Rethinking Serendipity in Recommender Systems

Denis Kotkov
Department of Computer Science
University of Helsinki
Helsinki, Finland
kotkov.denis.ig@gmail.com

Alan Medlar
Department of Computer Science
University of Helsinki
Helsinki, Finland
alan.j.medlar@helsinki.fi

Dorota Głowacka
Department of Computer Science
University of Helsinki
Helsinki, Finland
dorota.glowacka@helsinki.fi

ABSTRACT

Recommender systems suggest items, such as movies or books, to users based on their interests. These systems often suggest items that users are either already familiar with or could easily have found on their own without additional assistance. To overcome these problems, recommender systems aim to suggest serendipitous items. While there is a lack of consensus in the recommender systems research community on the definition of serendipity, it is often conceptualized as a complex combination of relevance, novelty and unexpectedness. However, the common understanding and original meaning of serendipity is conceptually broader, requiring serendipitous encounters to be neither novel nor unexpected. Recent work in the social sciences has highlighted the various ways that serendipity can manifest, leading to a more generalized definition of serendipity. We argue that the study of serendipity in recommender systems would benefit from considering items that are serendipitous under this more general definition, giving us a deeper understanding of the item characteristics and behavioral impact of serendipitous recommendations. These findings will help us to better optimize recommender systems for serendipity. In this paper, we explore various definitions of serendipity and propose a novel formalization of what it means for recommendations to be serendipitous. Lastly, we present an experimental design for how serendipity can be measured in a deployed recommender system.

CCS CONCEPTS

• Information systems → Recommender systems; **• Human-centered computing → Field studies**.

KEYWORDS

serendipity, recommender systems, experimental design

ACM Reference Format:
Denis Kotkov, Alan Medlar, and Dorota Głowacka. 2023. Rethinking Serendipity in Recommender Systems. In *ACM SIGIR Conference on Human Information Interaction and Retrieval (CHIIR '23), March 19–23, 2023, Austin, TX, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3576840.3578310

CHIIR '23, March 19–23, 2023, Austin, TX, USA

ACM ISBN 979-8-4007-0035-4/23/03…$15.00
https://doi.org/10.1145/3576840.3578310

1 INTRODUCTION

Recommender systems are software tools that suggest items of interest to users [16], where an item is a "*piece of information that refers to a tangible or digital object, such as a good, a service or a process*" [9]. By design, recommender systems often recommend items that are either popular or similar to what users have already consumed [4]. However, users are likely to already be familiar with these items or are capable of finding them without the help of a recommender system [10]. Furthermore, users tend to prefer recommendations that are novel and interesting [10]. To overcome these issues, recommender systems aim to suggest serendipitous items [10].

In the recommender systems literature, serendipity is usually defined as a complex concept which consists of one or more of the following components: relevance, novelty and unexpectedness. Relevance refers to an item that is beneficial to the user, and novelty indicates that the user is unfamiliar with the item [8, 10]. Unexpectedness has numerous definitions, the most common being that an item is dissimilar to what the user tends to consume [8]. This conceptualisation is at odds with what is commonly understood by serendipity. Merriam-Webster, for example, defines serendipity as "*the faculty or phenomenon of finding valuable or agreeable things not sought for*"[1], which **does not mention that a serendipitous encounter needs to be novel or unexpected to the discoverer.** Even outside of recommender systems, serendipity is a difficult term to translate between languages[2] and users may understand the concept in their own way [17].

Inconsistent definitions of serendipity create three key problems in recommender systems. First, **items considered serendipitous according to the common-usage definition might have different characteristics compared to those considered serendipitous in the recommender systems literature.** Without being able to identify such items, we cannot investigate their behavioral impact on users, nor can we optimize a recommender system to suggest them. Second, different definitions create confusion within the research community and limit our ability to compare results between experimental studies. Even among researchers with the same broad definition of serendipity, the lack of consensus on how to measure unexpectedness leads to similar issues [8]. Finally, recommender systems are user-facing systems and researchers often seek feedback from users. In the case of serendipity, it is unclear how researchers would translate between users' looser definitions of the term and the recommender systems definition, leading to incorrect conclusions [11].

In this paper, (1) we show that the definition of serendipity used in recommender systems is conceptually narrow compared to other

[1]https://www.merriam-webster.com/dictionary/serendipity
[2]https://www.todaytranslations.com/news/most-untranslatable-word/

definitions; (2) we propose a method to expand the definition of what it means for recommendations to be serendipitous; and (3) we present an experimental design for measuring serendipity in a deployed recommender system.

2 GENERALIZED SERENDIPITY

The term *serendipity* was coined by Horace Walpole in his letter to Sir Horace Mann in 1754. The author referenced the Persian fairy tale "The Three Princes of Serendip" when describing his recent discovery. In the story, the three princes of the country Serendip were exploring the world and "*making discoveries, by accidents & sagacity, of things which they were not in quest of*" [15]. This definition is also consistent with modern dictionary definitions, such as the one from the Britannica Dictionary[3]: "*luck that takes the form of finding valuable or pleasant things that are not looked for*".

Although there seems to be a certain level of agreement on the definition between dictionaries and the letter to Sir Horace Mann, the definition allows for multiple interpretations, covers a wide range of phenomena and can often cause confusion [20]. To illustrate the usage of the term in practice, we rely on the recent literature review by Ohid Yaqub [20], that presents a taxonomy of serendipity based on the archives of Robert K Merton, who spent decades collecting mentions of the term "serendipity" in magazines, newspapers, and journals. This collection later resulted in the popular book "The Travels and Adventures of Serendipity: A Study in Sociological Semantics and the Sociology of Science" [5].

According to Yaqub [20], there are four categories of serendipity: Walpolian, Mertonian, Bushian and Stephanian, which are summarized in Table 1. **Walpolian serendipity** is the discovery of things that the discoverer is not in quest of. For example, a man named Bob who does not plan to look for any gifts, visits a grocery store to buy food, comes across a toy which is a perfect gift for his daughter and buys it. Here, the toy is a serendipitous discovery. **Mertonian serendipity** extends Walpolian serendipity by including discoveries of things that the discoverer looks for but finds through an unexpected route. An example of this could be if Bob visits multiple toy stores to find a birthday gift for his daughter, but fails to find anything suitable. However, he comes across the kind of toy he is looking for in a grocery store while shopping for food. **Bushian serendipity** extends the definition further by including discoveries made by the discoverer when they do not have a specific problem in mind. In this case, while waiting for a bus, Bob's gaze falls on a shop window where he discovers the perfect toy and buys it for his daughter. Lastly, **Stephanian serendipity** includes discoveries that solve problems that appear after the discovery has been made. For example, Bob finds a vintage lighter while going through old boxes in his attic. Months later, he finds out that one of his friends collects vintage lighters, making it a perfect gift.

3 SERENDIPITY IN RECOMMENDER SYSTEMS

To understand how serendipity has been defined in recommender systems, we conducted a literature review in ACM Digital Library.

[3] https://www.britannica.com/dictionary/serendipity

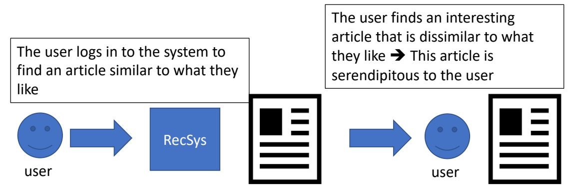

Figure 1: An example of a serendipitous item according to both RecSys and generalized definitions

Our search query was "[Title: serendipity] AND [[Title: recommender] OR [Title: recommendation]]", which resulted in 24 articles. We inspected these articles, two popular literature review articles on the topic [10, 23] and articles that cite the found papers.

Almost all the articles we found defined serendipity through its components: relevance, novelty and unexpectedness, which is consistent with [10]. In these definitions, relevance indicates that an item needs to be beneficial to the user, novelty – that the user should not be familiar with the item before consuming it, while unexpectedness has a number of definitions, such as that an item should be dissimilar to what the user usually consumes [8]. Each study required an item to correspond to one or more components mentioned above. For example, the first article we found, which explicitly defined serendipity in recommender systems, was published in 2006. The definition from the article required each of the three components to be present: "*serendipity in a recommender is the experience of receiving an unexpected and fortuitous item recommendation. But even if we remove that component, the unexpectedness part of this concept the novelty of the received recommendations is still difficult to measure*" [13]. Meanwhile, the authors of a recent method for recommending serendipitous courses to students at a university included only relevance and unexpectedness in serendipity: "*We use the definition of serendipity as user perceived unexpectedness of result combined with successfulness*" [14]. The only article we found that uses the generalized definition of serendipity in recommender systems is [18]: "*a serendipitous find relates to an unplanned yet interesting encounter*".

We note that the recommender systems definition of serendipity (RecSys definition) does not make any stipulations with respect to user goals, but assumes that users are looking for certain kinds of items, such as those they usually consume. Items can, however, be considered serendipitous regardless of what users are looking for and for what purpose, which can include browsing through items, influencing other users of the system or figuring out how the system works [6]. **Not knowing what the user's goals are and what the item will be used for will result in the misclassification of serendipitous and non-serendipitous items.** For example, in Figure 1, the article is serendipitous to the user, only because it is relevant to the user and dissimilar to their tastes, while the user is looking for articles similar to their tastes. However, in practice, the user does not always look exclusively for articles similar to their tastes.

There is no consensus on the definition of serendipity among users. They might be confused by the term [17], which is important to take into account when analyzing user feedback. We suggest

Table 1: A typology of serendipity [20]

		What type of solution did the discovery lead to?	
Is there a targeted line of inquiry?	Yes: Searching with a defined problem in mind	Solution of the given problem through an unexpected route: **Mertonian serendipity**	Solution of a different problem: **Walpolian serendipity**
	No: Searching with no particular problem in mind	Solution of a pre-existing problem through an unexpected route: **Bushian serendipity**	Solution is waiting for a problem: **Stephanian serendipity**

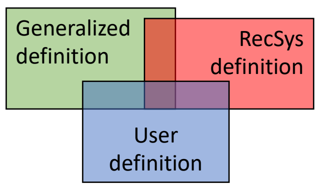

Figure 2: Euler diagram of items classified as serendipitous according to different definitions

that serendipity according to users and recommender system community capture some items that are serendipitous according to the generalized definition, but also misclassify other items (Figure 2).

4 GENERALIZED SERENDIPITY IN RECOMMENDER SYSTEMS

Based on the generalized serendipity definition, in recommender systems an item is serendipitous if it helps the target user to achieve any of the goals that are different from the goals they set out to achieve through a given recommender system at a given time (this includes users with no specific goals). An item is also considered serendipitous if it helps the user to achieve any goals that will appear in the future. We also define serendipity more formally as follows:

Assume that all goals of the target user at moment in time t_k is the set G_{t_k}. Meanwhile, the timeline is discrete $T = \{t_1, t_2, ..., t_n\}$, such that at each moment user goals are different compared to the goals in adjacent moments in time: $(G_{t_{k-1}} \neq G_{t_k}) \wedge (G_{t_k} \neq G_{t_{k+1}})$. Goals that the target user wants to achieve in a recommender system at t_k are $R_{t_k} \subseteq G_{t_k}$, while $R_{t_k} = \emptyset$ if the user does not interact with the system or has no specific goals during the interaction. An item recommended by the system to the target user is serendipitous if it helps them to achieve any goals from the set $S_{t_k} = \bigcup_{j=k}^{n} G_{t_j} \setminus R_{t_j}$.

We illustrate generalized serendipity with examples in Figure 3. Let us assume that a user logs into a recommender system which suggests articles. An article is serendipitous to the user according to Walpolian serendipity, if the user looks for an entertainment article but finds an article for work (Figure 3(a)) ("targeted search solves unexpected problem" [20]). According to Mertonian serendipity, an article is serendipitous if the user looks for an article for work using a search engine and fails to find it, but accidentally finds the article in a recommender system (Figure 3(b)) ("targeted search solves problem-in-hand via unexpected route" [20]). According to Bushian serendipity, an article is serendipitous if the user has no goals but when they visit the system, they come across an article

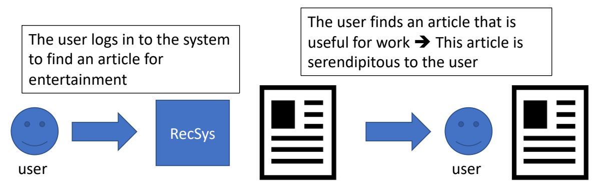

(a) Walpolian serendipity

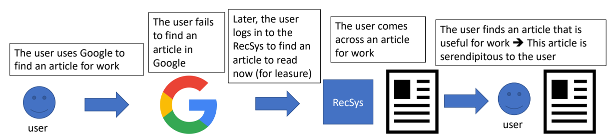

(b) Mertonian serendipity

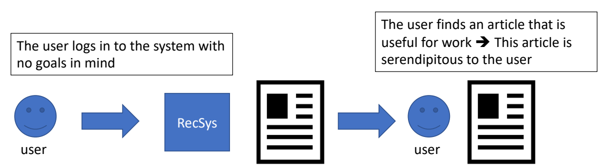

(c) Bushian serendipity

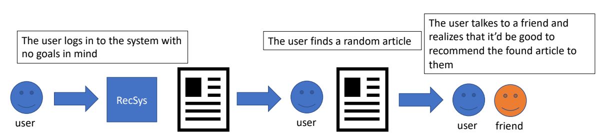

(d) Stephanian serendipity

Figure 3: Examples of serendipity in recommender systems, where the term RecSys corresponds to recommender system

for work (Figure 3(c)) ("untargeted search solves an immediate problem" [20]). According to Stephanian serendipity, an article is serendipitous if the user has no goals in mind, comes across an article but only later, when meeting a friend, realizes that that article would be useful to share with that friend (Figure 3(d)) ("untargeted search solves a later problem" [20]).

5 AN EXPERIMENTAL DESIGN FOR MEASURING SERENDIPITY

To date, researchers have been measuring serendipity by asking users questions only after the users have received recommendations [3, 8, 21]. For example, Kotkov et al. labeled a movie serendipitous if a user indicated that it was relevant, novel and unexpected to them only after watching the movie [8]. This experimental design does not allow for the collection of information on the initial goals of the user when they looked for recommendations, which is essential to labeling an item serendipitous according to the generalized definition.

In 2022, Smets et al. proposed an experimental setting for measuring generalized serendipity in recommender systems in a laboratory setting [18]. According to the proposed design, the user is given a particular task of finding certain items. For example, they could be asked to find articles to discuss at a meeting. The user can also label certain items as favorite. After the experiment, the user is given a survey regarding their favorite items and asked to identify if the items were serendipitous. This experimental design allows researchers to measure serendipity more precisely than the designs used in the past. However, as mentioned by the authors, the design is only applicable in a laboratory setting [18]. In this section, we build on the previous literature and propose an experimental design for measuring serendipity in a field experiment (in a functioning recommender system).

To assess serendipity in a recommender system, we propose the following procedure. As the user logs into the system, they should complete a survey about the goals of their visit. During each session, the user should have an opportunity to change their goals. After consuming each item, the user should complete the same survey indicating whether that item helped them to achieve any of their goals. The predefined goals in the survey should be identified prior to the study. The user should be able to add customized goals. The item is serendipitous to the user if it helps the user achieve at least one goal that is different from those the user has currently set.

5.1 Example

A user logs into a news recommender system. After logging in, the system requires the user to complete the following form (the user can select multiple goals and add new ones).

Please indicate all the goals of your visit:

- *Browse through articles*
- *Find articles within my interests*
- *Explore articles outside of my interests*
- *Provide feedback*
- *No goals*
- *Add a custom goal*

The user selects the goals: "Find articles within my interests" and "Browse through articles". The user also adds the customized goal "Find entertaining articles". The user browses through articles and reads one of them. After reading the article the user is asked to fill in the same form, except that the list of options now also contains "Find entertaining articles" and the instruction reads: "Please indicate all the goals that this article helped you to achieve". The user selects the goals: "Find entertaining articles","Browse through articles", adds and selects one more goal "Find an article to share with a friend". The system considers this article to be serendipitous because it helped the user to achieve a goal that the user initially did not plan to achieve.

The user continues browsing through articles, finds another article and reads it. After reading the article, the user indicates that it helped them to achieve the following goals: "Browse through articles" and "Find entertaining articles". The system considers this article non-serendipitous because it did not help the user achieve any new goals.

6 LIMITATIONS

Our approach has two main limitations. First, the proposed experimental design can only measure serendipity with the problem appearing before solution (Walpolian, Mertonian, Bushian definitions) and does not cover the situation where the problem appears after the solution (Stephanian definition) due to the difficulty of capturing such events. Second, asking users about serendipity might affect their experience of interacting with the system, so different interactions mechanisms should be investigated to make the process easier or otherwise worthwhile for the user.

7 CONCLUSION

In this paper, we showed that there is a mismatch between the RecSys definition and other valid definitions of serendipity. Furthermore, we proposed a method for expanding the definition of serendipity in recommender systems, and presented an experimental design for measuring serendipity in recommender systems. In recommender systems, serendipitous items have been claimed to achieve three benefits: to increase user satisfaction [1, 3, 8], to broaden user preferences [6, 8, 21, 22] and to overcome the overspecialization problem [1, 7]. Overspecialization happens when a user cannot discover new kinds of items as the recommender system only suggests items similar to what the user usually consumes [2]. Items that are serendipitous according to the generalized definition can potentially enhance these benefits and provide novel insights into serendipitous recommendations. We hypothesize that generalized serendipity is a more suitable basis to develop novel serendipitous recommenders than how serendipity is currently conceptualized in recommender systems and plan to conduct this experiment in future work.

We only covered the generalized and RecSys definitions of serendipity, but the term has also been defined in other fields. In information retrieval, for example, serendipity is often defined similarly to the generalized definition: "occurs when a user acquires useful information while interacting with a node of information for which there were no explicit *a priori* intentions" [19]. Researchers in this field also rely on user goals similar to those presented in this paper: "[w]e therefore define 'coming across information serendipitously' as 'finding useful or potentially useful information unexpectedly – either when not looking for information at all, when looking for information about something else or when looking for information with no particular aim in mind."' [12] However, fields outside of recommender systems are not within the scope of this article.

ACKNOWLEDGMENTS

We would like to thank the Academy of Finland Flagship programme: Finnish Center for Artificial Intelligence FCAI for supporting this work. We would also like to thank Ananth Mahadevan for discussions of the presented ideas.

REFERENCES

[1] Panagiotis Adamopoulos and Alexander Tuzhilin. 2014. On Unexpectedness in Recommender Systems: Or How to Better Expect the Unexpected. *ACM Transactions on Intelligent Systems and Technology* 5, 4 (2014), 1–32.

[2] Òscar Celma Herrada. 2009. *Music recommendation and discovery in the long tail.* Ph. D. Dissertation. Universitat Pompeu Fabra.

[3] Li Chen, Yonghua Yang, Ningxia Wang, Keping Yang, and Quan Yuan. 2019. How serendipity improves user satisfaction with recommendations? a large-scale user evaluation. In *The World Wide Web Conference*. 240–250.

[4] Michael D. Ekstrand, John T. Riedl, and Joseph A. Konstan. 2011. Collaborative Filtering Recommender Systems. *Found. Trends Hum.-Comput. Interact.* 4, 2 (Feb. 2011), 81–173. https://doi.org/10.1561/1100000009

[5] Harold M Green. 2004. The Travels and Adventures of Serendipity: a Study in Sociological Semantics and the Sociology of Science. *International Social Science Review* 79, 3-4 (2004), 174–176.

[6] Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl. 2004. Evaluating Collaborative Filtering Recommender Systems. *ACM Trans. Inf. Syst.* 22, 1 (Jan. 2004), 5–53. https://doi.org/10.1145/963770.963772

[7] Leo Iaquinta, Giovanni Semeraro, Marco de Gemmis, Pasquale Lops, and Piero Molino. 2010. *Can a recommender system induce serendipitous encounters?* InTech.

[8] Denis Kotkov, Joseph A Konstan, Qian Zhao, and Jari Veijalainen. 2018. Investigating serendipity in recommender systems based on real user feedback. In *Proceedings of the 33rd Annual ACM Symposium on Applied Computing*. 1341–1350.

[9] Denis Kotkov, Jari Veijalainen, and Shuaiqiang Wang. 2016. Challenges of Serendipity in Recommender Systems. In *Proceedings of the 12th International Conference on Web Information Systems and Technologies.*, Vol. 2. SCITEPRESS, 251–256.

[10] Denis Kotkov, Shuaiqiang Wang, and Jari Veijalainen. 2016. A survey of serendipity in recommender systems. *Knowledge-Based Systems* 111 (2016), 180–192.

[11] Yang Liu, Alan Medlar, and Dorota Głowacka. 2022. Lexical ambiguity detection in professional discourse. *Information Processing & Management* 59, 5 (2022), 103000.

[12] Stephann Makri, Jaffor Bhuiya, Jermaine Carthy, and Jermaine Owusu-Bonsu. 2015. Observing serendipity in digital information environments. *Proceedings of the Association for Information Science and Technology* 52, 1 (2015), 1–10.

[13] Sean M. McNee, John Riedl, and Joseph A. Konstan. 2006. Being Accurate is Not Enough: How Accuracy Metrics Have Hurt Recommender Systems. In *CHI '06 Extended Abstracts on Human Factors in Computing Systems* (Montréal, Québec, Canada) *(CHI EA '06)*. ACM, New York, NY, USA, 1097–1101. https://doi.org/10.1145/1125451.1125659

[14] Zachary A Pardos and Weijie Jiang. 2020. Designing for serendipity in a university course recommendation system. In *Proceedings of the tenth international conference on learning analytics & knowledge*. 350–359.

[15] Theodore G Remer. 1965. *Serendipity and the three princes: From the Peregrinaggio of 1557.* Norman, U. Oklahoma P, 20.

[16] Francesco Ricci, Lior Rokach, and Bracha Shapira. 2011. *Recommender Systems Handbook*. Springer US, Chapter Introduction to Recommender Systems Handbook, 1–35.

[17] Alan Said, Ben Fields, Brijnesh J. Jain, and Sahin Albayrak. 2013. User-centric Evaluation of a K-furthest Neighbor Collaborative Filtering Recommender Algorithm. In *Proceedings of CSCW'13* (San Antonio, Texas, USA). ACM, New York, NY, USA, 1399–1408. https://doi.org/10.1145/2441776.2441933

[18] Annelien Smets, Lien Michiels, Toine Bogers, and Lennart Björneborn. 2022. Serendipity in Recommender Systems Beyond the Algorithm: A Feature Repository and Experimental Design. In *16th ACM Conference on Recommender Systems*. CEUR Workshop Proceedings, 44–66.

[19] Elaine G Toms. 1998. Information exploration of the third kind: The concept of chance encounters. In *Information Exploration Workshop for SIGCHI*, Vol. 98. 85–92.

[20] Ohid Yaqub. 2018. Serendipity: Towards a taxonomy and a theory. *Research Policy* 47, 1 (2018), 169–179.

[21] Yuan Cao Zhang, Diarmuid Ó Séaghdha, Daniele Quercia, and Tamas Jambor. 2012. Auralist: Introducing Serendipity into Music Recommendation. In *Proceedings of WSDM'12*. ACM, New York, NY, USA, 13–22.

[22] Qianru Zheng, Chi-Kong Chan, and Horace H.S. Ip. 2015. An Unexpectedness-Augmented Utility Model for Making Serendipitous Recommendation. In *Advances in Data Mining: Applications and Theoretical Aspects*. Vol. 9165. Springer International Publishing, 216–230.

[23] Reza Jafari Ziarani and Reza Ravanmehr. 2021. Serendipity in recommender systems: a systematic literature review. *Journal of Computer Science and Technology* 36, 2 (2021), 375–396.