Exploring the Power of Weak Ties on Serendipity in Recommender Systems

Wissam Al Jurdi[1,2], Jacques Bou Abdo[3], Jacques Demerjian[2], and Abdallah Makhoul[1]

[1] University Bourgogne Franche-Comté, FEMTO-ST Institute, CNRS, Montbéliard, France
wissam.al_jurdi@univ-fcomte.fr, abdallah.makhoul@univ-fcomte.fr
[2] Lebanese University, LaRRIS, Faculty of Sciences, Beirut, Lebanon
jacques.demerjian@ul.edu.lb
[3] University of Cincinnati
Cincinnati, United States
bouabdjs@ucmail.uc.edu

Abstract. With our increasingly refined online browsing habits, the demand for high-grade recommendation systems has never been greater. Improvements constantly target general performance, evaluation, security, and explainability, but optimizing for serendipitous experiences is imperative since a serendipity-optimized recommender helps users discover unforeseen relevant content. Given that serendipity is a form of genuine unexpected experiences and recommenders are facilitators of user experiences, we aim at leveraging weak ties to explore their impact on serendipity. Weak ties refer to social connections between individuals or groups that are not closely related or connected but can still provide valuable information and opportunities. On the other hand, the underlying social structure of recommender datasets can be misleading, rendering traditional network-based approaches ineffective. For that, we developed a network-inspired clustering mechanism to overcome this obstacle. This method elevates the system's performance by optimizing models for genuine unexpected content. By leveraging group weak ties, we aim to provide a novel perspective on the subject and suggest avenues for future research. Our study can also have practical implications for designing online platforms that enhance user experience by promoting unexpected discoveries.

Keywords: recommender systems, clustering, serendipity, weak ties, social science

1 Introduction

Recommendation systems utilize complex algorithms to analyze large datasets and generate personalized recommendations for users. They have proven to be highly effective in various industries, including e-commerce and media [1], and have been shown to improve user engagement and satisfaction significantly [2]. Recommenders are typically designed utilizing data features such as users' search and interaction details. Designing an effective recommendation system that aligns with the main goal of recommenders remains a subjective and very challenging [3] task despite the abundance of research, convenient models, and diversified workflows, such as the recent release of Microsoft's best practices for building recommendation systems [4]. The main goal of recommender systems is to provide users with information that is tailored to their interests [5,6]. Despite the major agreement on this general definition, research pathways frequently prioritize other aspects that sometimes even impede the attainment of this goal, as the study by Herlocker et al. highlights regarding the accuracy improvement branches that stemmed from a concept in evaluation termed the "magic barrier", i.e., the point beyond which recommenders fail to become more accurate [7,8].

Nonetheless, there has been recent research that focuses on further refining recommender systems by addressing crucial issues in general performance [9,10], noise and evaluation [11,12,13], security [14], and explainability [15]. Despite progress, there is still a notable disparity: designing methods or frameworks that enhance user engagement and knowledge through chance discoveries and exploration of the unknown.

The concept of "strength of weak ties" is an influential social science theory that emphasizes the role of weak associations, such as acquaintances, in spreading information and creating opportunities through social networks. Weak ties are more likely to provide novel information than strong ties, such as close friends, who tend to share similar perspectives and resources [16]. As we have highlighted previously, recommender algorithms aim to suggest information likely to interest a user. Therefore, incorporating weak ties into recommender systems can enhance their effectiveness by presenting users with unforeseen recommendations that they may not have otherwise discovered, as one study by Duricic et al. recently hinted at [17]. This performance is not to be confused with general performance in terms of precision or accuracy; it is a little more complex to measure and set up in this case and requires methods beyond the conventional ones [7,11,13]. Surely, weak ties can also help overcome the cold-start problem, where new users have insufficient data to generate personalized recommendations. However, weak ties also pose privacy challenges, as users may not want to reveal their preferences or behavior to distant or unknown connections [18].

Our approach uses social network theory to measure the impact of weak connections between user groups on serendipity, making it a vital part of the recommender ecosystem. Narrowing the focus only on chance discoveries allows us to advance recommender enhancement for their primary objective. There are two types of recommender data sources: social and rating [19]. Social datasets have data about user relationships or interactions, such as friendships, likes, etc. Rating datasets have data about user ratings for items or services, such as stars, likes, etc. In our study, we target the rating-based datasets and recommenders. This research also builds on previous validation and serendipity approaches [11,13,20]. It adapts a new optimization framework to train recommenders oriented towards chance discoveries for users.

In the following section, we explore the latest research while also placing our own work in this context. Section 3 introduces our unique approach involving community-based data processing and cluster assessment. We present the experimental results and analysis in Section 4 and conclude with Section 5.

2 Background and Related Work

This section provides an overview of the research conducted on serendipity in recommender systems. Currently, there is no established method to specifically improve the chances of discovering new and unexpected recommendations and to increase user involvement in recommenders. This is especially true when it comes to making use of weak connections between clusters.

2.1 Serendipity in Recommenders

Some recent recommender system proposals aim to improve serendipity. For example, Kotkov et al. [21] proposed a new definition of serendipity in recommenders that considers items that are surprising, valuable, and explainable, arguing that the common understanding and original meaning of serendipity is conceptually broader, requiring serendipitous encounters to be neither novel nor unexpected. Others have proposed a multi-view graph contrastive learning framework that can enhance cross-domain sequential recommendation by exploiting serendipitous connections between different domains [22].

The study by Ziarani et al. [23] is crucial and reviews the overall serendipity-oriented approaches in recommender systems. The authors emphasize the significance of serendipity in generating attractive and practical recommendations in recommender systems. This reinforces our introduced concept regarding the direction and primary objective of recommenders in this work's introduction. The approaches covered in the study generally discuss serendipity enhancements by introducing randomness into the recommendation process, which can lead to the discovery of new and interesting items that the user may not have discovered otherwise. In addition, serendipity can be enhanced by incorporating diversity into the recommendation process, which can help reduce over-specialization and make recommendations more interesting and engaging. The study concludes that while there is no agreement on the definition of serendipity, most studies find serendipitous recommendations to be valuable and unexpected.

In a study about surprise in recommenders by Eugene Yan [24], the importance of a serendipity metric in recommenders is discussed. The author argues that while accuracy is an essential metric for recommendation systems, it is not the only metric that matters. Recommender systems that solely focus on accuracy can lead to information over-specialization, making recommendations boring and predictable. To address this issue, the author suggests incorporating serendipity as a criterion for making appealing and useful recommendations. Serendipity is defined as a criterion for making unexpected and relevant recommendations to the user's interests [23]. The usefulness of serendipitous recommendations is the main superiority of this criterion over novelty and diversity. The article highlights that serendipity can be measured using various metrics such as surprise, unexpectedness, and relevance [24]. The article further explains that serendipity-oriented recommender systems have been the focus of many studies in recent years. The author conducted a systematic literature review of previous studies on serendipity-oriented recommender systems. The review focused on the contextual convergence of serendipity definitions, datasets, serendipitous recommendation methods, and their evaluation techniques [23]. The results of the review indicate that the quality and quantity of articles in the serendipity-oriented recommender systems are progressing. In conclusion, incorporating serendipity as a criterion for making recommendations can help make them more appealing and useful. It can also help address issues related to information over-specialization and make recommendations more diverse.

One of the studies by Bhandari et al. [25] proposes a method for recommending serendipitous apps using graph-based techniques. The approach can recommend apps even if users do not specify their preferences and can discover apps that are highly diverse. The authors also introduce randomness into the recommendation process to increase the likelihood of discovering new and interesting items that the user may not have discovered otherwise. Therefore, similar to the studies covered in [23], this unique process of app recommendations also uses the same method of randomness.

2.2 Recommendations and social network connections

Another proposal by M. Jenders et al. [26] introduces a content-based recommendation technique with a focus on the serendipity of news recommendations. Serendipitous recommendations have the characteristic of being unexpected yet fortunate and interesting to the user and thus might yield higher user satisfaction. The authors explore the concept of serendipity in the area of news articles and propose a general framework that incorporates the benefits of serendipity and similarity-based recommendation techniques. An evaluation against other baseline recommendation models is carried out in a user study.

Based on the studies mentioned above, it is clear that enhancing serendipity is a crucial step in improving recommender systems. However, there is currently no established framework for achieving this goal aside from incorporating randomization into the system.

3 Community-based Mechanism

To establish weak-connection-based recommendations, there are multiple steps involved in utilizing recommender data. We can ideally establish two kinds of connections: social-based links [19] or non-social links inferred from user behavior. In our experimentation, we introduce the latter and develop an approach that could be expanded further if recommender datasets were to be enriched with more information, particularly those with social and rating-based components. The diagram shown in Figure 1 depicts the metadata level, where we aim to enhance the recommendations by processing data differently through the community-based mechanism. To achieve this, we use an approach inspired by networks theory involving grouping users and utilizing weak links between them and the communities (or groups) they belong to. We use techniques like Gower [27] to form initial user clusters and then create principal clusters to establish higher-level communities. Theoretically, this should help us optimize the recommendations to provide more relevant and unexpected suggestions. Next, we refine the training process for the recommender system; this includes modifying the cluster and principal group formation parameters and generating various versions of the potential "weak links" between groups, as depicted in Figure 1. The aim is to avoid any prejudice or overfitting towards a particular set of communities and links.

Fig. 1. This schematic illustrates a high-level difference between normal data processing and group-based processing.

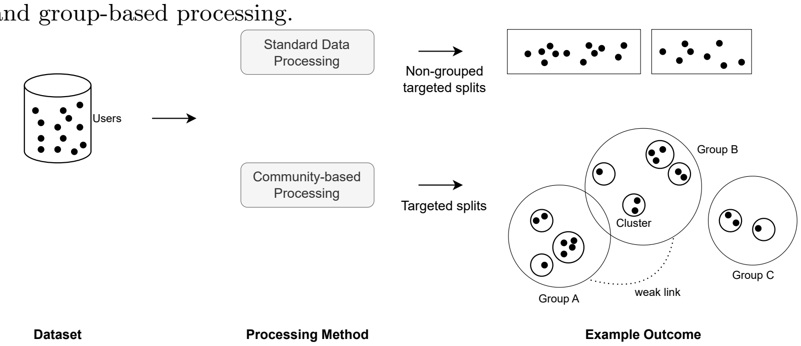

Serendipity-based Evaluation While accuracy is important for recommendation systems, it's not the only metric that matters. Incorporating serendipity, defined as making unexpected and relevant recommendations, can make recommendations more appealing and useful. Serendipity can be measured using metrics like surprise, unexpectedness, and relevance. Previous studies on serendipity-oriented recommender systems show that incorporating serendipity can help make recommendations more diverse and address issues related to information over-specialization.

Following the study by Eugene Yan [24], serendipity can be measured using the following formula:

$$serendipity(i) = unexpectedness(i) \times relevance(i) \tag{1}$$

Where $relevance(i) = 1$ if i is interacted upon and 0 otherwise. Alternatively, we use one of several approaches to measure the unexpectedness of recommendations [24,28]. This approach considers some distance metric (e.g., cosine similarity). We compute the cosine similarity between a user's recommended items (I) and historical item interactions (H). Lower cosine similarity indicates higher unexpectedness:

$$unexpectedness(I, H) = \frac{1}{I} \sum_{i \in I} \sum_{h \in H} cos(i, h) \tag{2}$$

The overall serendipity can be achieved by averaging all users (U) and all recommended items (I):

$$serendipity(i) = \frac{1}{count(U)} \sum_{u \in U} \sum_{i \in I} \frac{serendipity(i)}{count(I)} \tag{3}$$

The following section explains how we form user clusters and groups. As we measure serendipity on the group level, we use a recently proposed group-based validation technique [13] to track performance on smaller data portions, which helps avoid averaging results that may mask important effects. Therefore, changes in serendipity are measured by changes in its level within groups (e.g., group A in Fig. 1) rather than the overall user serendipity of equation 3.

User Clusters and Groups In this section, we cover the process of forming clusters and higher-level groups after discussing the method of evaluating groups of users in the previous section. Two levels are involved in this process - clustering users together at the first level and forming larger groups that can connect and include user groups with weak and strong links at the second level. We experiment with various versions of weak links between user clusters, as there can be multiple variations of higher-level cluster groups. This demonstrates the method's adaptability to accommodate different datasets.

To create the first level of user clusters for datasets like ML-100k, which often have both categorical and non-categorical data, we employ the Gower distance method [27] to produce a distance matrix. This approach involves calculating the distance between two entities based on their mixed categorical and numerical attribute values. We then use hierarchical clustering to refine the grouping further. For some given features $x_i = x_{i1}, ..., x_{ip}$ in a dataset, the Gower similarity matrix can be defined as:

$$S_{Gower}(x_i, x_j) = \frac{\sum_{k=1}^{p} s_{ijk} \delta_{ijk}}{\sum_{k=1}^{p} \delta_{ijk}} \tag{4}$$

For each feature $k = 1, ..., p$ a score s_{ijk} is calculated. A quantity δ_{ijk} is also calculated with a binary possible value depending on whether the input variables x_i and x_j can be compared. $S_{Gower}(x_i, x_j)$ is a similarity score, so the final result is converted through the following equation to achieve a distance metric: $d_{Gower} = \sqrt{1 - S_{Gower}}$. For numerical variables, the score can be calculated as a simple L1 distance between the two values normalized by the range of the feature R_k:

$$s_{ijk} = 1 - \frac{|x_{ik} - x_{jk}|}{R_k} \tag{5}$$

For categorical variables, the score will be 1 if the categories are the same and 0 if they are not:

$$S_{ijk} = 1 x_{ik} = x_{jk} \tag{6}$$

Several linkage methods exist to compute distance $d(s,t)$ between two clusters s and t using the distance matrix achieved with Equation 4. We utilize the general-purpose clustering algorithm proposed by Müllner [29]. The algorithm begins with a forest of clusters that have yet to be used in the hierarchy being formed. When two clusters s and t from this forest are combined into a single cluster u, s and t are removed from the forest, and u is added to the forest. When only one cluster remains in the forest, the algorithm stops, and this cluster becomes the root.

In the following section, we present experimental results for the method introduced above. The experiment has three main goals:

- Investigating the impact of recommending items via weak-linked groups.
- Determining whether optimizations in one group can impact others.
- Showcasing the effect of utilizing weaker connections alongside group linkage tuning and whether more favorable outcomes can be achieved.

4 Results and Discussions

In this section, we discuss the results of experiments carried out on two open-source datasets, namely the ML-100k [30] and the epinions [31]. Our work doesn't focus on a specific recommender algorithm, but rather on the experimentation process. We use LightGCN [32] as an example recommender. LightGCN is a simplified version of Neural Graph Collaborative Filtering (NGCF) that incorporates GCNs and is relatively new. We have created multiple versions of the code and experiment scenarios, all of which are also available in the source [33].

Our initial goal is to measure the impact on group serendipity. We plan to achieve this by selecting a formed group and tuning the recommender through the training process to allow favored recommendations to it from weakly-linked communities. The second objective is to determine whether this approach affects only the target group or any other group in the dataset that shares common users.

Figure 2 displays the average group serendipity obtained from one of our experiments on ML-100k. These groups were created through the process explained in Section 3. We observed a notable increase in the serendipity factor in multiple groups compared to the baseline process, ranging from 5% to 18%. This baseline process involved normal data processing and training using the same recommender parameters and tuning. However, only two out of the ten groups showed a decrease in the metric result.

Fig. 2. A comparison between the average serendipity value of a group and the baseline value achieved during regular data processing.

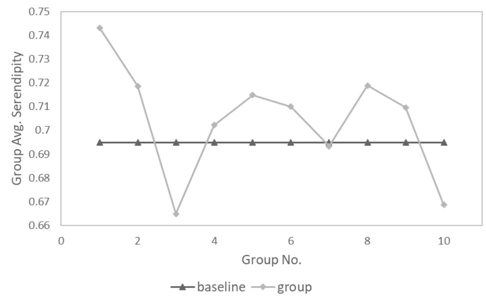

Table 1. The evaluation metric values for both the baseline and community-based processes.

Metric	Baseline	Group Avg.	Change (%)
Precision	0.2897	0.2288	-21.04
MAP	0.1424	0.0993	-30.26
NDCG	0.3716	0.2874	-22.65
Recall	0.2440	0.1866	-23.50
Coverage	0.3610	0.1741	-51.76

After analyzing the offline metric results of the system, it is evident that all of them experienced a decrease compared to the baseline run. This decrease can be attributed to the increase in serendipity, which leads to a corresponding decline in precision and recall. The results can be viewed in Table 1. However, we must keep in mind that offline evaluation is not enough to determine true relevance. Through online experimentation, we can accurately gauge the effectiveness of the model [24]. One of the interesting findings is a decrease in coverage. As explained in Section 2, increasing coverage (or introducing more randomness) in the dataset typically leads to an increase in serendipity during offline evaluation, which is a limitation of using this measure offline instead of in an A/B test. However, we took steps to minimize this effect by ensuring that our final recommendations were not biased and that we did not filter out items from the long tail. Online tests can improve the validation of the serendipity metric. In fact, several small tests have shown that users tend to converge more with recommenders that have lower accuracy and precision metrics [23,24]. Therefore, our results are in line with this trend.

Subsequently, an exploration is conducted to determine the potential impact of optimizing surprise in one group on the other group by utilizing the aforementioned approach. For the purpose of clarity, we have included the cluster-level outcomes (refer to Figure 1).

Fig. 3. Three scenarios that compare cluster serendipity to the baseline.

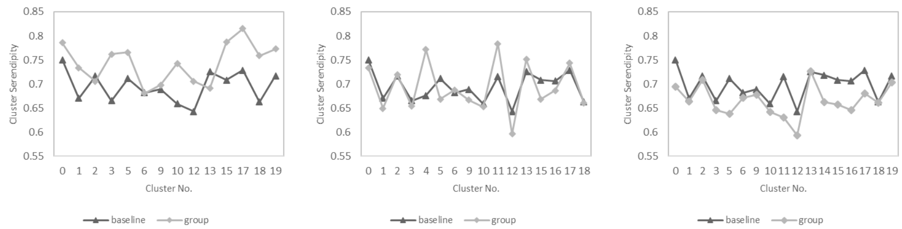

In Figure 3, we show the effect of the same serendipity metric but on the cluster level of the ML-100k dataset. The figure shows three cases when our approach is optimized to increase serendipity in one group while measuring the effect on the others. It can be noticed here how with no special tuning, the effect can be better (first figure), almost the same with small exceptions (middle figure), or worse (last figure). As mentioned in Section 3, the formation of user clusters and groups is sensitive, and there are multiple possibilities for weak links.

Fig. 4. The distribution of user serendipity metric values as group formations slightly vary.

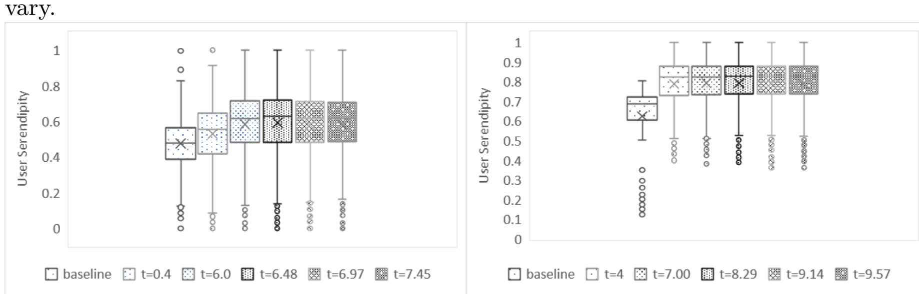

In Section 3, we discussed the hierarchical clustering method. This method can be used to simplify the dendrogram and assign data points to individual clusters. The assigned clusters are determined by a distance threshold, denoted as t. A smaller threshold will allow even the closest data points to form a cluster, while a larger threshold can result in too many clusters and few communities. By varying the value of t, we can produce different group representations that could affect the outcome of the metrics obtained in the initial stage of the experiments. To address this, we conduct multiple iterations that result in diverse weak-linked groups. Subsequently, we implement recommendations and re-evaluate the serendipity metric to determine any impact on the results.

The boxplot in Figure 4 displays the results. In the ml-100k scenario (on the left), we can observe an overall rise in user serendipity after the initial three iterations. This suggests an enhanced surprise element for most groups, as previously demonstrated in an experiment. We attain the highest value at approximately $t = 6.48$, which corresponds to the optimal distance between the groups formed. This distance has a positive impact on serendipitous recommendations for almost all groups. We were able to achieve the same outcome for the epinions dataset, although the parameter scale and the optimal distance between groups were slightly different. The best results were obtained with values between $t = 7$ and $t = 8.5$, and we found that further adjustments did not significantly improve results for most of the clusters. Using varied weak connections between groups can improve outcomes. Testing multiple scenarios helps find the best distance for each optimization run.

5 Conclusions and Future Work

This study emphasizes the significance of giving priority to chance discoveries for users in recommender systems. We developed a method based on social network theory which utilizes weak links to enhance recommender performance. As non-social links inferred from user behavior have not previously been utilized, we created a process for it in this work. This involves strategically forming communities rather than introducing random data. Our experiment yielded a positive result in enhancing the level of unexpectedness and surprise for users within the system. We demonstrated that recommending items through weak-linked communities among different users favors surprise and user engagement, as measured by a serendipity metric. This was achieved without any intentional randomness introduced into the data.

The clustering and grouping process can be improved by tuning with enhanced data for recommender systems, specifically data that reinforces social and rating-based behaviors within communities. Alternatively, the measure of serendipity is impacted by the items recommended from the long tail. While we took care to avoid any bias in suggesting long-tail items, it would be beneficial to conduct A/B tests to confirm convergence and not rely solely on offline tests. Finally, future research could explore the use of social clustering to validate whether different effects can be achieved on the serendipity of the model.

6 Acknowledgments

This work has been supported by the Lebanese University Research Program.

References

1. Dietmar Jannach and Michael Jugovac. Measuring the business value of recommender systems. *ACM Transactions on Management Information Systems (TMIS)*, 10(4):1–23, 2019.

2. Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. Deep learning based recommender system: A survey and new perspectives. *ACM computing surveys (CSUR)*, 52(1):1–38, 2019.
3. Tao Zhou, Zoltán Kuscsik, Jian-Guo Liu, Matúš Medo, Joseph Rushton Wakeling, and Yi-Cheng Zhang. Solving the apparent diversity-accuracy dilemma of recommender systems. *Proceedings of the National Academy of Sciences*, 107(10):4511–4515, 2010.
4. Andreas Argyriou, Miguel González-Fierro, and Le Zhang. Microsoft recommenders: Best practices for production-ready recommendation systems. In *Companion Proceedings of the Web Conference 2020*, pages 50–51, 2020.
5. Francesco Ricci, Lior Rokach, and Bracha Shapira. Recommender systems: Techniques, applications, and challenges. *Recommender Systems Handbook*, pages 1–35, 2021.
6. AmirHossein Nabizadeh, Alípio Jorge, and José Paulo Leal. Long term goal oriented recommender systems. *11th International Conference on Web Information Systems and Technologies*, 2015.
7. Jonathan L Herlocker, Joseph A Konstan, Loren G Terveen, and John T Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)*, 22(1):5–53, 2004.
8. Sean M McNee, John Riedl, and Joseph A Konstan. Being accurate is not enough: how accuracy metrics have hurt recommender systems. In *CHI'06 extended abstracts on Human factors in computing systems*, pages 1097–1101, 2006.
9. Tie-min Ma, Xue Wang, Fu-cai Zhou, and Shuang Wang. Research on diversity and accuracy of the recommendation system based on multi-objective optimization. *Neural Computing and Applications*, 35(7):5155–5163, 2023.
10. V Ramanjaneyulu Yannam, Jitendra Kumar, Korra Sathya Babu, and Bidyut Kumar Patra. Enhancing the accuracy of group recommendation using slope one. *The Journal of Supercomputing*, 79(1):499–540, 2023.
11. Wissam Al Jurdi, Jacques Bou Abdo, Jacques Demerjian, and Abdallah Makhoul. Critique on natural noise in recommender systems. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(5):1–30, 2021.
12. Wissam Al Jurdi, Jacques Bou Abdo, Jacques Demerjian, and Abdallah Makhoul. Strategic attacks on recommender systems: An obfuscation scenario. In *2022 IEEE/ACS 19th International Conference on Computer Systems and Applications (AICCSA)*, pages 1–8. IEEE, 2022.
13. Wissam Al Jurdi, Jacques Bou Abdo, Jacques Demerjian, and Abdallah Makhoul. Group validation in recommender systems: Framework for multi-layer performance evaluation. *arXiv preprint arXiv:2207.09320*, 2022.
14. Dhanya Pramod. Privacy-preserving techniques in recommender systems: state-of-the-art review and future research agenda. *Data Technologies and Applications*, 57(1):32–55, 2023.
15. Mohamed Amine Chatti, Mouadh Guesmi, and Arham Muslim. Visualization for recommendation explainability: A survey and new perspectives. *arXiv preprint arXiv:2305.11755*, 2023.
16. Mark S Granovetter. The strength of weak ties. *American journal of sociology*, 78(6):1360–1380, 1973.
17. Tomislav Duricic, Emanuel Lacic, Dominik Kowald, and Elisabeth Lex. Exploiting weak ties in trust-based recommender systems using regular equivalence. *arXiv preprint arXiv:1907.11620*, 2019.

18. Naren Ramakrishnan, Benjamin J Keller, Batul J Mirza, Ananth Y Grama, and George Karypis. When being weak is brave: Privacy in recommender systems. *arXiv preprint cs/0105028*, 2001.
19. Jyoti Shokeen and Chhavi Rana. Social recommender systems: techniques, domains, metrics, datasets and future scope. *Journal of Intelligent Information Systems*, 54(3):633–667, 2020.
20. Wissam Al Jurdi, Miriam El Khoury Badran, Chady Abou Jaoude, Jacques Bou Abdo, Jacques Demerjian, and Abdallah Makhoul. Serendipity-aware noise detection system for recommender systems. *Proceedings of the Information and Knowledge Engineering*, 2018.
21. Denis Kotkov, Alan Medlar, and Dorota Glowacka. Rethinking serendipity in recommender systems. In *Proceedings of the 2023 Conference on Human Information Interaction and Retrieval*, pages 383–387, 2023.
22. Yingheng Wang, Yaosen Min, Xin Chen, and Ji Wu. Multi-view graph contrastive representation learning for drug-drug interaction prediction. In *Proceedings of the Web Conference 2021*, pages 2921–2933, 2021.
23. Reza Jafari Ziarani and Reza Ravanmehr. Serendipity in recommender systems: a systematic literature review. *Journal of Computer Science and Technology*, 36:375–396, 2021.
24. Eugene Yan. Serendipity: Accuracy's unpopular best friend in recommender systems. *Towards Data Science, April*, 2020.
25. Upasna Bhandari, Kazunari Sugiyama, Anindya Datta, and Rajni Jindal. Serendipitous recommendation for mobile apps using item-item similarity graph. In *Information Retrieval Technology: 9th Asia Information Retrieval Societies Conference, AIRS 2013, Singapore, December 9-11, 2013. Proceedings 9*, pages 440–451. Springer, 2013.
26. Maximilian Jenders, T Lindhauer, Gjergji Kasneci, Ralf Krestel, and Felix Naumann. A serendipity model for news recommendation. In *KI 2015: Advances in Artificial Intelligence: 38th Annual German Conference on AI, Dresden, Germany, September 21-25, 2015, Proceedings 38*, pages 111–123. Springer, 2015.
27. John C Gower. A general coefficient of similarity and some of its properties. *Biometrics*, pages 857–871, 1971.
28. Adamopoulos Panagiotis. On unexpectedness in recommender systems: Or how to expect the unexpected. In *DiveRS 2011: Proceedings of the Workshop on Novelty and Diversity in Recommender Systems (DiveRS 2011), at the 5th ACM International Conference on Recommender Systems (RecSys 2011)*, pages 11–18, 2011.
29. Daniel Müllner. Modern hierarchical, agglomerative clustering algorithms. *arXiv preprint arXiv:1109.2378*, 2011.
30. F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4):1–19, 2015.
31. Matthew Richardson, Rakesh Agrawal, and Pedro Domingos. Trust management for the semantic web. In *International semantic Web conference*, pages 351–368. Springer, 2003.
32. Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. Neural graph collaborative filtering. In *Proceedings of the 42nd international ACM SIGIR conference on Research and development in Information Retrieval*, pages 165–174, 2019.
33. Wissam Al Jurdi and Jacques Bou Abdo. Github Repository: Optimizing Recommendations: A Contemporary Networks-inspired Approach, June 2023.