tr>
<tr><td>preference of best seller products</td><td>0.00</td></tr>
<tr><td>primary product category=Body</td><td>True</td></tr>
<tr><td>product subcategory=not available</td><td>True</td></tr>
<tr><td>preferred primary product category=Perfume</td><td>True</td></tr>
<tr><td>product is a best seller</td><td>1.00</td></tr>
<tr><td>preference of sale products</td><td>0.00</td></tr>
</table>
</td><td>
We think you'll love the Almond Shower Oil! Since you haven't purchased this product before, you're in for a treat. We know you tend to prefer best-selling products, and this one fits the bill. Although you've shown interest in the Almendra collection, we think you'll appreciate the Body category of this product even more. Plus, with an average price per product in your purchases being around €24, this shower oil's price point will feel comfortable for you. Give it a try and discover a new favorite!
</td></tr>
</table>

Table 2

The table includes example product recommendations from the dataset, their LIME explanation, and the generated DaLimeLlama explanation. The color mapping in the LIME explanation indicates the direction of influence the respective feature has on the predicted score. Orange denotes a positive impact, while blue highlights a negative impact.

[6] S. M. Lundberg, S.-I. Lee, A unified approach to interpreting model predictions, in: Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17, Curran Associates Inc., Red Hook, NY, USA, 2017, p. 4768–4777.

[7] K. Weber, L. Tinnes, T. Huber, A. Heimerl, M.-L. Reinecker, E. Pohlen, E. André, Towards demystifying subliminal persuasiveness: using xai-techniques to highlight persuasive markers of public speeches, in: Explainable, Transparent Autonomous Agents and Multi-Agent Systems: Second International Workshop, EXTRAAMAS 2020, Auckland, New Zealand, May 9–13, 2020, Revised Selected Papers 2, Springer, 2020, pp. 113–128.

[8] M. Dragoni, I. Donadello, C. Eccher, Explainable ai meets persuasiveness: Translating reasoning results into behavioral change advice, Artificial Intelligence in Medicine 105 (2020) 101840.

[9] I. Donadello, M. Dragoni, C. Eccher, Explaining reasoning algorithms with persuasiveness: a case study for a behavioural change system, in: Proceedings of the 35th Annual ACM Symposium on Applied Computing, 2020, pp. 646–653.

[10] S. P. Erdeniz, T. N. Trang Tran, A. Felfernig, S. Lu-

bos, M. Schrempf, D. Kramer, P. P. Rainer, Employing nudge theory and persuasive principles with explainable ai in clinical decision support, in: 2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), 2023, pp. 2983–2989. doi:10.1109/BIBM58861.2023.10385315.

[11] A. Das, P. Rad, Opportunities and challenges in explainable artificial intelligence (xai): A survey, arXiv preprint arXiv:2006.11371 (2020).

[12] R. Carli, A. Najjar, D. Calvaresi, Risk and exposure of xai in persuasion and argumentation: The case of manipulation, in: International Workshop on Explainable, Transparent Autonomous Agents and Multi-Agent Systems, Springer, 2022, pp. 204–220.

[13] P. Mount, W. Yuen, Can upward brand extensions be an opportunity for marketing managers during the covid-19 pandemic and beyond? (2021).

[14] K. Tejasvi, V. L. Vinya, J. N. Padmaja, R. Begum, M. Jabbar, Explainable artificial intelligence (xai) for managing customer needs in e-commerce: A systematic review, Role of Explainable Artificial Intelligence in E-Commerce (2024) 17–31.

[15] S. Lubos, T. N. T. Tran, S. P. Erdeniz, M. E. Mansi, A. Felfernig, M. Wundara, G. Leitner, Concentrating on the impact: Consequence-based explanations in recommender systems, arXiv preprint arXiv:2308.16708 (2023).

[16] W. Saeed, C. Omlin, Explainable ai (xai): A systematic meta-survey of current challenges and future opportunities, Knowledge-Based Systems 263 (2023) 110273.

[17] B. Min, H. Ross, E. Sulem, A. P. B. Veyseh, T. H. Nguyen, O. Sainz, E. Agirre, I. Heintz, D. Roth, Recent advances in natural language processing via large pre-trained language models: A survey, ACM Comput. Surv. 56 (2023). URL: https://doi.org/10.1145/3605943. doi:10.1145/3605943.

[18] OpenAI, Gpt-4 technical report, 2023. arXiv:2303.08774.

[19] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, further authors, Llama 2: Open foundation and fine-tuned chat models, 2023. arXiv:2307.09288.

[20] AI@Meta, Llama 3 model card (2024). URL: https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md.

[21] L. Reynolds, K. McDonell, Prompt programming for large language models: Beyond the few-shot paradigm, in: Extended Abstracts of the 2021 CHI Conference on Human Factors in Computing Systems, CHI EA '21, Association for Computing Machinery, New York, NY, USA, 2021. URL: https://doi.org/10.1145/3411763.3451760. doi:10.1145/3411763.3451760.

[22] L. Beurer-Kellner, M. Fischer, M. Vechev, Prompting is programming: A query language for large language models, Proc. ACM Program. Lang. 7 (2023). URL: https://doi.org/10.1145/3591300. doi:10.1145/3591300.

[23] J. Lin, X. Dai, Y. Xi, W. Liu, B. Chen, X. Li, C. Zhu, H. Guo, Y. Yu, R. Tang, W. Zhang, How can recommender systems benefit from large language models: A survey, 2023. URL: https://doi.org/10.48550/arXiv.2306.05817. arXiv:2306.05817.

[24] L. Wu, Z. Zheng, Z. Qiu, H. Wang, H. Gu, T. Shen, C. Qin, C. Zhu, H. Zhu, Q. Liu, H. Xiong, E. Chen, A survey on large language models for recommendation, 2023. URL: https://doi.org/10.48550/arXiv.2305.19860. arXiv:2305.19860.

[25] I. Silva, L. Marinho, A. Said, M. C. Willemsen, Leveraging chatgpt for automated human-centered explanations in recommender systems, in: Proceedings of the 29th International Conference on Intelligent User Interfaces, IUI '24, Association for Computing Machinery, New York, NY, USA, 2024, p. 597–608. URL: https://doi.org/10.1145/3640543.3645171. doi:10.1145/3640543.3645171.

[26] S. Lubos, T. N. T. Tran, A. Felfernig, S. Polat Erdeniz, V.-M. Le, Llm-generated explanations for recommender systems, in: Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization, UMAP Adjunct '24, Association for Computing Machinery, New York, NY, USA, 2024, p. 276–285. URL: https://doi.org/10.1145/3631700.3665185. doi:10.1145/3631700.3665185.

[27] M. T. Ribeiro, S. Singh, C. Guestrin, " 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, 2016, pp. 1135–1144.

[28] C. J. Burges, From RankNet to LambdaRank to LambdaMART: An Overview, Technical Report MSR-TR-2010-82, 2010. URL: https://www.microsoft.com/en-us/research/publication/from-ranknet-to-lambdarank-to-lambdamart-an-overview/.

[29] R. Cialdini, Influence: The Psychology of Persuasion, Harper Business, 2021.

[30] S. Gkika, G. Lekakos, The persuasive role of explanations in recommender systems., in: BCSS@PERSUASIVE, 2014, pp. 59–68.