Long-term User Engagement in Recommender Systems: A Review

Swathi Edem[1], B V Ram Naresh Yadav[2]

[1] *Department of Computer Science and Engineering, Jawaharlal Nehru Technological University Hyderabad, Telangana, India*
[2] *Department of Computer Science and Engineering, Jawaharlal Nehru Technological University Hyderabad, Telangana, India*

E-mail address: swathiedem@gmail.com, bvramnaresh@jntuh.ac.in

Received ## Mon. 20##, Revised ## Mon. 20##, Accepted ## Mon. 20##, Published ## Mon. 20##

Abstract: The use of recommender systems is considered to be highly significant in our day-to-day life. The feed streaming approach has been used to a considerable degree inside the recommender system because of its effectiveness. Through the use of the feed streaming setting, users are given the opportunity to engage in an interactive manner of recommendation inside never-ending feeds. The concept of user stickiness, which goes beyond the conventional concept of instant analysis and is often evaluated by long-term user involvement, should get a higher amount of attention from a powerful recommender system. Specifically, this is due to the fact that user stickiness is not within the purview of conventional measures. A primary objective of recommender systems (RS) is to improve the level of user participation and communication that occurs inside the platform. In spite of this, there is a lack of comprehension of the extent of this link and the ways in which RS might potentially improve continuous user engagement with the platform to a certain degree. The present study endeavours to analyse the role that RS plays in turning users' short-term engagement with the RS into engagement with the platform that is maintained over time. This is done in order to solve the knowledge gap that has been identified. In order to investigate these concerns, we first construct a conceptual framework by doing a literature research on the relevant literature in the fields of user engagement and recommendation systems. This is done in order to get the necessary information. Within the scope of this research, we talk about metrics for the evaluation as well as open challenges that are presently being encountered in this aspect of the field.

Keywords: long-term metrics, user engagement, user satisfaction, Recommender Systems(RS).

1. INTRODUCTION

The use of recommender systems is on the rise because they make it easier to find and use information in many different areas of our lives, including shopping [11], food [14], travel [13], social platforms [17,18], media [12], and news[15,16]. In recent years, recommendation algorithms that prioritize instantaneous user responses like likes and clicks have achieved remarkable success [12, 19]. But it's becoming more and more obvious that putting too much weight on short-term engagement might cause clickbait or pigeon-holing effects, which are bad for users' experience in the long run [20, 21, 22]. Algorithm designers on recommendation systems have started optimising for other objectives that are better aligned with the long-term user experience, after seeing the pitfalls associated with excessive attention on short-term metrics. For instance, according to Wu et al.[58], a recommender system's ultimate objective is to encourage users to return to the platform more frequently rather than only meeting their needs during the current session.

It is challenging to optimise for long-term satisfaction since the expected long-term outcomes occur less often, are more complicated, and change slowly over time, as opposed to short-term engagement metrics. One can question if there are simpler options to optimize that are more indicative of the eventual outcome. An enhanced correlation between the suggestion and a more straightforward target for optimization is preferred. Optimising for long-term satisfaction may be tough due to the complexity and slower change of desired outcomes when compared to short-term engagement measures.

In this research, the distinctions between recommender systems that depend on explicit user rating input and those that rely on implicit user action feedback are presented in a clear and demonstrable manner. The ability to predict the activities that users will do in the future is more effective than the ability to predict explicit rankings for the purpose of enhancing engagement. It is not enough to just focus on predicting implicit behaviors in order to improve user engagement. This article will discuss the long-term metrics that should be examined in order to improve user engagement and satisfaction. Higher degrees of negative user feedback, including negative action rates and browsing effort, are shown by the fact that certain research indicate that recommenders that are primarily concerned with predicting implicit actions may not be as accurate as recommenders that are based on ratings. This finding hints to the need of doing more research to investigate different approaches to combining both positive and negative feedback, with a particular emphasis on penalizing products that have received negative feedback from users.

By combining the recommender systems and including online user interaction, the levels of user engagement achieved are comparable to those achieved by using an implicit-action-based recommender alone. However, it did not result in an increase in user browsing, suggesting a compromise between the objectives of user engagement and satisfaction. The observed outcome is expected to be relevant to other algorithms that simulate explicit or implicit feedback data, since it likely reflects the underlying characteristic of both forms of feedback signals rather than the particular methods used. Additional investigation is required to validate this. Nevertheless, it suggests that in order to precisely assess user happiness, it is essential to take into account elements that extend beyond just ratings or behaviors.

2. TAXONOMY OF RECOMMENDER SYSTEMS

A. *Classifications of Recommender Systems*

There are four main types of recommendation systems- Collaborative filtering, content based filtering, knowledge based methods, hybrid methods [43]. The Figure 1 gives the illustration of broad classification of Recommender Systems.

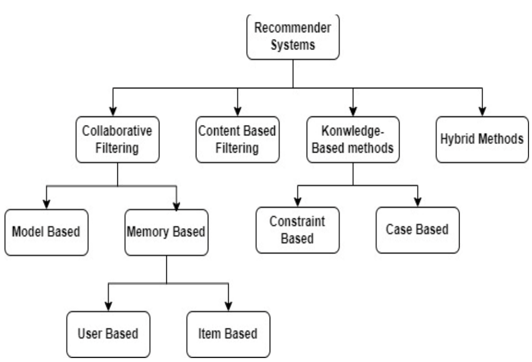

Figure 1. The primary types of Recommender Systems

Collaborative Filtering (CF) approaches: CF techniques function by collecting and assessing the actions of the user. The goal of CF approaches is to predict correctly the preferences of the user by assessing the user's data, which includes, among other things, the goods that have been seen in the past and the individual's purchasing history. Memory-based collaborative filtering, is predicated on the assumption that individuals who have preferences that are comparable in the past are likely to continue to have preferences that are similar in the future. The another name for this model is the user-item filtering model. The ratings of items may be predicted by using this technique, which takes the ratings of persons or organisations that are located in close proximity to one another into the account. Filters in model-based CF make use of a statistical or machine learning model in order to uncover and exploit on hidden relationships and patterns within the data. This is done to achieve the goal of model-based CF. It is done in this manner in order to discover patterns and connections that have been disguised and to make use of them. The training data, which is constituted of before the interactions between users and items, is used by these models to provide predictions about the preferences of users for unfamiliar products. This is done to enhance the accuracy of the forecasts. Two separate categories may be distinguished from one another when it comes to the memory-based collaborative filtering technique:

A collaborative filtering system based on users: A recommendation system that is known as collaborative filtering is one that analyses the preferences of users and provides ideas for goods that are enjoyed by users who are similar to the person who is being questioned.

A collaborative filtering system that is based on items and produces suggestions for things that are analogous to those that the active user has shown interest in is called an item-based collaborative filtering system on items.

Content-Based Rating Systems: These systems operate on the assumption that user preferences may be anticipated by analysing the user's prior experiences with goods, such as their history of viewing and purchasing things with the product in question. A content-based recommendation system is designed to provide suggestions to users that are similar to the things that they have previously interacted with. This is the system's primary purpose.

A recommender system becomes knowledge-based when it generates recommendations by analysing the user's specific queries rather than relying solely on their rating history. It could ask the user to provide a criteria for the desired outcome or a set of instructions. The system searches through its database of items and returns similar results [35]. KBR systems operate by gathering user preferences through dialogue and providing recommendations based on either predefined rules or similarity metrics that match the user's preferences. We can categorise the first approach as constraint-based recommendation, while the second approach falls under case-based recommendation.

A hybrid recommendation: The framework utilises a blend of content-based recommendation systems (CBRS) and collaborative filtering recommendation systems (CFRS) to achieve optimal performance by mitigating the limitations of traditional recommendation techniques [36,47].Maintaining the Integrity of the Specifications

B. *Long-term Participation of Users in the Recommendation Process*

When it comes to recommendation, one of the most essential metrics is long-term user engagement, which is often reflected as the degree to which people are committed to a certain product. In most cases, when we are presented with a product, we anticipate that people would either spend more time on it or utilize it as often as they possibly can. The session-based recommender systems have been extensively used in real-world applications, such as the recommendation of short-form films and the recommendation of news [36, 40, 45, 46]. The amount of suggested things that users consume during each visit is something that we are especially interested in growing, as is the frequency with which users visit the product. There is a significant amount of difficulty involved in optimizing these two indicators since it is difficult to tie them to a single proposal. In the event that a user increases the number of times they visit our website, for instance, we are unable to determine which suggestion is responsible for the increase in frequency.

C. *The Key Factors for Long-term User Engagement*

In order to effectively tackle the aforementioned issues, it is necessary to adopt certain techniques throughout the implementation of recommender systems. The long-term involvement of users in recommender systems is contingent upon certain crucial criteria [46]. Below are few crucial factors to take into account:

- Personalization: Recommender systems should consider user preferences and feedback to provide the customized feedback. The personalization could be achieved in many ways like user profiling, user modeling, using collaborative filtering techniques, etc. User long term engagement could be improved by providing the more relevant recommendations to the user based users' preferences[47].
- Accuracy: The accuracy of recommendations is crucial for user satisfaction and engagement. Recommender systems should strive to minimize errors and provide recommendations that match users' interests and needs as closely as possible.
- Diversity: Diverse recommendations keep the user attentive, interested and engaged to the recommender system. Consequently, it is essential to sometimes propose random recommendations in order to boost the level of user satisfaction and keep users engaged over the long run. Some research is done to recommend a random item, the social neighbors can be analysed or emotion of the user can be used [44].
- Transparency and Explainability: Users should be aware of the reasons why they are receiving specific suggestions. The transparent process of recommendation and explanations about the recommendations could help in the user confidence build up and user satisfaction which indirectly leads to the user's long term engagement with RS.
- Serendipity: It is important for recommender systems to periodically include random recommendations or unexpected suggestions that create a positive response from consumers. The unexpected occurrences can increase the user engagement and satisfaction via the exposure of users to products that they may not have otherwise come across.
- User Feedback Incorporation: The incorporation of systems that allow users to submit feedback on recommendations contributes to the enhancement of the accuracy and relevance of recommendations offered in the future. It should be possible for users to rate goods, offer explicit feedback, or modify their preferences over the course of time.
- Context Awareness: Recommender systems must consider the contextual factors while generating suggestions, such as the user's location, time of day, device, and social milieu. User-specific recommendations are more likely to be relevant and engaging to the user [48].
- Long-term Learning and Adaptation: In order to continually gain information from user interactions and modify their recommendations over time, recommender systems should be continuously learning. The ability of recommender systems to sustain relevance and interest over the long term may be achieved by the incorporation of feedback and the adaptation to changing user preferences.
- Seamless Integration: To make it effective, recommender systems need to be able to effortlessly integrate themselves into the workflows or routines that users engage in on a daily basis. The user experience must be smooth and require less work to encourage long-term engagement.
- Privacy and Trust: Gaining trust of customers is important aspect to increase the long-term user engagement. This possible only by creating trust in customers that their personal information will be handled in a responsible manner and that their privacy will be protected. Building and maintaining user trust may be facilitated by providing openness on the utilisation of data and putting in place protective measures for users' privacy.

By considering these factors, recommender systems can enhance long-term user engagement, leading to increased user satisfaction.

3. PREPARE RELATED WORK

A. *Analysing user behaviour on recommendation systems*

Before Extensive research has been conducted on analysing user behaviour in recommender systems across several disciplines such as human-computer interaction, marketing, and information retrieval. The distribution of user interests in long-tail and specialty material was investigated by Goel et al. [23]. User behaviour on recommender systems is affected by user preferences, algorithmic suggestions, and other factors such as personality qualities. Knijnenburg et al.[24] and Xiao et.al. [25] shed light on the processes that influence user experience in recommender systems. Anderson et al. [26] investigated the influence of suggestions on the variety of content consumed by users. In order to better assess the impact of recommendations, Villermet et al. [27] suggested differentiating between algorithmic and human behaviour when listening to music online. Models were proposed to comprehend the evolving user preferences by utilising a mix of structural and probabilistic methodologies[28, 29, 22]. The influence of individual and contextual elements on user behaviour in recommendation systems was demonstrated by Karumur et al.[30] and Xiao and Benbasat [25]. Another area of study involves creating simulations or doing field research to assess user behaviours while considering possible confounding variables[20,31,32]. The feedback loop between recommendation algorithms and user behaviours is studied by Hansen et al. [33] and Zhou et al. [34] in relation to patterns of consumption on video and music streaming services. While many studies have focused on how to measure user engagement with recommender systems, few have sought to understand users' long-term experiences by tracking their sequential and developing behaviours.

B. *RL for learning the long-term metrics*

In the discipline of artificial intelligence and machine learning, reinforcement learning (RL) is a subfield that focuses on teaching computers to make a series of decisions in an unpredictable and sometimes dynamic environment through the process of trial and error. In the context of recommender systems, RL can be used to optimize long-term user engagement by learning from user feedback and adapting to their changing preferences over time. This can involve rewarding recommended items that are well-received by users and penalizing items that are not, in order to incentivize the system to provide recommendations that are more likely to be relevant and engaging to individual users. Recommender systems, as shown in Figure 2, may be represented as an agent that interacts with users, who function as the environment. After each suggestion request is fulfilled by the agent, we may log the feedback and status changes from users. This data can be used to compute a reward and update the agent's current state. Utilizing reinforcement learning (RL) will result in the development of a recommendation policy that maximizes user engagement over an extended period of time.

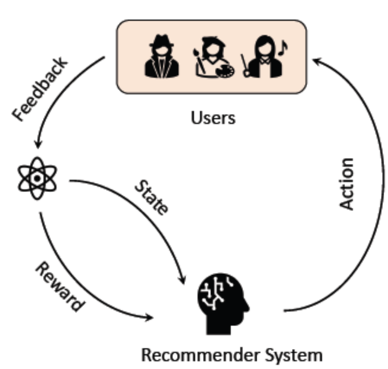

Figure 2. Recommender systems based on reinforcement learning

RL can also be used to optimize the exploration and exploitation of new recommendations, balancing the need to introduce new items[1]. Some works use reinforcement learning to look for the best weights in order to make users happy over the long run. Off-policy reinforcement learning is used by Han et al. [4] to find the best weights for an advertiser's expected click-through rate and bid price. Because interacting with young agents too much will ruin the user experience, they create an environment simulator to get feedback from users while they train their model offline. But the real recommendation world is too complicated for the simulator to fully represent it. The RL model that is built on the simulator will actually be hard to use online, which will make the experience worse for the users. Pei et al. [3] suggest using reinforcement learning to find the best way for a platform to make the most money. To make the model easier to understand, they use an evolutionary strategy to solve the problem. This means that the proposed methods can only improve the profile of the current suggestion. Short-horizon reinforcement learning is a type of machine learning algorithm that can improve long-term metrics in recommendation systems by focusing on immediate goals and learning from them to reach better outcomes in the long run. In situations in which the long-term goals are not fully defined or projected, this strategy is especially helpful since it enables the system to adjust to changes in user behaviour and preferences over the course of time applications. The technique has been applied to various domains, including music recommendation, news recommendation, and movie recommendation, where it has shown to outperform other traditional methods [2]. Various simulators like Sim2Rec [5], ADAREC [7] etc. have been used to test the RL based approaches which mainly focused on long-term user engagement.

The decision-making strategy known as Sim2Rec [5] is designed to maximise the quality of long-term user engagement by optimising real-world user involvement. It combines a simulation-based framework with machine learning techniques to simulate user behavior and make data-driven recommendations. This approach has been used in various industry applications, such as e-commerce and personalized medicine, to improve user engagement and retention. Sim2Rec can help businesses understand user preferences and tailor their recommendations to improve user experience and increase revenue. ADAREC [7] is an advanced recommendation system that optimises long-term user involvement through adaptive, sequential decision-making processes, rather than just generating item-to-item recommendations. The goal of ADAREC is to increase user happiness and foster long-term engagement by learning from user feedback and adjusting to shifting preferences.

There is a body of research on modelling users' return behaviour, which is now seen as a crucial indicator of long-term user engagement. In their study, Du et al. [9] used low rank models and a self-exciting point procedure to identify the user-item consumption patterns that occurred repeatedly over time. In their most recent work, the authors predict when user events will return by using marked temporal point processes and intensity functions. Kapoor et al. [8] suggested a Cox's proportional hazard function that takes use of survival analysis in order to estimate when consumers would return to utilising free internet services. This was necessary in order to create an accurate prediction. Subsequently, the authors used a hidden semi-Markov model in order to conduct an analysis of the time between a user's sequential consuming habits in accordance with their underlying psychological states. Display ad conversion delay forecasting was made possible by Chapelle's [10] survival analysis method. However, none of the referenced research focus on offline analysis or user return forecasts. When user input is highly dependent on service quality, some fail to consider how simulated customer feedback could be used to improve a service system.

4. METHODOLOGIES AND ANALYSIS

A. *Policy Learning:*

For the purpose of enhancing user involvement over the long term, a significant number of research works are focused on policy. The process of finding the ideal approach, or policy, that an agent ought to adopt in order to maximise its cumulative reward in a particular environment is referred to as policy learning in the field of reinforcement learning. This kind of learning is known as reinforcement learning, and it involves an agent interacting with its surroundings by performing actions and getting feedback in the form of rewards or penalties. It is the objective of the agent to acquire a policy that will enable them to choose actions that will result in the greatest potential cumulative reward over the course of time period.

Break is suggested as a means of promoting and maintaining the user over a longer amount of time [42]. According to this strategy, encouraging the user to take a break from the RS is a way to boost user satisfaction, which in turn leads to more engagement over a longer period of time.

We are able to use cutting-edge strategies in order to make the policy flexible enough to accommodate shifting preferences among users. A context encoder is used by ADAREC inside the policy network. This encoder makes it possible for RL rules to recognise various patterns of user activity [7].

Within the context of a batch reinforcement learning (RL)-based multi-task fusion, a conservative-OP estimator is developed.

Both a Batch RL framework and an online exploration component are included into the framework, which is referred to as BatchRL-MTF. The former makes use of batch reinforcement learning to train an optimum recommendation strategy from the fixed batch data offline for long-term user happiness, whilst the latter investigates alternative high-value actions immediately in order to overcome the local optimal dilemma[38]. It is recommended that a short-horizon policy improvement (SHPI) be implemented, which approximates policy-induced changes in user behaviour with respect to different sessions[2]. Optimising long-term user engagement with the usage of FeedRec [37]. There are two components that make up FeedRec: 1) a Q-Network, which is constructed in hierarchical LSTM, is responsible for simulating complicated user behaviours; and 2) an S-Network, which has the capacity to mimic the environment, provides assistance to the Q-Network, and eliminates the instability of convergence in policy learning.

B. *Point of Interest(POI) Recommendation:*

According to the short-term preferences, the next point of interest (POI) that a user will visit is impacted by the objects and venues that the user has recently visited in the trajectory that is now being followed. For instance, a person may go to a pub immediately after eating eaten at a restaurant the previous night. An individual's long-term preferences are a representation of their general interests, which are derived from their historical trajectories. In conclusion, there is a tendency for short-term preferences to fluctuate often throughout the course of time, but long-term choices tend to remain relatively consistent.

In the early stages of POI recommendation research, the primary emphasis was placed on assessing the preferences of users via the use of Collaborative Filtering (CF), particularly algorithms based on Matrix Factorization (MF). Only the consumers' choices that are static may be modelled using these approaches. These sorts of recommenders are unable to reflect the dynamism of user preferences, thus for instance, when a user who lives in India visits to Landon for a vacation, they may still recommend points of interest (POIs) that are situated in India. approaches that are based on deep learning have lately shown promising results in a variety of recommendation systems. Some examples of these approaches are embedding learning [51,53,54], neural CF [55], deep latent factor model [56], and metric learning [57].

For the purpose of modelling the long-term periodicity, DeepMove[52] makes use of a deep neural network that is equipped with two attention processes. For the purpose of modelling preferences over a short period of time, both DRCF and DeepMove have utilised RNN-based techniques. A gated mechanism that models both long-term and short-term interests is presented by STGN [58], which is also an effort to simulate both types of interests. This mechanism falls within the LSTM architecture.

RNN-based approaches are becoming more prevalent in the area of next-POI recommendation [52,5,60,61]. This is a direct result of the success that RNN has shown in sequential data modelling [62]. As an example, the ST-RNN model [61] is an extension of the RNN that models local temporal and geographical settings. Through the use of GRU's gate mechanism, CARA [60] is able to record the changing preferences of consumers. Both the LSTM-based and gated LSTM frameworks are used by TMCA [59] and STGN [58] in order to acquire knowledge about spatial-temporal contexts, respectively. To capture the sequential transition, Deep-Move [52] develops a multi-modal recurrent neural network (RNN). A geo-dilated recurrent neural network (RNN) is used for short-term preference learning, whereas a nonlocal network is used for long-term preference modelling using LSTPM[39].

C. *Explore and Exploitation methods:*

Exploration refers to the process of trying out different actions to gain more information about the environment and improve the agent's understanding of which actions lead to favorable outcomes. Exploitation, on the other hand, involves using the knowledge or information gained from past experiences to select actions that are expected to yield the highest immediate reward. Exploitation aims to maximize the short-term reward by choosing actions that the agent believes are currently the best based on its existing knowledge. Balancing exploration and exploitation is a key challenge in reinforcement learning, as the agent needs to strike a balance between trying out new actions to learn more about the environment (exploration) and selecting actions that are known to be good based on past experiences (exploitation).

The recommendation is based not just on the predicted number of clicks that the user will make immediately, but also on the anticipated number of clicks that will come from the user's subsequent return. The exploitation for immediate click, the exploitation for projected future clicks, and the exploration of unknowns for model estimate are the three competing variables that are taken into consideration while developing a bandit-based solution for online learning. This method is created on the basis of this idea.

D. *Reward Fomulation*

Reward formulation in reinforcement learning involves designing and defining the reward signal that the agent receives from the environment based on its actions. The reward signal serves as feedback to the agent, guiding it to learn a policy that maximizes the cumulative reward over time. Reward formulation is a critical aspect of reinforcement learning, as it directly influences the behavior of the agent and ultimately determines the success of the learning process.

A surrogate for long-term user experience was offered in a study [41], which suggested that it is possible to develop a surrogate by using data on previous user interactions with the recommender system and then utilising that data to model the user's upcoming and current preferences. This makes it possible to represent the long-term user experience by a fictitious object, which can then be utilised to develop the system and measure performance metrics. It has been suggested that PrefRec should be used to develop a reward function in order to boost long-term user engagement. Through the use of the preferences, it automatically trains a reward function in a way that is end-to-end. Following that, the reward function is used to provide learning signals for the purpose of training the recommendation policy [6].

In light of the findings of our investigation, we have created a list of the methodology and metrics used in previous publications, which may be seen in Table 1. We found only nine publications that we considered completely relevant to our subject matter. The proportion of publications that used a certain approach is shown in figure 3, which can be found here. The number of research articles that are published as a result of each year is shown in figure 4. We are able to draw the conclusion that there is less effort being done in the direction of optimising the long-term user involvement in the recommender system by looking at Fig.3 and Fig.4. In light of this, it is imperative that further study be conducted in this particular area.

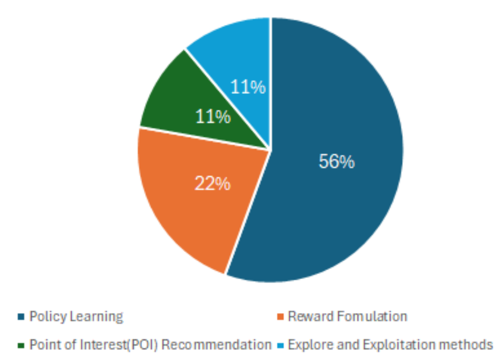

FIGURE 3. THE SUMMARY OF THE METHODS USED IN THIS STUDY

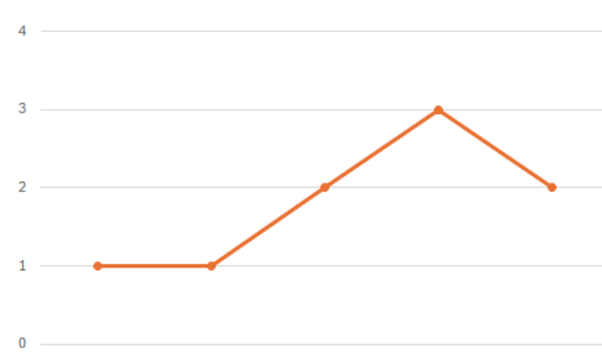

FIGURE 4. THE NUMBER OF PUBLISHED PER YEAR USED IN THIS STUDY

5. METRICS

After While there isn't a specific set of universally agreed-upon long-term metrics for user engagement in recommender systems, there are several commonly used metrics that can provide insights into long-term user engagement. Table 1 shows what metrics are adopted in the existing methods and in this section, those metrics are described. These metrics are often adapted based on the specific goals and context of the recommender system [45]. Here are some examples:

- Retention Rate: The proportion of users who continue to utilise the recommender system over a certain time period is referred to as the retention rate inside the system. The capacity of the system to retain the interest and engagement of users over an extended period of time is shown by this statistic.
- Session Length: The average duration of user sessions within the recommender system. Longer session lengths suggest deeper user engagement and interest in the recommended content.

TABLE-1: SUMMARY OF THE METHODS FOCUSED ON LONG-TERM USER ENGAGEMENT USING REINFORCEMENT LEARNING ALGORITHMS

Model	Year	Method	Dataset/Domain	Evaluation Metric
Lotka-Volterra dynamical system [42]	2023	Policy Learning	MovieLens 1M dataset, Goodreads dataset.	Mean long-term engagement rate (LTE)
ADAREC [7]	2023	Policy Learning	Real-world E-commerce dataset.	The cumulative retention reward, the users' average return days, the users' return probability on the next day
PrefRec[6]	2022	Reward Fomulation	Customized dataset of short-form videos.	Session depth, visiting frequency Sampling (NCIS)
Multi-Task Learning model(MTL) + Multi-Task Fusion model (MTF) [38]	2022	Policy Learning	A real-world short video dataset.	App dwell time (ADTime),User positive-interaction rate (UPIRate)
Reward surrogates in an RL-based recommender system [41]	2022	Reward Fomulation	Industrial recommendation platform.	Overall user visiting frequency,User visiting frequency from low frequency user segment (based on pre-experiment visiting frequency), Number of homepage visits, Number of satisfied consumptions
FeedRec[37]	2021	Policy Learning	Real-world E-commerce dataset.	Depth, ReturnTime
Short Horizon Policy Improvement (SHPI)[2]	2021	Policy Learning	A real-world private recommendation dataset and an HIV treatment domain	Click, Long Term Reward(LTR)
Long- and Short-Term Preference Modeling	2020	Point of Interest(POI) Recommendation	The Foursquare check-in dataset and Gowalla dataset	Normalized Discounted Cumulative

(LSTPM) [39]				Gain and Recall@K
r2 Bandit[40]	2017	Explore and Exploitation methods	A real-world Yahoo frontpage news dataset	Cumulative clicks over, Click-through rate (CTR),Average return time,Return rate,Improved user ratio, No return count

- Frequency of Interactions: This refers to the frequency with which consumers engage with the recommender system over a period of time, such as via clicks, views, or purchases. An increased interaction frequency is indicative of a user's continued continued involvement.
- Churn Rate: The rate at which users disengage or stop using the recommender system over time. A lower churn rate indicates higher user retention and long-term engagement.
- Conversion Rate: The proportion of people who buy or subscribe to a service after recommendations. Higher conversion rates show suggestion efficacy in motivating user behaviours.
- Customer Lifetime Value (CLV): The estimated value that each user contributes to the system over their lifetime. CLV helps assess the long-term impact of user engagement on the system's success and profitability.
- Customer Satisfaction Scores: Feedback or ratings provided by users about their satisfaction with the recommender system. Positive satisfaction scores indicate successful long-term engagement.

These metrics should be adjusted according to the recommender system's unique aims, the suggested content's type, and user interactions. To get a complete picture of user engagement over time, it's a good idea to incorporate both quantitative metrics and qualitative feedback.

6. THE CHALLENGES FOR LONG-TERM USER ENGAGEMENT IN RECOMMENDER SYSTEMS

The use of recommender systems has become more important across a variety of online marketplaces. They want to give consumers with personalised suggestions that are based on the interests and actions of the same users; however, the key challenge lies in maintaining long-term user engagement [50]. Common Difficulties in Achieving Long-Term User Engagement:

- Limited User Feedback: Often users provide sparse feedback or ratings on recommended items; this complicates the system's ability to accurately understand their preferences.
- Dynamic User Preferences: Users' preferences can change over time, leading to discrepancies between their current interests and the system's recommendations; this can result in frustration and disengagement.
- Cold Start Problem: When new users join a recommender system, there is a lack of historical data to personalize recommendations for them; this initial phase can deter user engagement before it even begins [49].
- Frequent changes in the User Preferences: It’s natural that a user interests change time to time. Sometimes, due to external factors user interests get changed. But in these cases, it is very difficult to predict the future preferences.
- Lack of randomness in recommendations: sometimes user may get bored if the recommendations are of the same types. So it is important to include variety types of recommendations.
- Information Overload: Though the recommendations are given based on user preferences, user may loose the satisfaction if the the recommendations are more. Too many recommendations will lead to the user confusion and dissatisfaction.

To address these challenges, recommender systems must continually adapt and evolve by incorporating innovative algorithms and techniques; this ongoing development is essential to ensure that users remain actively engaged with the system over extended periods. By understanding and addressing the factors that hinder long-term user engagement, recommender systems can enhance user satisfaction and loyalty, ultimately leading to greater success for online platforms and businesses.

7. CONCLUSION

We provided an exhaustive and current analysis of long-term user engagement with recommender systems. We placed a strong emphasis on the significance of user engagement and satisfaction for recommender systems, as well as the difficulties that are associated with this path. Based on our study we conclude that the reinforcement learning algorithms are suitable to optimize the long term user engagement. Therefore we presented a literature survey on the role of Reinforcement learning algorithms to increase the user long term engagement in Recommender Systems. There isn't a specific set of universally agreed-upon long-term metrics for user engagement in recommender systems. We presented the metrics considered in the previous works to evaluate the long-term engagement. We feel the research in this direction is very less and significant progress is needed. In conclusion, we are confident that our survey will aid the researcher in comprehending fundamental concepts and future developments to increase the long-term user engagement.

REFERENCES

[1] Lixin Zou, Long Xia, Zhuoye Ding, Jiaxing Song, Weidong Liu, and Dawei Yin. 2019. Reinforcement Learning to Optimize Long-term User Engagement in Recommender Systems. In Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '19). Association for Computing Machinery, New York, NY, USA, 2810-2818. https://doi.org/10.1145/3292500.3330668

[2] Mazoure, B., Mineiro, P., Srinath, P., Sedeh, R. S., Precup, D., & Swaminathan, A. (2021). Improving long-term metrics in recommendation systems using short-horizon reinforcement learning. arXiv preprint arXiv:2106.00589. https://doi.org/10.48550/arXiv.2106.00589

[3] Pei, C., Yang, X., Cui, Q., Lin, X., Sun, F., Jiang, P., ... & Zhang, Y. (2019). Value-aware recommendation based on reinforced profit maximization in e-commerce systems. arXiv preprint arXiv:1902.00851. https://doi.org/10.48550/arXiv.1902.00851

[4] J. Han, Y. Yu, F. Liu, R. Tang and Y. Zhang, "Optimizing Ranking Algorithm in Recommender System via Deep Reinforcement Learning," 2019 International Conference on Artificial Intelligence and Advanced Manufacturing (AIAM), Dublin, Ireland, 2019, pp. 22-26, doi: 10.1109/AIAM48774.2019.00011

[5] X. -H. Chen et al., "Sim2Rec: A Simulator-based Decision-making Approach to Optimize Real-World Long-term User Engagement in Sequential Recommender Systems," 2023 IEEE 39th International Conference on Data Engineering (ICDE), Anaheim, CA, USA, 2023, pp. 3389-3402, doi: 10.1109/ICDE55515.2023.00260.

[6] Wanqi Xue, Qingpeng Cai, Zhenghai Xue, Shuo Sun, Shuchang Liu, Dong Zheng, Peng Jiang, Kun Gai, and Bo An. 2023. PrefRec: Recommender Systems with Human Preferences for Reinforcing Long-term User Engagement. In Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23). Association for Computing Machinery, New York, NY, USA, 2874-2884. https://doi.org/10.1145/3580305.3599473

[7] Xue, Z., Cai, Q., Zuo, T., Yang, B., Hu, L., Jiang, P., ... & An, B. (2023). AdaRec: Adaptive sequential recommendation for reinforcing long-term user engagement. arXiv preprint arXiv:2310.03984. https://doi.org/10.48550/arXiv.2310.03984

[8] Komal Kapoor, Mingxuan Sun, Jaideep Srivastava, and Tao Ye. 2014. A hazard based approach to user return time prediction. In Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '14). Association for Computing Machinery, New York, NY, USA, 1719-1728. https://doi.org/10.1145/2623330.2623348

[9] Nan Du, Yichen Wang, Niao He, and Le Song. 2015. Time-sensitive recommendation from recurrent user activities. In Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2 (NIPS'15). MIT Press, Cambridge, MA, USA, 3492-3500. https://dl.acm.org/doi/10.5555/2969442.2969629

[10] Olivier Chapelle. 2014. Modeling delayed feedback in display advertising. In Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD '14). Association for Computing Machinery, New York, NY, USA, 1097-1105. https://doi.org/10.1145/2623330.2623634

[11] Chatzidimitris, T., Gavalas, D., Kasapakis, V., Konstantopoulos, C., Kypriadis, D., Pantziou, G., & Zaroliagis, C. (2020). A location history-aware recommender system for smart retail environments. Personal and Ubiquitous Computing, 24, 683-694. https://doi.org/10.1007/s00779-020-01374-7

[12] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In Proceedings of the 10th ACM Conference on Recommender Systems (RecSys '16). Association for Computing Machinery, New York, NY, USA, 191-198. https://doi.org/10.1145/2959100.2959190

[13] Mihajlo Grbovic and Haibin Cheng. 2018. Real-time Personalization using Embeddings for Search Ranking at Airbnb. In Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD '18). Association for Computing Machinery, New York, NY, USA, 311-320. https://doi.org/10.1145/3219819.3219885

[14] Wang, Y., Ning, Y., Liu, I., & Zhang, X. X. (2021). Food discovery with Uber Eats: recommending for the marketplace. https://www.uber.com/en-IN/blog/uber-eats-recommending-marketplace

[15] Guanjie Zheng, Fuzheng Zhang, Zihan Zheng, Yang Xiang, Nicholas Jing Yuan, Xing Xie, and Zhenhui Li. 2018. DRN: A Deep Reinforcement Learning Framework for News Recommendation. In Proceedings of the 2018 World Wide Web Conference (WWW '18). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 167-176. https://doi.org/10.1145/3178876.3185994

[16] Jiahui Liu, Peter Dolan, and Elin Rønby Pedersen. 2010. Personalized news recommendation based on click behavior. In Proceedings of the 15th international conference on Intelligent user interfaces (IUI '10). Association for Computing Machinery, New York, NY, USA, 31-40. https://doi.org/10.1145/1719970.1719976

[17] Z. Wang, J. Liao, Q. Cao, H. Qi and Z. Wang, "Friendbook: A Semantic-Based Friend Recommendation System for Social Networks," in IEEE Transactions on Mobile Computing, vol. 14, no. 3, pp. 538-551, 1 March 2015, doi: 10.1109/TMC.2014.2322373.

[18] X. Xie, "Potential Friend Recommendation in Online Social Network," 2010 IEEE/ACM Int'l Conference on Green Computing and Communications & Int'l Conference on Cyber, Physical and Social Computing, Hangzhou, China, 2010, pp. 831-835, doi: 10.1109/GreenCom-CPSCom.2010.28.

[19] Shuai Zhang, Lina Yao, Aixin Sun, and Yi Tay. 2019. Deep Learning Based Recommender System: A Survey and New Perspectives. ACM Comput. Surv. 52, 1, Article 5 (January 2020), 38 pages. https://doi.org/10.1145/3285029

[20] Allison J. B. Chaney, Brandon M. Stewart, and Barbara E. Engelhardt. 2018. How algorithmic confounding in recommendation systems increases homogeneity and decreases utility. In Proceedings of the 12th ACM Conference on Recommender Systems (RecSys '18). Association for Computing Machinery, New York, NY, USA, 224-232. https://doi.org/10.1145/3240323.3240370

[21] Tien T. Nguyen, Pik-Mai Hui, F. Maxwell Harper, Loren Terveen, and Joseph A. Konstan. 2014. Exploring the filter bubble: the effect of using recommender systems on content diversity. In Proceedings of the 23rd international conference on World wide web (WWW '14). Association for Computing Machinery, New York, NY, USA, 677-686. https://doi.org/10.1145/2566486.2568012

[22] Xing Zhao, Ziwei Zhu, and James Caverlee. 2021. Rabbit Holes and Taste Distortion: Distribution-Aware Recommendation with Evolving Interests. In Proceedings of the Web Conference 2021 (WWW '21). Association for Computing Machinery, New York, NY, USA, 888-899. https://doi.org/10.1145/3442381.3450099

[23] Sharad Goel, Andrei Broder, Evgeniy Gabrilovich, and Bo Pang. 2010. Anatomy of the long tail: ordinary people with extraordinary tastes. In Proceedings of the third ACM international conference on Web search and data mining (WSDM '10). Association for Computing Machinery, New York, NY, USA, 201-210. https://doi.org/10.1145/1718487.1718513

[24] Knijnenburg, B. P., Willemsen, M. C., Gantner, Z., Soncu, H., & Newell, C. (2012). Explaining the user experience of recommender systems. User modeling and user-adapted interaction, 22, 441-504. https://doi.org/10.1007/s11257-011-9118-4

[25] Xiao, B., & Benbasat, I. (2007). E-Commerce Product Recommendation Agents: Use, Characteristics, and Impact. MIS Quarterly, 31(1), 137-209. https://doi.org/10.2307/25148784

[26] Ashton Anderson, Lucas Maystre, Ian Anderson, Rishabh Mehrotra, and Mounia Lalmas. 2020. Algorithmic Effects on the Diversity of Consumption on Spotify. In Proceedings of The Web Conference 2020 (WWW '20). Association for Computing Machinery, New York, NY, USA, 2155-2165. https://doi.org/10.1145/3366423.3380281

[27] Quentin Villermet, Jérémie Poiroux, Manuel Moussallam, Thomas Louail, and Camille Roth. 2021. Follow the guides: disentangling human and algorithmic curation in online music consumption. In Proceedings of the 15th ACM Conference on Recommender Systems (RecSys '21). Association for Computing Machinery, New York, NY, USA, 380-389. https://doi.org/10.1145/3460231.3474269

[28] Ansari, A., Li, Y., & Zhang, J. Z. (2018). Probabilistic topic model for hybrid recommender systems: A stochastic variational Bayesian approach. Marketing Science, 37(6), 987-1008. https://doi.org/10.1287/mksc.2018.1113

[29] Francesco Sanna Passino, Lucas Maystre, Dmitrii Moor, Ashton Anderson, and Mounia Lalmas. 2021. Where To Next? A Dynamic Model of User Preferences. In Proceedings of the Web Conference 2021 (WWW '21). Association for Computing Machinery, New York, NY, USA, 3210-3220. https://doi.org/10.1145/3442381.3450028

[30] Karumur, R. P., Nguyen, T. T., & Konstan, J. A. (2018). Personality, user preferences and behavior in recommender systems. Information Systems Frontiers, 20, 1241-1265. https://doi.org/10.1007/s10796-017-9800-0

[31] David Holtz, Ben Carterette, Praveen Chandar, Zahra Nazari, Henriette Cramer, and Sinan Aral. 2020. The Engagement-Diversity Connection: Evidence from a Field Experiment on Spotify. In Proceedings of the 21st ACM Conference on Economics and Computation (EC '20). Association for Computing Machinery, New York, NY, USA, 75-76. https://doi.org/10.1145/3391403.3399532

[32] Zhang, J., Adomavicius, G., Gupta, A., & Ketter, W. (2020). Consumption and performance: Understanding longitudinal dynamics of recommender systems via an agent-based simulation framework. Information Systems Research, 31(1), 76-101. https://doi.org/10.1287/isre.2019.0876

[33] Christian Hansen, Rishabh Mehrotra, Casper Hansen, Brian Brost, Lucas Maystre, and Mounia Lalmas. 2021. Shifting Consumption towards Diverse Content on Music Streaming Platforms. In Proceedings of the 14th ACM International Conference on Web Search and Data Mining (WSDM '21). Association for Computing Machinery, New York, NY, USA, 238-246. https://doi.org/10.1145/3437963.3441775

[34] Renjie Zhou, Samamon Khemmarat, and Lixin Gao. 2010. The impact of YouTube recommendation system on video views. In Proceedings of the 10th ACM SIGCOMM conference on Internet measurement (IMC '10). Association for Computing Machinery, New York, NY, USA, 404-410. https://doi.org/10.1145/1879141.1879193

[35] Uta M, Felfernig A, Le VM, Tran TNT, Garber D, Lubos S, Burgstaller T. Knowledge-based recommender systems: overview and research directions. Front Big Data. 2024 Feb 26;7:1304439. doi: 10.3389/fdata.2024.1304439. https://doi.org/10.3389/fdata.2024.1304439

[36] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: a survey of the state-of-the-art and possible extensions," in IEEE Transactions on Knowledge and Data Engineering, vol. 17, no. 6, pp. 734-749, June 2005, doi: 10.1109/TKDE.2005.99.

[37] Luo Ji, Qi Qin, Bingqing Han, and Hongxia Yang. 2021. Reinforcement Learning to Optimize Lifetime Value in Cold-Start Recommendation. In Proceedings of the 30th ACM International Conference on Information & Knowledge Management (CIKM '21). Association for Computing Machinery, New York, NY, USA, 782-791. https://doi.org/10.1145/3459637.3482292

[38] Qihua Zhang, Junning Liu, Yuzhuo Dai, Yiyan Qi, Yifan Yuan, Kunlun Zheng, Fan Huang, and Xianfeng Tan. 2022. Multi-Task Fusion via Reinforcement Learning for Long-Term User Satisfaction in Recommender Systems. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22). Association for Computing Machinery, New York, NY, USA, 4510-4520. https://doi.org/10.1145/3534678.3539040

[39] Sun, K., Qian, T., Chen, T., Liang, Y., Nguyen, Q. V. H., & Yin, H. (2020). Where to Go Next: Modeling Long- and Short-Term User Preferences for Point-of-Interest Recommendation. Proceedings of the AAAI Conference on Artificial Intelligence, 34(01), 214-221. https://doi.org/10.1609/aaai.v34i01.5353

[40] Qingyun Wu, Hongning Wang, Liangjie Hong, and Yue Shi. 2017. Returning is Believing: Optimizing Long-term User Engagement in Recommender Systems. In Proceedings of the 2017 ACM on Conference on Information and Knowledge Management (CIKM '17). Association for Computing Machinery, New York, NY, USA, 1927-1936. https://doi.org/10.1145/3132847.3133025

[41] Yuyan Wang, Mohit Sharma, Can Xu, Sriraj Badam, Qian Sun, Lee Richardson, Lisa Chung, Ed H. Chi, and Minmin Chen. 2022. Surrogate for Long-Term User Experience in Recommender Systems. In Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '22). Association for Computing Machinery, New York, NY, USA, 4100-4109. https://doi.org/10.1145/3534678.3539073

[42] Saig, E., & Rosenfeld, N. (2023, July). Learning to suggest breaks: sustainable optimization of long-term user engagement. In International Conference on Machine Learning (pp. 29671-29696). PMLR.

[43] Zahrawi, Mohammad. (2021). Implementing Recommender Systems using Machine Learning and Knowledge Discovery Tools. Knowledge-Based Engineering and Sciences. 2. 44-53. 10.51526/kbes.2021.2.2.44-53. https://doi.org/10.51526/kbes.2021.2.2.44-53

[44] Castells, P., Hurley, N., Vargas, S. (2022). Novelty and Diversity in Recommender Systems. In: Ricci, F., Rokach, L., Shapira, B. (eds) Recommender Systems Handbook. Springer, New York, NY. https://doi.org/10.1007/978-1-0716-2197-4_16

[45] Pearl Pu, Li Chen, and Rong Hu. 2011. A user-centric evaluation framework for recommender systems. In Proceedings of the fifth ACM conference on Recommender systems (RecSys '11). Association for Computing Machinery, New York, NY, USA, 157-164. https://doi.org/10.1145/2043932.2043962

[46] Ricci, F., Rokach, L., Shapira, B. (2011). Introduction to Recommender Systems Handbook. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P. (eds) Recommender Systems Handbook. Springer, Boston, MA. https://doi.org/10.1007/978-0-387-85820-3_1

[47] Burke, R. Hybrid Recommender Systems: Survey and Experiments. User Model User-Adap Inter 12, 331-370 (2002). https://doi.org/10.1023/A:1021240730564

[48] Adomavicius, G., Tuzhilin, A. (2011). Context-Aware Recommender Systems. In: Ricci, F., Rokach, L., Shapira, B., Kantor, P. (eds) Recommender Systems Handbook. Springer, Boston, MA. https://doi.org/10.1007/978-0-387-85820-3_7

[49] Panteli A, Boutsinas B. Addressing the Cold-Start Problem in Recommender Systems Based on Frequent Patterns. Algorithms. 2023; 16(4):182. https://doi.org/10.3390/a16040182

[50] Fayyaz Z, Ebrahimian M, Nawara D, Ibrahim A, Kashef R. Recommendation Systems: Algorithms, Challenges, Metrics, and

Business Opportunities. Applied Sciences. 2020; 10(21):7748. https://doi.org/10.3390/app10217748

[51] Feng, S.; Li, X.; Zeng, Y.; Cong, G.; Chee, Y. M.; and Yuan, Q.2015. Personalized ranking metric embedding for next new poi recommendation. In IJCAI, 2069–2075.

[52] Feng, J.; Li, Y.; Zhang, C.; Sun, F.; Meng, F.; Guo, A.; and Jin, D. 2018. Deepmove: Predicting human mobility with attentional recurrent networks. In WWW

[53] Zhang, S.; Yin, H.; Wang, Q.; Chen, T.; Chen, H.; and Nguyen, Q.V. H. 2019. Inferring substitutable products with deep network embedding. IJCAI 4306–4312.

[54] Shi, C.; Hu, B.; Zhao, W. X.; and Philip, S. Y. 2018. Heterogeneous information network embedding for recommendation. TKDE 357–370.

[55] He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural collaborative filtering. In WWW, 173–182.

[56] Cheng, W.; Shen, Y.; Zhu, Y.; and Huang, L. 2018. Delf: A dualembedding based deep latent factor model for recommendation. In IJCAI, 3329–3335.

[57] Tay, Y.; Anh Tuan, L.; and Hui, S. C. 2018. Latent relational metric learning via memory-based attention for collaborative ranking. In WWW, 729–739.

[58] Zhao, P.; Zhu, H.; Liu, Y.; Xu, J.; Li, Z.; Zhuang, F.; Sheng, V. S.; and Zhou, X. 2019. Where to go next: A spatio-temporal gated network for next poi recommendation. In AAAI, 5877–5884.

[59] Li, R.; Shen, Y.; and Zhu, Y. 2018. Next point-of-interest recommendation with temporal and multi-level context attention. in ICDM, 1110–1115.

[60] Manotumruksa, J.; Macdonald, C.; and Ounis., I. 2017. A deep recurrent collaborative filtering framework for venue recommendation. In CIKM, 1429–1438.

[61] Liu, Q.; Wu, S.; Wang, L.; and Tan, T. 2016. Predicting the next location: a recurrent model with spatial and temporal contexts. In AAAI, 194–200.

[62] Hidasi, B.; Karatzoglou, A.; Baltrunas, L.; and Tikk, D. 2015. Session-based recommendations with recurrent neural networks. arXiv preprint arXiv:1511.06939.