

Explaining Social Recommendations Using Large Language Models

Md Ashaduzzaman¹, Thi Nguyen¹, and Chun-Hua Tsai¹

College of Information Science & Technology, University of Nebraska Omaha,
Omaha NE 68182, USA

Abstract. This paper introduces an innovative approach to recommender systems through the development of an explainable architecture that leverages large language models (LLMs) and prompt engineering to provide natural language explanations. Traditional recommender systems often fall short in offering personalized, transparent explanations, particularly for users with varying levels of digital literacy. Focusing on the Advisor Recommender System, our proposed system integrates the conversational capabilities of modern AI to deliver clear, context-aware explanations for its recommendations. This research addresses key questions regarding the incorporation of LLMs into social recommender systems, the impact of natural language explanations on user perception, and the specific informational needs users prioritize in such interactions. A pilot study with 11 participants reveals insights into the system’s usability and the effectiveness of explanation clarity. Our study contributes to the broader human-AI interaction literature by outlining a novel system architecture, identifying user interaction patterns, and suggesting directions for future enhancements to improve decision-making processes in AI-driven recommendations.

Keywords: Recommender Systems, LLMs, Prompt Engineering, Natural Language Explanations, Human-AI Interaction

1 Introduction

Social recommendation systems facilitate decision-making processes by filtering out irrelevant information, thereby assisting users in online shopping or forming social connections on media platforms [17]. However, the opacity of artificial intelligence (AI) systems enables powerful predictions but compromises direct explainability, which can undermine user trust and experience [2]. In response to this challenge, recent research in recommender systems has focused on developing more sophisticated interfaces, including enhancing the explainability and controllability of recommendations [19, 13, 18]. Despite these advancements, existing designs and solutions rely on predefined rules that lack interactive elements with users and adapt to their desires, needs, and literacy levels. Consequently, there remains a pressing need for a dynamic, user-centered approach to delivering AI-based recommendations that can accommodate users’ diverse backgrounds and foster transparency and informed decision-making.

The rapid development of Large Language Models (LLMs), such as Bard, GPT-4, LLaMA, and others, has ushered in a new era characterized by both challenges and opportunities [10, 14]. These models have demonstrated exceptional proficiency in comprehending and generating human-like text. Prompt engineering provides reusable solutions to common challenges in output generation and interaction when working with LLMs. It enables the creation of more accurate, context-specific, and nuanced responses from LLMs, particularly supporting advancements in conversational AI [12]. To facilitate the delivery of recommendations accompanied by natural language explanations, we propose an architecture for an explainable recommender system that leverages the human-like conversational capabilities of generative AI. Our proposed architecture enables users to receive personalized suggestions and gain insight into the rationale underlying these recommendations.

Specifically, this paper aims to answer the research question of *how to utilize large language models (LLMs) to enhance explanatory capabilities within social recommendation systems*. To this end, we designed and developed a prototype of an *Advisor Recommender System*, which integrates two content-based recommendation models - text similarity and topical similarity - to generate personalized advisor recommendations for graduate students based on their research interests. Moreover, our system enables users to engage in follow-up inquiries regarding the given recommendations, i.e., allowing users to ask questions about the recommended advisors, and the system could explain why a particular scholar or faculty is recommended to them. These communications are based on LLM-powered responses providing detailed, context-sensitive explanations. We hypothesize that these interactive explanatory capabilities can foster enhanced user comprehension and trust in the recommender system.

To assess the efficacy of our proposed system and the user experience, we undertook pilot user studies involving 11 graduate students majoring in computer sciences or information systems. Through a qualitative thematic analysis of the data collected, three primary themes emerged: firstly, the types of information sought by participants when engaging with the system’s recommendations; secondly, the usability and feasibility of integrating a chatbot within the recommendation system; and thirdly, how participants interacted with the chatbot to obtain explanatory rationales for the recommendations. Through these user feedback, we identified potential avenues for system improvement to enhance user satisfaction and facilitate informed decision-making. Based on these experimental findings, we then discuss the design implications and insights of using LLMs along with explainable recommendation systems.

2 Related Works

The concept of collaborative filtering was first introduced by Goldberg et al. [7] in 1992 as a method to filter mail by tracking individuals’ responses to the documents they read. Resnik and Varian [15] first introduced the term ‘recommender system’ in 1997 to describe systems that not only filter out less relevant

items but also suggest particularly interesting ones. In October 2006, Netflix’s release of a movie rating dataset [3] and subsequent challenge spurred increased research interest in recommender systems. Advances in machine learning deliver AI systems that operate independently but cannot explain their decisions. To bridge this gap, DARPA initiated the explainable artificial intelligence (XAI) program in May 2017 [8]. The black-box nature of AI systems enables powerful predictions but lacks direct explainability. XAI offers significant potential to enhance trust and transparency in AI-based systems. Explanation in recommendations is crucial for justifying outcomes, increasing visibility to swiftly identify and correct errors, and continuously improving the system [2]. Parra and Brusilovsky [13] sought to enhance the controllability and transparency of hybrid recommendations by integrating sliders for fusion control and a Venn diagram to visualize the outcomes. Several research studies have focused on enhancing explainability by augmenting interactive hybrid recommender interfaces with various types of explanations [19, 18]. Tsai et al. [19] discussed four types of explanatory interfaces: a two-way bar chart for illustrating publication similarity; a topical radar for showing topic similarity; an enhanced strength network for visualizing co-authorship networks; a venn tag for depicting CN3 interest similarity. Several studies have utilized NLP models to generate textual explanations for specific recommendations [9, 6]. P5 [6] demonstrates that learning multiple recommendation-related tasks simultaneously is feasible through a unified sequence-to-sequence framework by treating these tasks as prompt-based natural language tasks, integrating user-item information and features with personalized prompt templates as inputs, highlighting a promising approach for unified, instruction-based recommendation. However, a significant limitation of these NLP-related models is their inability to generate human-like text. Recent LLMs like GPT-4 [1] excel at generating coherent and contextually relevant text, and are adept at holding conversations, answering questions, providing detailed explanations, and engaging in nuanced, context-aware dialogues, surpassing earlier versions in complexity and understanding. Our research is driven by the potential to leverage the advanced capabilities of recent LLMs to develop an explainable recommender system.

3 Explainable Recommender System Architecture

In this section, we describe the proposed architecture of an explainable recommender system, as illustrated in Figure 1. The architecture is based on an advisor recommender system that utilizes LLM for generating explanations. Initially, users must input research interest keywords into the system. These keywords are processed by both a Text Similarity and a Topic Similarity model. Each model identifies the top three advisors and displays them to the user. The results along with similarity scores, LDA topics, and other relevant model data are saved for the prompt engineering process. When a user requests an explanation of the recommendations, the chatbot interacts with the feature engineering process to create a suitable prompt. This prompt is then sent to the LLM, which provides




Fig. 1. Proposed architecture of Advisor Recommendation System leveraging LLM for explanation.

an explanation for the user query. The response is delivered to the user, and the conversation history is stored to enhance future interactions through refined prompt engineering. Details of each step are outlined below:

3.1 System Implementation and Data

An advisor recommender system designed for prospective graduate students seeking a research supervisor can significantly streamline the process of finding the right mentor for advanced studies. This system matches students with potential supervisors based on shared research interests, publications, and other relevant criteria. It generates a list of potential supervisors whose research aligns closely with the student's interests. It reduces the time and effort needed for students to find compatible supervisors. By aligning students with supervisors who share their research interests, the system promotes more effective and fulfilling research collaborations. The recommendations produced by our proposed system originate from information obtained from the Citation Network Dataset [16]. This citation information is pulled from several sources including DBLP, ACM, and the Microsoft Academic Graph (MAG), among others. The dataset includes information from the conference proceedings, detailing the conference papers (author(s), title, keywords, and abstract) and providing extensive profiles on the authors and attendees, including names, publications, affiliations, and positions. While the dataset includes a vast quantity of data (80435 authors), only a limited portion of it was selected randomly for our research. We extracted the names, IDs, and research keywords from the publications of 2,500 researchers who we are considering as potential advisors for the user. We selected a lim-

Advisor Recommender System

Enter your Name

Ashad

Enter keywords of your reseach interest

data mining

Submit

Top Text Similarity Recommendations:

Ranking	Name
1	Cristóbal Romero
2	Ramesh Natarajan
3	Yew-Kwong Woon

Top Topic Similarity Recommendations:

LDA_rank	LDA_Name
1	Zoran Obradovic
2	Wei-Hua Wang
3	Isidore Rigoutsos

Welcome Ashad! Would you like recommendation explanation?

Example: 1. Tell me the research interests of the recommended advisor based on cosine

Fig. 2. User interface of the Advisor Recommendation System.

ited number of researchers to ensure the system responds quickly and to avoid complexity, as our main goal is to develop a system that can clearly explain its recommendations.

We developed the Advisor Recommender system using the Streamlit framework. We hosted this app file, along with the relevant datasets and trained models, in a GitHub repository. After setting up the repository, we connected it to the Streamlit cloud, allowing us to deploy the app seamlessly. GitHub repository can be found at ¹. Figure 2 displays the user interface of the Advisor Recommender system. It prompts users to enter their name and research interest keywords. Users may type any research interests related to computer and information science. If users have multiple research interests, they can enter them in the text field, separated by commas or semicolons. When a user clicks the submit button after entering their name and research interest keywords, the interface interacts with text and topic similarity-based models to recommend the top three advisors for each model. Initially, it prompts the user to wait while the top-recommended advisors are being retrieved. The output is then presented in two columns of information. On the left side of the image, the top three advisors are listed based on text (Cosine) similarity, and on the right side, based on topic (LDA) similarity, displayed in a table format. The first column of each table shows the ranking based on similarity score, and the second column lists the name of the advisor. At the bottom of the interface, an AI chatbot is integrated, asking whether an explanation of the recommendations is needed, followed by a user input field for queries. This chatbot captures user input through an interac-

¹ https://github.com/asadayon/Advisor_Recommendation

tive interface, generates a relevant prompt using prompt engineering techniques, communicates with LLM via API, and delivers the response back to the user. Users can ask questions related to advisors, recommendations, and models.

3.2 Recommendation Models

Text Similarity Model Previous research [19, 18] employed text similarity models based on scholars’ publications to recommend attendees with similar interests. This model operates by initially generating count vectors from the keywords associated with all authors’ publications within our comprehensive publication database. These count vectors serve as numerical representations of the keywords, allowing us to quantify and analyze the textual data effectively. Subsequently, we employ cosine similarity to measure the degree of similarity between the keywords input by the user (representing their research interests) and the keywords associated with other authors’ publications. Cosine similarity provides a score ranging from 0 to 1, where a score closer to 1 signifies a higher degree of similarity, indicating closely aligned research interests between the user and the author in question. Based on the cosine similarity scores, the system then identifies and recommends the top three authors whose publication keywords most closely match the user’s specified interests. These authors are suggested as the top three recommended advisors, underpinning the system’s goal to facilitate meaningful academic collaborations and mentorships grounded in shared research interests.

Topic Similarity Model Given the abstract of the article, a list of authors, and known past collaborators, Topic Models, such as Latent Dirichlet Allocation (LDA), can generate a list of highly relevant authors who would be suitable reviewers for this article [5]. By categorizing keywords from a vast array of publications into 30 distinct topics, the system effectively maps out the research landscape into thematic clusters, each defined by a set of representative keywords. Authors in the database are analyzed to ascertain their thematic engagement across these topics, with each author assigned a probability score reflecting their involvement in each topic. This creates a thematic profile for each author. When a user inputs interest keywords, the system dynamically matches these against the predefined topics to identify the top three topics most aligned with the user’s interests, quantified by probability scores. Using the gensim library, the system compares the user’s topic preferences with the thematic profiles of authors to recommend the top three authors whose research interests best match the user’s. This approach ensures that recommendations are grounded in thematic, rather than merely textual, similarities, fostering more relevant and intellectually aligned collaborations.

3.3 Generating Explanations

LLM models are trained on vast amounts of text data, allowing them to learn language patterns, context, and semantics at a sophisticated level. These mod-

els excel in various natural language processing tasks, including text generation, summarization, translation, and question-answering. Our system’s chatbot, powered by the LLM model (specifically, GPT-3.5 Turbo), provides explanations for the system’s recommendations in a conversational and user-friendly manner. When supplied with context such as recommended outputs and pertinent background knowledge as input prompts, the model crafts explanations that are both informative and easily understandable. This approach enhances user engagement by encouraging deeper exploration of the recommendations and offering clear insights into the underlying rationale.

Prompt engineering involves carefully crafting and refining the input prompts to steer the responses of these models, ensuring that the output is accurate, relevant, and coherent [11]. This technique is essential for fully utilizing the capabilities of LLM models, making them more versatile and effective across various fields. The prompt must be clear and precise to guide the AI in generating relevant and accurate responses. Ambiguities or vague descriptions can lead to unpredictable or irrelevant outputs. Including the right context in the prompt can help the AI understand the scope and depth of the response needed. The prompt should be aligned with the end goal while it’s generating text or providing explanations.

In the proposed prompt engineering architecture, we specify that the system is designed to operate as an advisory recommendation system, elucidating the rationale behind recommending specific advisors. We then gather pertinent data such as the top recommended advisors, their similarity scores, and the LDA topics derived from both the Text and Topic similarity models. Additionally, the system retrieves research interest keywords of the top recommended advisors from its storage. To enhance user interaction, we incorporate contextual details such as the user’s name, research interests, recommendations from both models, advisors’ research interests, similarity scores, and topics into prompts. We also stipulate that the system should refrain from providing responses beyond its designated scope to mitigate hallucination. This augmented prompt framework is employed to facilitate communication between the LLM and the chatbot, ensuring effective handling of user queries.

4 Experiment

We conducted pilot user studies of the Advisor recommendation system involving 11 participants (P1 through P11), all of whom are graduate students with multiple interests in computer science related fields, except for one undergraduate. These participants had prior experience working with research advisors. All participants resided in Omaha, United States, and were invited via email to voluntarily participate in the pilot studies, which involved using an advisor recommendation system followed by a post-study questionnaire. Upon agreeing to participate, each was sent a Zoom invitation. The user studies varied in length from 13 to 20 minutes. Each session began with a brief presentation by the first author about the advisor recommendation system, followed by provid-

ing the participants with a link to access the system. After using the system, participants engaged in a semi-structured interview as part of the post-study questionnaire. All sessions were recorded on the cloud via Zoom, and the interviews were conducted in English. Consent for recording was obtained from each participant.

Interview recordings were transcribed using Zoom. We employed an inductive approach to generate coding categories [4]. Two researchers independently reviewed the interview transcripts and manually identified initial codes. These initial codes were then discussed collectively by the research team, leading to the formulation of a more comprehensive code list. The finalized codes were organized into three main themes. These themes include the types of information users wish to follow up on based on system recommendations, the usability of integrating an AI-powered chatbot, and user responses to the chatbot’s reasoning capabilities and its ability to impart new knowledge. Quotations have been lightly edited for clarity.

5 Findings

The Findings section is structured around three central themes: firstly, the types of information students seek when receiving recommendations from the system. Secondly, the usability of integrating a chatbot within the recommendation system. Lastly, the nature of user interactions with the chatbot when seeking explanations for the recommendations.

5.1 Theme 1: Information Needs of Advisor Seeking

There are different types of information that the users want to know after getting the recommended advisor lists from the system. The most in-demand type of information is closely related to the top recommended advisors, including their academic institutions or affiliations, the publications that they had on the requested topics, and any personal website if available. For instance, user P5 posed questions after getting the list of advisors from the chatbot: *“Give a bio about [X]”* and *“Where is this person [the advisor] affiliated?”*. User P3 inquired: *“What are the papers he published?”*. Meanwhile, user P6 requested: *“Can you share his profile with me?”*. However, our current system design only incorporates the advisors’ publication research keywords as a means of prompt engineering. Consequently, the chatbot is unable to respond to queries about other aspects of the advisors’ profiles. The chatbot responded with: *“I’m not sure about [X]’s current affiliation...”*, *“I apologize, but I do not have access to a specific list of publications for individual advisors...”* and *“I’m unable to provide real-time or specific information about...”*. The majority of users (9 out of 11) believe that information on advisor affiliation is essential for the recommendation system, as it facilitates the subsequent task of contacting them. User P1 commented during the post-study questionnaires, *“Unless and until it shows the affiliation of that particular advisor to some university or similar institution, it is like... how can*

we use this recommendation?”. According to this user, the recommendation is not useful unless the affiliation is provided. User P2 commented, *“If I found a source [weblink] about [advisor], it would be much better... I wish it could be implemented in the system.”* According to this user, a personal website or link to the recommended advisor is very useful. When User P3 was asked about their experience while using the system, he commented, *“It is quite good, but one recommendation I can give is to add the description of the particular advisor, and what are the papers and the publications...”* The user suggested that the system would be better if it also provided descriptions of advisors and their publication lists.

5.2 Theme 2: User Perceptions of Usability

The term usability in this subsection refers to the system’s ability to provide users with practical information that helps them effectively find advisors aligned with their research interests. Among 11 participants, there are 7 users who think that the chatbot is able to provide them with helpful information when finding advisors for their research. *“It is helpful to narrow down the supervisor names based on my research interest.”*, user P7 added when being asked if they perceived any benefit from using the system. P8 reported that after requesting an explanation of the model used, he received a convincing answer. He commented, *“I feel like I’m having a human-like chat with the system,”* indicating that he believed he was conversing with a real person. However, the remaining 4 users suggested that it would be more beneficial if more detailed information about individual advisors were provided. P2 expressed confusion while using the system, stating, *“I was a bit confused about how much it knows; I didn’t realize the scope was really limited. I expected the advisor system to provide more detail.”* The user, P2, suggests that the system should provide more detailed information, as the current scope is too limited. When asked about interface improvements to enhance usefulness and interactivity, most users were satisfied with the current interface. However, we received several specific suggestions. These included adding tooltips with quick instructions to guide users on how to begin and the functionalities available. Additionally, it was suggested that user inputs should remain visible at the top of the screen while allowing the conversation section below to be scrollable. This would prevent users from having to scroll all the way back up to adjust research keywords as needed. Again, incorporating thumbs-up and thumbs-down icons would enable users to rate the responses they receive. Another point of praise was the system’s quick response time, as noted by three users, indicating that users appreciate prompt replies from the chatbot and dislike waiting for responses.

5.3 Theme 3: Transparency and Explainability

For each user study, the users are encouraged to ask the chatbot any questions that come to their minds related to the recommended list of advisors. There are 7 out of 11 users who explicitly asked the chatbot to give an explanation on its

recommendations or previous answers, such as *"Why was [X] recommended as number 1?"*, *"Why [X] is recommended in both similarities"*, *"...explain (to) me in simple terms"*, while the rest of the users asked follow-up questions in relevance to the previous conversation they had with the chatbot. Users found the concept of explaining recommendations both interesting and intuitive. P9 commented, *"Yeah, it was very interesting... when I purchase anything on Amazon, it suggests similar products. However, it never explains the basis on which those products are recommended. When I searched for my topic, like image processing, this system not only gave recommendations but also explained why these advisors were suggested. So, that was really a good idea."* Thus, the user found it very interesting to receive reasons for the recommendations. User P10 commented, *"In this scenario, where any recommendation given by [the system] can be verified by, you know, the algorithm inside it, I think it's very unique and helpful."* This implies that they appreciate transparency in how the recommendations are generated, as it allows for greater trust and reliability in the system's suggestions. P11 commented, *"I really like the way it was trying to explain, even to any user, even unaware of the technical aspect, but still can get a glimpse of what the model is trying to recommend. It gives you more confirmation, more clarity on why something is being recommended or prescribed to you, instead of being blindly recommended to you,"* The main point of the comment is that user appreciates the system's ability to explain its recommendations in a way that is accessible to all users, regardless of their technical knowledge. In fact, knowing that there is a logical explanation behind the recommendations gives users more confidence and trust to adopt them into their decision-making process. Some users noted that the answers were clear, short, and concise, making them easy to follow. One user specifically requested explanations in simple terms, which the system provided. This suggests that users generally prefer explanations that are both concise and straightforward.

6 Discussion

In addressing the research question on how to utilize LLMs to enhance explanatory capabilities within social recommendation systems, our study demonstrated the effectiveness of leveraging LLMs in providing personalized, transparent explanations. Through prompt engineering, our Advisor Recommender System provided personalized explanations, citing shared research interests and high cosine similarity scores when users inquired about specific advisor recommendations. While the system successfully addressed basic queries, limitations were noted when detailed affiliation information was lacking. These findings highlight the potential of LLMs in contextualizing recommendations and improving user understanding, while also indicating areas for further refinement in prompt design.

Usability emerged as a critical factor in the adoption and effectiveness of the system. Participants appreciated the system's ability to narrow down potential advisors based on their research interests. However, there were notable concerns

about the depth of information provided. The mixed feedback on the system’s knowledge scope highlights the importance of setting clear expectations for users about the system’s capabilities and limitations. Enhancing the utility and satisfaction of recommendation systems requires comprehensive data integration, which can be accomplished by implementing prompt engineering to provide detailed information about recommended advisors. Enhancements such as tooltips, persistent visibility of user inputs, and interactive elements like rating icons could further improve the user experience and engagement.

A significant number of users have explicitly requested explanations for the recommendations or previous answers provided by the chatbot, highlighting a strong desire for transparency and understanding of the system’s suggestions. Comments from participants P9, P10, and P11 underscore the importance of explanation and transparency in recommendation systems, with examples illustrating common frustrations when recommendations lack accompanying explanations and the value users place on understanding algorithms for credibility. However, there’s a concern about whether users genuinely comprehend the explanations provided, suggesting a potential gap between expressed interest and actual understanding. Further research is needed to investigate users’ comprehension levels and develop strategies to support them in formulating informed inquiries effectively, thus enhancing their learning from AI recommendations.

7 Conclusion

Our research on the Advisor Recommender System demonstrates the value of integrating LLMs to provide natural language explanations within recommender systems. The findings from our pilot study indicate that users value transparency and personalized explanations, which enhance their understanding and trust in the recommendations. While the system shows promising results in user interaction and satisfaction, there are opportunities for further refinement, particularly in enriching the data sources. Future work will focus on these improvements to better meet user needs and further advance the field of human-AI interaction. We also plan to assess the extent to which the explanation of recommendations enhances user knowledge and how individuals with different levels of digital literacy understand these explanations.

Acknowledgements This material is based upon work supported by the National Science Foundation under Grant No. 2153509.

References

- [1] Josh Achiam et al. “Gpt-4 technical report”. In: *arXiv preprint arXiv:2303.08774* (2023).
- [2] Amina Adadi and Mohammed Berrada. “Peeking inside the black-box: a survey on explainable artificial intelligence (XAI)”. In: *IEEE access* 6 (2018), pp. 52138–52160.

- [3] James Bennett, Stan Lanning, et al. “The netflix prize”. In: *Proceedings of KDD cup and workshop*. Vol. 2007. New York. 2007, p. 35.
- [4] Virginia Braun and Victoria Clarke. “Using thematic analysis in psychology”. In: *Qualitative research in psychology* 3.2 (2006), pp. 77–101.
- [5] Uttam Chauhan and Apurva Shah. “Topic modeling using latent Dirichlet allocation: A survey”. In: *ACM Computing Surveys (CSUR)* 54.7 (2021), pp. 1–35.
- [6] Shijie Geng et al. “Recommendation as language processing (rlp): A unified pretrain, personalized prompt & predict paradigm (p5)”. In: *Proceedings of the 16th ACM Conference on Recommender Systems*. 2022, pp. 299–315.
- [7] David Goldberg et al. “Using collaborative filtering to weave an information tapestry”. In: *Communications of the ACM* 35.12 (1992), pp. 61–70.
- [8] David Gunning and David Aha. “DARPA’s explainable artificial intelligence (XAI) program”. In: *AI magazine* 40.2 (2019), pp. 44–58.
- [9] Lei Li, Yongfeng Zhang, and Li Chen. “Generate neural template explanations for recommendation”. In: *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2020, pp. 755–764.
- [10] Yiheng Liu et al. “Summary of chatgpt-related research and perspective towards the future of large language models”. In: *Meta-Radiology* (2023), p. 100017.
- [11] Yao Lu et al. “Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity”. In: *arXiv preprint arXiv:2104.08786* (2021).
- [12] Ggaliwango Marvin et al. “Prompt Engineering in Large Language Models”. In: *International Conference on Data Intelligence and Cognitive Informatics*. Springer. 2023, pp. 387–402.
- [13] Denis Parra and Peter Brusilovsky. “User-controllable personalization: A case study with SetFusion”. In: *International Journal of Human-Computer Studies* 78 (2015), pp. 43–67.
- [14] Alec Radford et al. “Better language models and their implications”. In: *OpenAI blog* 1.2 (2019).
- [15] Paul Resnick and Hal R Varian. “Recommender systems”. In: *Communications of the ACM* 40.3 (1997), pp. 56–58.
- [16] Jie Tang et al. “ArnetMiner: Extraction and Mining of Academic Social Networks”. In: *KDD’08*. 2008, pp. 990–998.
- [17] Jiliang Tang, Xia Hu, and Huan Liu. “Social recommendation: a review”. In: *Social Network Analysis and Mining* 3 (2013), pp. 1113–1133.
- [18] Chun-Hua Tsai and Peter Brusilovsky. “Explaining recommendations in an interactive hybrid social recommender”. In: *Proceedings of the 24th international conference on intelligent user interfaces*. 2019, pp. 391–396.
- [19] Chun-Hua Tsai and Peter Brusilovsky. “The effects of controllability and explainability in a social recommender system”. In: *User Modeling and User-Adapted Interaction* 31 (2021), pp. 591–627.