International Scientific Journal of Engineering and Management (ISJEM)
ISSN: 2583-6129
Volume: 04 Issue: 03 | March – 2025
DOI: 10.55041/ISJEM02367
An International Scholarly || Multidisciplinary || Open Access || Indexing in all major Database & Metadata

Resolving the Cold-Start Issue in Recommender Systems with Reinforcement Learning

Anunay Singh[1] , Anveet Pal[2] , Ashish Baghel[3]

[1]Department of Computer Science & Engineering, Medi-Caps University, Indore (M.P.) [2]Department of Computer Science & Engineering, Medi-Caps University, Indore (M.P.) [3]Department of Computer Science & Engineering, Medi-Caps University, Indore (M.P.)

Abstract - The cold-start problem faced by recommender systems is a serious problem, mainly because of the lack of historical data for new users or items. Traditional recommendation techniques, such as collaborative filtering and content-based filtering, are prone to fail in making good recommendations under such conditions. This paper explores the use of reinforcement learning (RL) as a remedy for cold-start problems based on active learning methods and multi-armed bandit models. We propose a novel RL-based approach that learns user preferences incrementally from interaction and improves recommendations in an exploration-exploitation setting. The setting shows improved performance in personalization and user engagement compared to baseline methods. This paper also provides the first-order implications of the cold-start problem in the real world and assesses the challenges faced by industries impacted by this problem.

Keywords - Recommender systems, cold-start problem, reinforcement learning, Markov decision processes, deep reinforcement learning, user modelling

1. Overview

Recommender systems are an essential building block of many industries, such as e-commerce, streaming, and online courses. Lack of interaction history for new items or users, however, leads to low-quality recommendations. This problem, also referred to as the cold-start problem, can be classified into three classes: user cold-start, item cold-start, and system cold-start. Although explicit feedback mechanisms and hybrid methods have been explored, reinforcement learning is a promising choice by dynamically adjusting to user interactions in real time. Solving the cold-start problem can significantly contribute to user experience, engagement, and revenue across many applications.

2. Problem Statement

The most fundamental challenge that is inherent in the cold-start problem is the lack of interaction history; thus, it is challenging to generate personalized recommendations. Past user-item interaction history is used by conventional recommendation systems, which is a data set that is zero for newly added items or users. The lack of information results in inaccurate recommendations, hence low user engagement and high churn rates. The most important research question that this paper seeks to explore is: How can reinforcement learning be employed to effectively alleviate the cold-start problem while inducing high recommendation accuracy rates?

3. Real-Life Situations Impacted by Cold-Start

The cold-starting issue impacts varied regions, including:

Shopping websites online have problems in giving proper product recommendations to new customers, which decrease the conversion rates.

• Streaming Services: Services like Netflix and Spotify face challenges in suggesting movies or songs to new users who have an insufficient viewing history.

• Online Learning: New learners in learning management systems have a hard time receiving focused course recommendations.

• Job Portals: Fresh candidates entering the workforce are provided poor job recommendations by virtue of non-existence of past application histories.

• Health: Patient history is needed for tailored health counsel and is difficult for new patients.

4. Challenges Faced in Cold-Start Mitigation

Numerous significant challenges occur in the solution of the cold-start problem:

• Sparsity of Data: Historical data scarcity makes it hard to train the models. Mostly, recommender systems rely on user history; however, new users or items might have zero or sparse historical interactions, leading to poor recommendations.

• User Engagement: Failing to offer personalized recommendations can lessen user engagement. Users who don't receive relevant recommendations at the right time may show lack in interest in using the platform, resulting in increased churn rates.

• Bias in Recommendations: Excessive dependence on standardized recommendations might result in biased recommendations. Most systems resolve cold-start problems by suggesting extremely popular items; however, this reduces diversity and distinctiveness, ultimately resulting in a monotonous user experience.

• Computational Complexity: RL models are computationally intensive. In contrast to other recommendation models, RL-based models consume a lot of computational resources because they need to continuously train and learn.

• Scalability: Large-scale recommendation systems must be able to process new items and users in an efficient manner. As the user base and product grows, the recommendation algorithms must scale effectively and accurately.

5. Relevant Methodologies

Present cold-start strategies mainly comprise:

• Explicit Feedback Collection: The method of gathering early user preferences via onboarding surveys. While the method can help alleviate cold-start problems, it is largely based on the users providing correct information, which is not always the case.

• Hybrid Models: Using a mix of collaborative and content-based filtering. Hybrid models use both interaction history and item metadata to make recommendations, but they also degrade when user data is sparse.

• Active Learning and Multi-Armed Bandits: Emphasizing discovering user preferences by deliberate questioning. Multi-Armed Bandit models allow for enhancing recommendation quality by balancing exploration (acquiring information about new users) and exploitation (using existing data to enhance recommendations).

Recent studies have shown that reinforcement learning (RL)-based methods, such as Deep Q Networks (DQN) and Thompson Sampling, are capable of improving recommendation quality well by balancing exploration and exploitation methods. For example, Mnih et al. (2015) introduced a deep reinforcement learning model that performed more effectively in game settings, and similar approaches have been applied to recommender systems to enable dynamic learning from feedback. In addition, studies have suggested that reinforcement learning can enhance cold-start recommendations through the use of pre-trained models and rapid adaptation to new information.

6. Proposed Methodology

Our approach presents an RL-based recommendation model that solves the cold-start problem with the following elements:

• User Clustering: Segmentation of users based on shared demographic and contextual factors to support initial recommendations.

• Exploration-Exploitation Mechanism: Using an Upper Confidence Bound (UCB) model for dynamically updating recommendations.

• Deep Q-Networks (DQN): Applying deep reinforcement learning methods to improve recommendations by learning from implicit feedback.

• User Interaction-Based Adaptation: Inbuilt profiling of users through clicks, dwell time, and user signals of engagement.

7. Experimental Setup

•Datasets: The test is performed on real-world publicly available datasets.

• Performance is measured in terms of Precision@K, Recall@K, and Normalized Discounted Cumulative Gain (NDCG).

• Baseline Models: Comparison to collaborative filtering, content-based filtering, and hybrid models.

8. Findings and Analysis

Experimental results show that reinforcement learning-based models outperform traditional models in handling cold-start users. The exploration-exploitation trade-off ensures that users receive progressively more relevant recommendations over time, leading to increased engagement and retention. The clustering process significantly reduces computational overhead, thereby ensuring the approach is scalable for large data.

9. Conclusion and Future Research Trends

This study emphasizes the potential of reinforcement learning to solve the cold-start problem. The future work will be focused on incorporating federated learning to offer privacy-aware recommendations and meta-learning approaches to speed up new user adaptation . The use of graph neural networks (GNNs) can further enhance contextual awareness of user-item interactions. Future studies can investigate reinforcement learning techniques applied in multi-modal recommendation systems that support text, visual, and audio modes to enable personalized experience. Additionally, investigating the use of reinforcement learning to reinforce causal inference techniques can help develop more robust and interpretable recommendation systems. Lastly, investigating real-time adaptation of reinforcement learning-based recommender systems based on streaming data can further enhance their responsiveness and accuracy in dynamic settings.

References

[1] F. Ricci, L. Rokach, and B. Shapira, *Recommender Systems Handbook*, 2nd ed. Springer, 2015.
[2] A. I. Schein, A. Popescul, L. H. Ungar, and D. M. Pennock, "Methods and metrics for cold-start recommendations," in *SIGIR Proceedings*, 2002, pp. 253–260.
[3] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. MIT Press, 2018.
[4] X. He, L. Liao, H. Zhang, and T. Chua, "Neural collaborative filtering," in *WWW Proceedings*, 2017, pp. 173–182.
[5] V. Mnih, K. Kavukcuoglu, D. Silver, et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
[6] Y. Koren, R. Bell, and C. Volinsky, "Matrix factorization techniques for recommender

systems," *Computer*, vol. 42, no. 8, pp. 30–37, 2009.

[7] A. Feder and Y. Shokef, "Federated learning: Privacy-preserving personalized recommendations," *IEEE Transactions on Big Data*, vol. 6, no. 1, pp. 45–56, 2020.

[8] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine Learning*, vol. 47, no. 2-3, pp. 235–256, 2002.

[9] S. Rendle, "Factorization machines," in *IEEE ICDM Proceedings*, 2010, pp. 995–1000.

[10] F. M. Harper and J. A. Konstan, "The MovieLens datasets: History and context," *ACM Transactions on Interactive Intelligent Systems (TIIS)*, vol. 5, no. 4, pp. 1–19, 2015.

[11] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl, "Evaluating collaborative filtering recommender systems," *ACM Transactions on Information Systems (TOIS)*, vol. 22, no. 1, pp. 5–53, 2004.

[12] T. Lattimore and C. Szepesvári, *Bandit Algorithms*, Cambridge University Press, 2020.

[13] X. He, H. Zhang, and T. Chua, "Neural collaborative filtering with autoencoder," in *WWW Proceedings*, 2018, pp. 1573–1582.

[14] Y. Wang, X. He, L. Nie, and T. Chua, "Trust-aware collaborative deep learning for recommendation systems," in *ACM SIGIR Proceedings*, 2019, pp. 1303–1312.

[15] T. Hofmann, "Probabilistic latent semantic analysis," in *UAI Proceedings*, 1999, pp. 289–296.

[16] D. Blei, A. Ng, and M. Jordan, "Latent Dirichlet allocation," *Journal of Machine Learning Research*, vol. 3, pp. 993–1022, 2003.

[17] J. McAuley, C. Targett, Q. Shi, and A. van den Hengel, "Image-based recommendations on styles and substitutes," in *ACM SIGIR Proceedings*, 2015, pp. 43–52.

[18] M. Quadrana, A. Karatzoglou, and M. Cremonesi, "A survey on sequence-aware recommender systems," *ACM Computing Surveys (CSUR)*, vol. 51, no. 4, pp. 1–36, 2018.

[19] R. He and J. McAuley, "VBPR: Visual Bayesian personalized ranking from implicit feedback," in *AAAI Conference on Artificial Intelligence*, 2016, pp. 144–150.

[20] J. Tang, X. Hu, and H. Liu, "Social recommendation: A review," *Social Network Analysis and Mining*, vol. 3, no. 4, pp. 1113–1133, 2013.

[21] M. D. Ekstrand, J. T. Riedl, and J. A. Konstan, "Collaborative filtering recommender systems," *Foundations and Trends in Human-Computer Interaction*, vol. 4, no. 2, pp. 81–173, 2010.

[22] S. Liu, X. He, and H. Zhang, "Personalized recommendation via integrating user historical preferences with deep learning," *IEEE Transactions on Knowledge and Data Engineering*, vol. 33, no. 1, pp. 232–245, 2021.

[23] B. M. Kim, Q. Wang, and H. Park, "Deep reinforcement learning for large-scale personalized recommendation systems," *Neural Computing and Applications*, vol. 34, no. 8, pp. 5725–5740, 2022.

[24] Z. Zhao, Y. Xin, and J. Tang, "Graph neural networks for recommendation systems: A survey," *IEEE Transactions on Knowledge and Data Engineering*, vol. 35, no. 4, pp. 2345–2360, 2023.