UDC 004.85

11. ARTIFICIAL INTELLIGENCE AS A MODERN TOOL FOR ENHANCING RECOMMENDER SYSTEMS

Klimenko I.V., Master's degree student, group 476701

Belarusian State University of Informatics and Radioelectronics, Minsk, Republic of Belarus

Liakh Y.V. – Senior Lecturer

Annotation. This article explores modern recommender systems built upon Artificial Intelligence, Machine Learning, and Deep Learning technologies. The author discusses how these methods are employed by digital platforms to enhance user experiences and improve business outcomes. The pivotal role of personalisation, driven by extensive user behaviour analysis, is underscored, and techniques to process large, diverse datasets are examined. Special attention is given to how Deep Learning approaches and Reinforcement Learning enable the dynamic adaptation of recommendations. Additionally, the article addresses ethical considerations such as the formation of «filter bubbles» and the need for algorithmic transparency. Conclusively, it highlights that AI-powered solutions have become integral to recommender systems, offering competitive advantages and meeting user needs in an era of information overload.

Keywords. Recommender systems, machine learning, deep learning, personalisation, user behavior analysis, e-commerce, streaming services, social media, reinforcement learning.

Users of digital platforms increasingly face information overload, struggling to quickly select relevant content among vast options offered by streaming services, e-commerce platforms, and social media. Recommender systems effectively address this issue by suggesting personalised content based on user preferences and behaviours [1].

Initially, recommenders employed simple collaborative filtering methods based on user ratings, suggesting items liked by similar individuals. However, growing complexity of user data demanded more sophisticated solutions, making artificial intelligence (AI), particularly machine learning (ML) and deep learning, essential components of modern recommender systems [2].

Netflix and Amazon exemplify the success of AI-powered recommender systems in driving user engagement and business outcomes. Approximately 80% of Netflix content viewed stems from ML-driven recommendations, profoundly influencing consumption patterns, while Amazon attributes about 35% of its sales to its recommendation engine [1–2]. These figures highlight the substantial commercial and user-experience benefits of advanced recommendation technologies in competitive digital markets.

Modern recommenders are built on machine learning techniques that improve accuracy by analysing historical interactions – such as clicks, views, purchases, and ratings – to predict future preferences. A key ML method is matrix factorisation, popularised during the Netflix Prize competition. This approach decomposes large, sparse user-item interaction matrices into lower-dimensional latent features for users and items. These latent factors reveal hidden characteristics, such as a user's affinity for particular genres or an item's mood, allowing the system to predict unseen ratings or interactions effectively. Netflix's use of matrix factorisation markedly improved its recommendation accuracy, setting industry benchmarks and demonstrating ML's ability to uncover subtle relationships within expansive datasets.

However, as digital platforms evolved, the nature of user data became increasingly complex and diverse, encompassing not just ratings but also implicit signals, sequential actions, and multimodal information (text, images, audio). Traditional ML approaches like matrix factorisation, while powerful, faced limitations in capturing the intricate, non-linear relationships present in such rich data. This necessity paved the way for the adoption of deep learning techniques. Deep learning employs artificial neural networks with multiple layers, which are inherently capable of learning complex hierarchical patterns and representations directly from raw data without extensive manual feature engineering. These networks can process vast amounts of information and model highly non-linear dependencies, making them particularly well-suited for the nuances of human behaviour reflected in digital interactions. The ability of deep learning models to handle various data types simultaneously – tabular, sequential, visual, textual – further expanded the possibilities for creating more holistic and accurate recommender systems.

Platforms like YouTube and Netflix extensively integrate deep neural networks within their recommendation systems, leveraging their power to model complex user behaviour. For instance, Recurrent Neural Networks (RNNs), designed specifically to handle sequential data, are utilised by YouTube to analyse user history and temporal patterns in video consumption. By understanding the order in which users watch videos, RNNs can effectively predict subsequent interests and capture evolving preferences over time. Simultaneously, Convolutional Neural Networks (CNNs), which excel at processing grid-like data such as images, significantly enhance recommendations by analysing visual information. Elements like video thumbnails, movie posters, or product images contain rich signals that influence user choice; CNNs can extract visual features from these images to identify aesthetic preferences and recommend visually similar or appealing content [1].

Netflix provides a compelling example of combining multiple deep learning approaches for sophisticated personalisation. They employ CNNs for visual content analysis, understanding the appeal of different types of imagery, and integrate RNN-based methods to capture temporal viewing patterns – recognising, for example, that viewing habits might differ significantly between weekdays and weekends, or even change based on the time of day. This allows for highly context-sensitive recommendations that adapt not just to what a user likes, but when and how they tend to consume content. The computational demands of training such complex models on massive datasets are substantial, but the resulting improvements in recommendation quality and user satisfaction justify the investment for major platforms.

Furthermore, recommender systems are increasingly incorporating reinforcement learning (RL), a paradigm distinct from supervised and unsupervised learning, which significantly enhances their adaptive capability. RL frames the recommendation problem as a sequence of decisions over time, where the system learns to optimise its actions (recommendations) based on feedback received from the environment (user interactions like clicks, views, or purchases). This allows systems to dynamically adapt their strategies by continuously analysing real-time user responses, learning optimal policies to maximise long-term user engagement or satisfaction. A prevalent implementation in this domain is the multi-armed bandit algorithm, a class of RL techniques that strategically balance the exploration of novel or less-known content with the exploitation of known user preferences. Platforms such as Netflix utilise RL techniques, including bandit algorithms, to periodically introduce unfamiliar content into recommendations. This carefully balances the need to expose users to novelty and facilitate discovery against the safety of recommending familiar favourites, thereby maintaining long-term user engagement, satisfaction, and preventing recommendations from becoming overly monotonous [1]. The ability of RL to learn from interaction rewards makes it particularly suited for environments where user preferences are dynamic and explicit feedback is scarce.

Ultimately, machine learning, deep learning, and reinforcement learning collectively elevate recommender systems far beyond their initial capabilities. They work synergistically to ensure highly personalised, accurate, and context-aware content recommendations. These advanced AI-driven methods are crucial not only for enhancing the user experience but also for maintaining a competitive advantage and driving user satisfaction in today's information-rich digital landscape. They fundamentally transform traditional recommenders into dynamic, highly personalised systems capable of effectively handling massive datasets, adapting flexibly to changing user behaviours, and significantly enhancing prediction accuracy and overall user engagement.

Personalisation is central to the effectiveness of recommender systems, allowing digital platforms to offer content uniquely tailored to each user. At its core, personalisation involves building detailed user profiles that reflect individual preferences and behaviours. These profiles are often represented mathematically as vectors or embeddings in a high-dimensional space, capturing subtle patterns derived from a user's interactions with content or products across the platform. The sophistication of these profiles directly impacts the quality of the recommendations generated.

Creating accurate user profiles depends heavily on analysing extensive user behaviour data, often on a massive scale. Examples include explicit signals like ratings and reviews, but more commonly rely on implicit feedback such as clicks, viewing history, purchase records, likes, shares, scrolling behaviour, dwell time, and even mouse movements. Each interaction provides valuable information for recommender systems, helping AI algorithms detect nuanced preferences and behavioural trends. For example, a streaming platform might recognise that a user frequently watches documentaries related to science on weekday evenings, while preferring light-hearted comedy series during weekend afternoons. These subtle patterns allow for fine-grained personalisation beyond simple genre preferences.

Implicit feedback plays a particularly critical role in user behaviour analysis, especially given that explicit ratings are often sparse. Unlike explicit feedback, implicit signals are derived indirectly from user actions (or inactions). Actions like skipping content quickly, repeatedly returning to certain pages or product categories, or spending considerable time viewing specific items provide powerful clues about underlying user preferences and intent. AI models are adept at leveraging these often noisy but abundant implicit signals, significantly enhancing recommendation accuracy even when explicit ratings are largely absent. Interpreting this implicit data correctly is a key challenge and area of ongoing research.

To maintain recommendation quality and accuracy in the face of evolving user tastes and content landscapes, platforms frequently employ rigorous evaluation methods, most notably A/B testing. This method involves comparing two or more different recommendation algorithms or strategies simultaneously by deploying them to distinct segments of the user base. Through these controlled experiments, platforms like Netflix and Amazon continuously monitor and evaluate performance metrics such as user engagement (e.g., click-through rates, viewing time), conversion rates (e.g., purchases, subscriptions), and user satisfaction surveys. Regular evaluation through A/B testing is crucial for iteratively improving algorithms and adapting recommendation systems effectively to evolving user preferences and contexts [3, 5]. This continuous refinement cycle is essential for staying relevant.

Personalisation, driven by continuous behaviour analysis and rigorous testing, ensures recommender systems maintain high relevance and adaptivity, delivering engaging, tailored experiences in real-time. Streaming services like Spotify and YouTube extensively use AI-driven systems to enhance user experiences by navigating vast content libraries and offering personalised suggestions.

Spotify heavily relies on AI, notably in "Discover Weekly" and "Daily Mix", combining collaborative filtering (analysing similar users' habits) with content-based methods using deep learning to analyse raw audio characteristics (tempo, energy, genre). Similarly, YouTube employs deep neural networks to recommend personalised video content at scale, analysing viewing histories, interactions, search queries, and audio-visual features in real-time, continuously adapting recommendations to maximise engagement and time on the platform. Consequently, streaming platforms effectively manage large content catalogues, maintaining high user satisfaction through deeply personalised experiences.

In e-commerce, AI-driven recommender systems are crucial for personalised shopping experiences impacting sales and loyalty. Amazon attributes significant retail success (around 35% of sales) to its sophisticated recommendation engine, using advanced collaborative filtering and ML techniques to analyse Browse history, past purchases, viewed items, and cart activity to effectively recommend relevant products throughout the user journey, such as in “Customers who bought this item also bought” and “Inspired by your Browse history” sections, and as tailored suggestions on homepages and in emails [2].

AI in e-commerce extends beyond interaction data to include image and text analysis. Deep learning techniques like CNNs analyse product images to identify visual similarities, aiding users in discovering aesthetically similar products even without specific search terms, particularly useful in fashion and home décor. Additionally, NLP techniques analyse customer reviews and product descriptions, extracting key features, attributes, and user sentiments. This textual understanding refines recommendations by matching products based on detailed characteristics and perceived quality, significantly boosting user engagement by simplifying discovery, reducing bounce rates, and increasing sales conversions. AI-driven recommenders transform generic online shopping into highly personalised journeys, consistently aligning suggestions with individual preferences and needs.

Social media platforms like TikTok, Facebook, and Instagram fundamentally depend on AI-powered recommender systems to curate and personalise content streams. These systems analyse vast amounts of user-generated content (posts, photos, videos) and real-time interaction data (likes, shares, comments, viewing time) at scale. Based on this, they deliver personalised posts, videos, friend suggestions, and targeted ads tailored to individual interests and predicted engagement, aiming to maximise user time on the platform by continuously presenting compelling and relevant content.

While AI demonstrably improves recommender systems' accuracy and personalisation, driving key metrics for platforms like Netflix, Amazon, and TikTok [1-2], social media applications raise ethical concerns, notably the risk of "filter bubbles" or "echo chambers". By consistently showing users content aligned with past behaviour and inferred beliefs, algorithms might limit exposure to diverse perspectives, potentially reinforcing biases and hindering exposure to different viewpoints. These ethical considerations have prompted calls for greater transparency regarding recommendation algorithm operation. Companies and researchers are exploring ways to balance personalisation with content diversity, serendipity, and fairness, ensuring systems do not unintentionally amplify societal biases, spread misinformation, or contribute to polarisation. Responsibly managing the influence of these systems on user content consumption remains critical, with researchers also emphasising incorporating transparency mechanisms, enabling users to understand and potentially influence recommendations.

In conclusion, AI has profoundly reshaped recommender systems, significantly improving effectiveness through greater accuracy and deeper personalisation via sophisticated machine learning, deep learning, and reinforcement learning. AI analyses vast datasets to predict user preferences with remarkable precision, enhancing user experiences across digital platforms. The successes of Netflix, Amazon, and TikTok highlight AI recommenders' ability to effectively personalise content, increasing user engagement and driving critical business outcomes like sales, retention, and loyalty [1-2].

As these systems increasingly integrate into our digital lives, addressing ethical considerations and promoting transparency are paramount. Finding the right balance between personalisation, diversity, and user control is essential for responsible development. Ultimately, AI's role in recommender systems is indispensable, continuously evolving to help users navigate information overload. For digital platforms striving for competitiveness and high user satisfaction, AI-driven recommendation is a fundamental necessity.

References:

1. Netflix content recommendation system – product analytics case study [Electronic resource]. – 2024. – Mode of access: https://hellopm.co/netflix-content-recommendation-system-product-analytics-case-study. – Date of access: 24.03.2025.

2. Amazon’s recommendation algorithm drives 35% of its sales [Electronic resource]. – 2020. – Mode of access: https://evdelo.com/amazons-recommendation-algorithm-drives-35-of-its-sales. – Date of access: 24.03.2025.

3. Necula, S.-C., Păvăloaia, V.-D. AI-Driven Recommendations: A Systematic Review of the State of the Art in E-Commerce / S.-C. Necula, V.-D. Păvăloaia, // Applied Sciences. – 2023. – Vol. 13, no. 9. – 5531. – doi: 10.3390/app13095531. – Mode of access: https://www.mdpi.com/2076-3417/13/9/5531. – Date of access: 24.03.2025.

4. See what’s next: how Netflix uses personalization to drive billions in revenue [Electronic resource]. – 2024. – Mode of access: https://www.rebuyengine.com/blog/netflix. – Date of access: 24.03.2025.

5. Chong, D. Deep dive into Netflix’s recommender system [Electronic resource] / D. Chong. – 2024. – Mode of access: https://medium.com/data-science/deep-dive-into-netflixs-recommender-system-341806ae3b48. – Date of access: 24.03.2025.

6. Cohen, B. Amazon’s secret to AI-powered product recommendations [Electronic resource] / B. Cohen. – 2020. – Mode of access: https://www.mdm.com/article/technology/ecommerce/amazons-secret-to-ai-powered-product-recommendations. – Date of access: 24.03.2025.