Reinforcement Learning for Personalized News Recommendation: A Deep Q-Network Approach

ADITYA CHAUHAN | ARJUN VERMA | MD AMIR KAMAL | RAJ ARYAN SINGH

Mentor : Shobha Rai

ABSTRACT

In the rapidly evolving digital landscape, where the sheer volume of news content can be overwhelming, personalized recommendation systems have emerged as a critical tool to enhance user experience and engagement. This project presents a novel approach to news recommendation by leveraging the power of Reinforcement Learning (RL), specifically Deep Q-Networks (DQN). Trained on the extensive Microsoft News Dataset (MIND), our system addresses the challenges of extracting meaningful features from news texts and categories, formulating the problem within the RL framework, and enabling offline training.

The DQN agent learns to recommend relevant news categories based on the user's browsing history, dynamically adapting to user feedback through an online learning process. By incorporating user interactions and preferences, the system continuously refines its recommendations, providing a personalized and engaging news consumption experience. This project demonstrates the potential of RL techniques in tackling complex recommendation tasks and paves the way for further advancements in intelligent content curation.

BACKGROUND

Recommendation systems have been a vital component of various online platforms, aiming to provide personalized content and enhance user experience. Traditional approaches, such as collaborative filtering and content-based filtering, rely on historical user-item interactions and content similarities, respectively. However, these methods often struggle to capture the dynamic nature of user preferences and may fail to adapt to evolving interests effectively.

Reinforcement Learning (RL) has emerged as a promising technique for recommendation tasks, as it allows systems to learn from user feedback and adapt dynamically. RL models the recommendation process as a sequential decision-making problem, where an agent (the recommender system) interacts with an environment (the user and their preferences) by taking actions (recommending news articles) and receiving rewards (user feedback). By optimizing for long-term rewards, RL algorithms can learn to make recommendations that maximize user satisfaction over time.

Our project leverages the powerful capabilities of DQN to build an intelligent News Recommender System tailored to individual user preferences. By harnessing the rich information present in the Microsoft News Dataset (MIND), our system tackles the challenges of feature engineering, problem formulation, and offline training, paving the way for a personalized and engaging news consumption experience.

ARCHITECTURE

Our News Recommender System comprises three key components: the DQN Agent, Environment, and Feature Engineering Pipeline. The DQN Agent, leveraging reinforcement learning, utilizes deep neural networks to map user browsing history to relevant news categories. Through trial-and-error interactions, it refines recommendations, aided by techniques like experience replay and target network updates to stabilize learning and balance exploration-exploitation trade-offs.

The Environment, built around the Microsoft News Dataset (MIND), simulates user interactions, maintaining user states and responding to agent recommendations. By incorporating personalized ranking, diversity, and randomization, it ensures engaging content presentation, providing a rich user experience.

The Feature Engineering Pipeline transforms raw data into actionable features for the DQN Agent. It employs text preprocessing, tokenization, and advanced techniques like TF-IDF and word embeddings to extract semantic and contextual information from news articles. Moreover, it encodes user browsing histories and news categories, ensuring compatibility with the agent's neural network architecture.

By seamlessly integrating these components, our system offers personalized news recommendations that adapt to individual user preferences. The DQN Agent learns iteratively from user feedback, while the Environment simulates real-world interactions, and the Feature Engineering Pipeline enhances data representation. This holistic approach enables the system to continually improve recommendation quality, providing users with relevant and engaging content tailored to their interests.

MIND DATASET

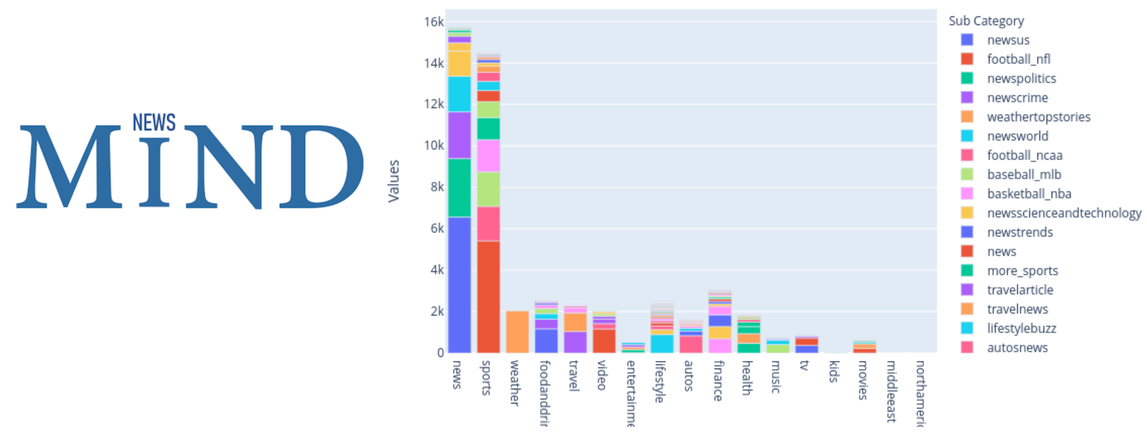

METHODOLOGY

Data Preprocessing: Data preprocessing is integral to our News Recommender System's reliability and performance. We start by handling missing values and removing irrelevant or redundant features from the Microsoft News Dataset (MIND). This ensures high-quality data, reducing noise and improving recommendation accuracy. Text preprocessing follows, including tokenization, stop-word removal, and stemming/lemmatization to reduce text dimensionality and capture essential semantic information. TF-IDF and word embeddings further transform news articles into numerical formats suitable for the DQN agent's neural network, capturing semantic and contextual information. Encoding user browsing histories, news categories, and other features ensures compatibility with the agent's architecture.

Model Architecture and Training: Our News Recommender System centers on the Deep Q-Network (DQN) agent, which intelligently recommends news based on user browsing history and preferences. The DQN agent's deep neural network approximates the action-value function, estimating long-term rewards for recommending news categories in specific user states. Training utilizes experience replay and target network updates to stabilize learning and improve data efficiency. The agent explores actions, receiving rewards based on user interactions to refine recommendations over time. Techniques like epsilon-greedy exploration balance exploitation and exploration for effective adaptation to user preferences. Hyperparameters like learning rate and batch size are optimized to ensure convergence and performance.

In summary, our system preprocesses data to ensure quality and transforms it into formats compatible with the DQN agent's architecture. The agent's neural network learns from experiences, refining recommendations based on user feedback. By balancing exploration and exploitation, our system continuously adapts to user preferences, delivering personalized news recommendations with high accuracy and reliability.

RESULT

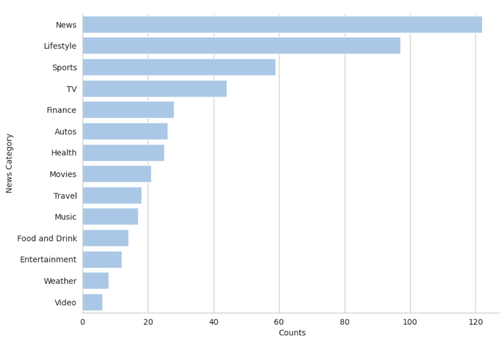

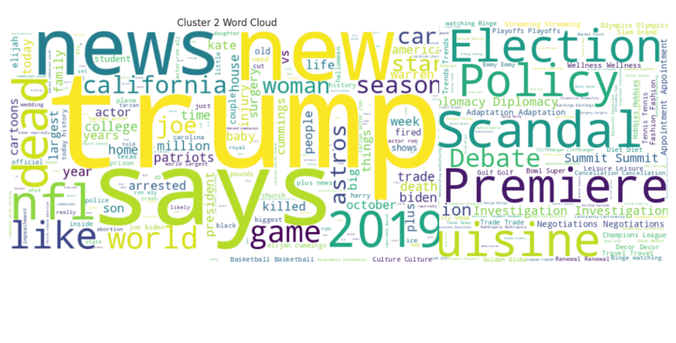

Topics recommended to two users along with their corresponding word clouds.

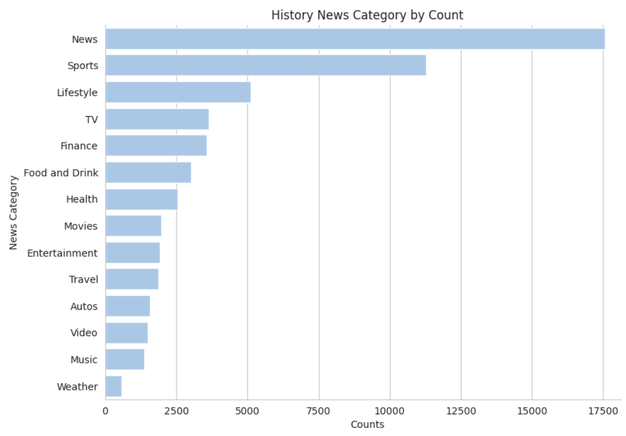

Backend for our project

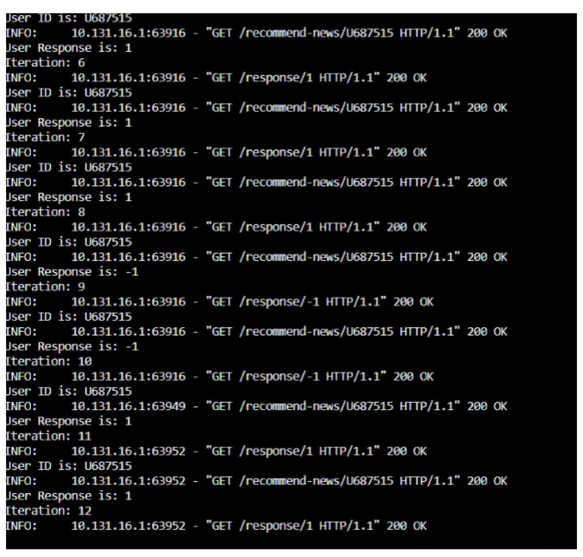

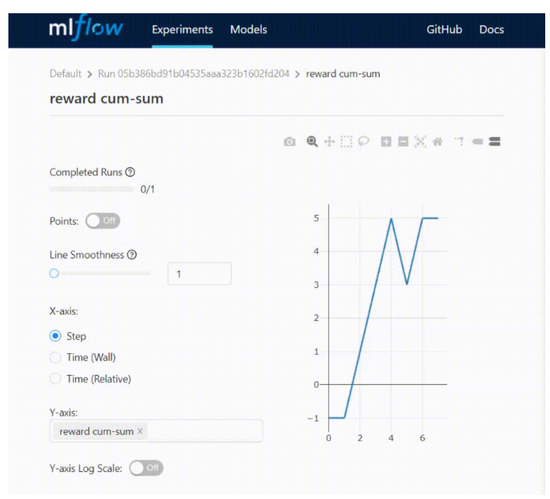

ML Flow metrics

CONCLUSION

While the current implementation has yielded promising results, the potential for further improvements remains vast. Ultimately, this project exemplifies the power of combining advanced machine learning techniques with rich datasets to develop intelligent systems that cater to individual user preferences. By bridging the gap between the vast information landscape and personalized content curation, our News Recommender System paves the way for more engaging and satisfying news consumption experiences, empowering users to navigate the digital world with ease and efficiency.

FUTURE WORK

Future Work and Potential Improvements: While the current News Recommender System has shown promising results, there are areas for improvement:

Exploring Different Models: Beyond the current DQN model, exploring Transformer, BERT, or GPT models could enhance news article and user history representation, potentially improving recommendation accuracy.

Incorporating User Feedback: Implementing a feedback loop for user interactions, whether through explicit ratings or implicit feedback like click events, can personalize recommendations further.

Handling Cold-Start Problems: Addressing the cold-start problem for new users or news articles can improve recommendation quality, especially during initial recommendations before sufficient data is available.

References:

[1] 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.CHE, 167–176. https://doi.org/10.1145/3178876.3185994

[2] Edirisinghe, Kelum. (2020). Reinforcement Learning For Personalized Recommendations. 10.13140/RG.2.2.32213.45286.

[3]Li, Ziyan. (2024). Application And Optimization of Deep Reinforcement Learning in News Recommendation. Highlights in Science, Engineering and Technology. 85. 389-395. 10.54097/n8t6ty60.

[4]Bangari, S., Nayak, S., Patel, L., & Rashmi, K. T. (2021, March). A review on reinforcement learning based news recommendation systems and its challenges. In 2021 International Conference on Artificial Intelligence and Smart Systems (ICAIS) (pp. 260-265). IEEE.