Reinforcement Learning with Knowledge Graphs for Personalized News Recommendation

Shyamsundar Pralhadrao Magar[1], *Metti* Vinodh Kumar[2*], *Manjula*. B M[3], *Leeladhar* Gudala[4]

[1]Department of Information Technology, Trinty College of Engineerng & Research, Pune, India
[2]Department of Electronics and Communication Engineering, St.Peter's Institute of Higher Education and Research, Chennai, India
[3]Department of Electronics and Communication Engineering, Nitte Meenakshi Institute of Technology, Nitte (Deemed to be University), Bengaluru, India
[4]Independent researcher

Abstract. In recent years, the rapid growth of online media and personalized news recommendations have become essential for reducing information overload and improving user engagement. However, traditional models significantly rely on semantic similarity and suffered to capture dynamic user preferences. Moreover, the extensive dissemination of misinformation poses serious limitations, necessitating recommendation systems that are both adaptive and knowledge aware. To overcome these limitations, a Reinforcement Learning with Knowledge Graphs (RL-KG) framework for personalized news recommendations. Initially, data was gathered from the MIND dataset, which consist of millions of users' news interactions enriched with titles, abstracts, and entity annotations. During preprocessing, text normalization and entity linking are employed for semantic clarity and structured knowledge representation. Furthermore, feature extraction integrates Bidirectional Encoder Representations from Transformer (BERT) embeddings for contextual semantics with Knowledge Graph embeddings for entity-level reasoning. Moreover, the Reinforcement Learning (RL) model was utilized for sequential decision-making procedures in which user histories form states, candidate news represents actions, and rewards provide engagement. Finally, the recommendation module integrates hybrid features and RL policies to deliver personalized, diverse news suggestions. The experimental results demonstrate that RL-KG achieves higher accuracy (96.77%), precision (95.55%), recall (96.45%), and F1-score (93.76%) than the existing model Recommendation for Mitigation (Rec4Mit).

1 Introduction

The rapid expansion of online news outlets has resulted in information overload, whereby readers must sift hundreds of articles every day. Personalized News Recommendation (PNR) systems have been developed to calibrate news articles according to user needs [1], thereby increasing their satisfaction and engagement [2]. The initial iterations of these PNR systems were based on collaborative filtering models and content-based filtering models, while existing methods have transitioned towards deep learning models that enable them to summarize contextual semantics as well as capture different and multi-dimensional user behavior [3]. Despite the aforementioned advancements, numerous challenges in building effective personalized news recommendation systems continue to inspire research [4]. First, user preferences change constantly and according to current events; therefore, static models ultimately fail to adequately capture user interests. Second, despite recognizing that news articles are typically very short, different from other article types, and fast-moving [5], once a high-assured, appropriate feature representation is found. Thirdly, fake or misleading news poses an even greater contribution to unreliable information; and there is a risk that a model based upon 'clicks,' an optimizing mechanism, cannot be generally trusted. Demonstrating diversity, transparency, and fairness when users use PNR systems remains an unresolved research topic. Taken together, these challenges and issues motivate the development of an efficient recommendation framework that is adaptive to uncertain user preferences, knowledge awareness, and trustworthy [6]. By integrating natural language processing with structured knowledge, systems can more accurately reflect user intentions and provide contextualized content [7]. Additionally, using reinforcement learning means that recommendations can adapt to changes in user behavior, and knowledge graphs allow for some degree of factual grounding and reliability [8]. These are the key features of recommendation systems. Most similar contemporary models have transitioned from shallow methods, ostensibly based on personalized item perspectives, to deep learning–based architectures (CNNs, RNNs, and Transformers) that learn specific representations of news [2]. Even newer models have progressed to complex frameworks all approaches argue to leverage the appropriate use of attention mechanisms and sequential models to personalize recommendations.

* Corresponding author: mettivinodhkumar@gmail.com

More recent work has shown a combination of knowledge-aware models that include entity embeddings [9], either on their own or as extensions of neural models, or recommender systems based on some form of reinforcement learning that frames the recommendation problem as sequential decision-making on a user's behalf. However, most of the model's reports optimize either personalization or optimizing credibility through some support of knowledge, but not both, leaving opportunities for hybrid Reinforcement Learning with Knowledge Graphs (RL-KG) types of models. The key contribution of this study follow.

- RL-KG framework is a reinforcement learning-based approach that leverages knowledge graph embeddings to provide real-time knowledge-aware personalized news recommendations. The incorporation of knowledge graphs allows our recommendations to dynamically adapt to evolving user interests to maximize content relevance while ensuring factual consistency and diversity of the recommendations.
- hybrid approach to feature extraction that combines Bidirectional Encoder Representations from Transformer (BERT) semantic embeddings with knowledge-graph-based entity embeddings. This ensures that we can extract not only contextual meaning but also meaningful structured factual relations, which improves understanding and explainability over previous methods.
- The RL-KG framework we employed was designed to be flexible and scalable to allow it to be implemented in real-life personalized news delivery systems. The RL-KG framework has many future deployment possibilities, including potentially in the future transformer-based temporal modeling for time-based recommendations and privacy-preserving federated learning for privacy protection when implementing personalized news recommendations.

This followed paper was organized follows. Related works and model drawbacks are presented in Section 2, an overall explanation of the proposed model is presented in Section 3, results are demonstrated in section 4, and Section 5 completes the paper.

2 Literature review

Wang et al. [10] presented personalized news recommendations to mitigate fake news. Initially, the PolitiFact dataset was collected from the Fake Newsnet dataset, which contains diverse features of news content, social content, and true and fake news articles. Further, the Rec4Mit model was projected, where each user's news was ordered by timestamp, context-target pairs were created, and only true news was targeted. Subsequently, the event-veracity disentangler was built on an adversarial autoencoder framework that parts event-related and veracity-related data using a decoder, encoder, veracity decoder, and veracity classifier. Furthermore, the event transition and detection module uses attention-based transition and SoftMax-based detector to capture latent events. Later, the next news prediction module was used to estimate the likelihood of each candidate news item, and the Adam optimizer deals with negative sampling. Moreover, this model separates event information from news veracity, which generates true news. However, this model depends on accurate event detection, which makes it complex.

Lu et al. [11] demonstrated news recommendation using news representation and user preference learning. Data were collected from the Microsoft News Dataset (MIND), which contains the number of logs, users, and news from different topics. Furthermore, an Aspect-driven News Recommender System (ANRS) architecture was created using three techniques, and a news aspect-encoder level extracts level-aspect features from articles. Later, a user level aspect encoder was encode employed to preference user level, and predictor of click calculated the probability of click the user clicking on new items. Furthermore, negative sampling and regularization terms were used to optimize the training model to ensure the uniqueness and diversity of embeddings. Subsequently, the CNN filter extracts features, and k refers to the number of aspects. Moreover, this model successfully captures fine-grained aspect-level information. However, an articles of large number news have released the challenge of choosing appropriate news items daily.

Chen et al. [12] presented an interactive recommendation system using knowledge-guided reinforcement learning. Initially, data were collected from the Book-Crossing dataset, which contained users' demographic information. Furthermore, a knowledge-guided deep reinforcement learning (KGRL) framework was proposed. Then, a knowledge graph was constructed based on an entity relation entity tuple to capture semantic relations, which was then modified into a user-specific adjacency matrix. Later, an actor-critic network framework was implemented to preserve the local knowledge system that guides decision and attention making mechanism was employed to captured the semantics of long-term among items. Moreover, this model combines knowledge graphs, and reinforcement earning captures long-term relations. However, the cold-start problem and long training times are due to the actor–critic framework.

Di Wu et al. [13] demonstrated a recommendation system using a deep knowledge-based reinforcement-learning network. data were collected from a Monielens-1M movie related dataset containing users, interactions, and knowledge graph entities. Further, a knowledge-preference-aware reinforcement learning network (KPRLN) was projected, which features preference of user fine grained. Later, this model learned the feature preference entity relation of each user and constructs paths among the user's interaction of historical items in the graph knowledge. Furthermore, learning of deep reinforcement uses cluster expansion, which explores and updates weights in the knowledge graph and design feedback. Later, in the convolution layer, the attention mechanism focuses on entity relation pairs and propagates higher-order interests that contain fine-

grained features in the knowledge graph. Moreover, this model effectively learns dynamic features and improves the recommendation explainability. However, there is high computational cost and difficulty in dealing with large datasets.

Zhang et al. [14] presented an interactive and explainable recommendation model using knowledge graph hierarchical reinforcement learning. Initially, data were collected from the Movielens-20M dataset which contains user ratings and interactions. Furthermore, knowledge graph-enhanced hierarchical reinforcement learning (KGHR) was developed for explainable and interactive recommendations. Later, two components, a high-level recommendation agent (HRA) and a low-level recommendation agent (LRA), used double-Q for ranking and recommending the best item. further, a Graph Attention Network (GAT) and transformers are used for modeling user interactions. Moreover, LRA uses a policy network to perform path reasoning and guide candidate sampling with beam search, and then the two agents are trained together using experience replay and target networks to avoid divergence and stabilize learning. Moreover, both enhance sample efficiency and provide explainable recommendations. However, sequential recommendations cannot be achieved.

3 Methodology

This study proposes a methodology that merges Reinforcement Learning (RL) and Knowledge Graphs (KGs) to create a personalized news recommendation system that is adaptive, trustworthy, and explainable. First, we preprocessed the MIND dataset using text normalization and entity linking to provide semantic trust and structured knowledge interoperability. The features utilize BERT embeddings for contextual semantics and KG embeddings for meaningful entity-level thinking representations of personalized hybrid recommendation features. Construct a recommendation model based on sequential decision-making, in which RL can effectively learn a policy that considers the balancing aspects of engagement, diversity, and factual credibility. As a result, users receive personalized, timely, and knowledge-infused news feeds that match their personalized knowledge and develop over time. A representation of the proposed model as shown in Figure 1.

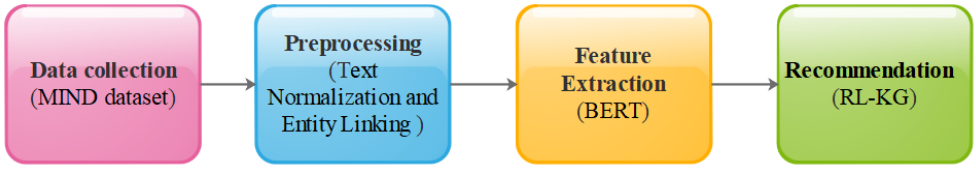

Fig. 1. Pictorial representation of proposed RL-KG.

3.1 Dataset

This dataset uses the Microsoft News Dataset, a large-scale benchmark dataset widely used in personalized news recommendations. MIND contains millions of records of user–news interactions collected from Microsoft News [11], which covers user interaction data from a number of news categories such as politics, sports, business, entertainment, and health. Each news article comes with metadata that include title, abstract, category, and entity annotations to support not only semantic modeling but also knowledge graph-based modeling. Indeed, logs of user interactions capture sequential click behaviors, which reveal the evolving preferences of users. The MIND dataset contains a rich set of entity annotations, making it easy to use in conjunction with Knowledge Graphs and a great base for RL-based personalized recommendations.

3.2 Preprocessing

Preprocessing is highly relevant for preparing unprocessed news data for personalized recommendations. In our work, we took two key preprocessing steps, Text Normalization and Entity Linking, to ensure that the data that will be used for downstream feature extraction will be structured, clean, and semantically valuable.

3.2.1 Text normalization

Raw news articles typically contain noise, including stop words, punctuation, and variable word forms. To standardize the textual input, tokenization, lowercasing, stopword removal, and lemmatization are applied. Each article's text is then transformed into a vector form, where TF - IDF weighting was conducted to assess the importance terms. The TF-IDF score for term t in document d is shown in Equation (1):

$$TF - IDF(t,d) = TF(t,d) \times log\ \frac{N}{DF(t)} \quad (1)$$

Where $\mathrm{TF}(t,d)$ is the frequency of term t in document d, N is the total number of documents, and IDF was the documents number that contain t. This ensures the frequent but less informative words have lower weights, whereas context-rich terms are more important.

3.2.2 Entity linking

In addition to text, entity linking is conducted to associate the important terms in news articles with external Knowledge Graph (KG) nodes. This detects entities such as people, places, and organizations, and maps them to KG concepts. Each entity is embedded to generate a vector through graph learning. The

embedding of an entity is updated using neighborhood aggregation, as shown in Equation (2):

$$h_i^{(l+1)} = \sigma(\sum_{j \in \mathcal{N}(i)} W^{(l)} h_j^{(l)}) \quad (2)$$

Where $h_i^{(l+1)}$ is the entity representation at layer l, $\mathcal{N}(i)$ is the neighborhood of entity i, and $W^{(l)}$ is the trainable weight matrix. Normalization and entity linking contribute to feature extraction by ensuring that the input features maintain semantic soundness and structured relationships between entities, thus providing a solid foundation for feature extraction and reinforcement learning.

3.3 Feature extraction

Feature extraction is crucial for transforming raw news data into semantically rich and knowledge-augmented forms, which map a wide variety of raw news content-derived feature vectors useful for personalized recommendation. The proposed system uses two complementary methodologies: BERT embeddings for textual representation and KG embeddings using entity-based reasoning. The concatenation of the two feature sets creates a hybrid feature space capable of capturing linguistic semantics and structured knowledge.

3.3.1 BERT embeddings for semantic features

For the textual components of news articles, Bidirectional Encoder Representations from Transformers (BERT) are used. In contrast to embedding techniques that only keep track of the left or right context, BERT uses a left or right context when generating contextualized vectors, creating highly contextualized vectors. First, the news article title and abstract were tokenized, and the titles and abstracts were entered into the pre-trained BERT model, which generated embedding at the word level. These are then aggregated to create a sentence vector representation, as shown in Equation (3):

$$h_i^{(l)} = TransformerLayer(h_i^{(l-1)}, \{h_j^{(l-1)} \mid j \in context(i)\}) \quad (3)$$

Where $h_i^{(l)}$is the input embedding and the transformer layer applies self-attention over the token's context. This ensures that important terms such as "pandemic" or "election results" are emphasized in the final representation.

3.3.2 Knowledge graph embeddings for entity features

Structured knowledge integration extracts entities with Named Entity Recognition (NER) [15] and links them with existing Knowledge Graphs (DBpedia, Wikidata). Each entity is embedded using graph representation learning methods such as TransE or GNNs. For GNN-based embeddings, the representation of an entity is updated according to the features of its neighbor's Equation (4):

$$h_i^{(l+1)} = \sigma(W^{(l)} \cdot AGG(\{h_j^{(l)} : j \in \mathcal{N}(i)\})) \quad (4)$$

Where $j \in \mathcal{N}$ denotes the neighbors, $W^{(l)}$ is a trainable weight matrix, and σ is a nonlinear activation function. This ensures that the embedding of "COVID-19" captures its relational context with entities such as "vaccines," "Pfizer," and "WHO."

3.3.3 Fusion of features

Ultimately, BERT-based semantic embeddings and KG-based entity embeddings are fused into one representation that is processed through a dense projection layer or attention mechanism to control the contributions from both BERT and KG embeddings. Therefore, a hybrid feature vector that contains the word embeddings deep linguistic semantics from BERT, as well as explicit relationships between entities from the KG, allows the recommendation algorithm to receive deep context-rich and knowledge-aware insights. Moreover, with BERT and KG embeddings together, its feature is to guarantee sufficient and explainable representations from both types of embeddings to support the framework for reinforcement-learning-based personalized news recommendations.

3.4 News recommendation system

Following the feature extraction process, where each news article is represented through hybrid embedding incorporating KGs and semantic features from BERT, the system proceeds into the recommendation phase under an RL-based process. In this phase, the news recommendation task is modelled as a sequential decision-making problem. Instead of clicks, the conceptual agent interacts with users to optimize long-term satisfaction. At each time step t, the user's state S_t is constructed using their historical interactions, which are stored in the hybrid embeddings of previously clicked articles. To emphasize more relevant interactions, we use an attention mechanism with higher weights assigned to impacting articles, as shown in Equation (5):

$$S_t = f(\sum_{i=1}^{n} \alpha_i \cdot (h_i^{BERT} \oplus h_i^{KG})) \quad (5)$$

Where h_i^{BERT} and h_i^{KG} are the semantic and KG embeddings of the i^{th} article, respectively, α_i is the attention weight, and $\oplus$ represents concatenation. This ensures that the state captures both semantic context and factual knowledge. The action space, denoted by A_t describes the set of candidate news articles that the agent can recommend. The candidates are in the same embedding space such that the RL agent can compare and select the best article. The learning agent uses a reward function to beacon its decision making. It receives a positive reward if the user clicks on, or spends time reading or sharing the article if the article is

irrelevant, redundant, or not factually correct; then, the agent penalty will be activated. The agent's target was to maximizing the summation of the discounted rewards, as shown in Equation (6).

$$J(\pi) = \mathbb{E}_{\pi}[\sum_{t=0}^{T} \gamma^t R_t] \quad (6)$$

Where π is the policy, R_t is the reward at time step t, and γ^t is a discount factor that balances immediate engagement with long-term satisfaction. To learn the optimal policy, the agent employs Deep Q-Networks (DQN) or actor-critic algorithms. In the DQN framework, the Q-value function, which estimates the expected utility of recommending a particular article, is updated iteratively, as shown in Equation (7).

$$Q(S_t, A_t; \theta) \approx R_t + \gamma \underset{A'}{max} Q(S_{t+1}, A'; \theta^-) \quad (7)$$

Where θ and θ^- were parameters of the online and the target networks, respectively. Through repeated interactions, the agent improves its decision-making process and gradually learns to recommend news items that maximize both engagement and reliability. Consequently, the system produces a personalized and ranked recommendation list for each user. Unlike static recommenders, this output is dynamic, knowledge aware, and explainable. For instance, if a user frequently interacts with political news, the system may recommend fresh but related articles such as “election debates” or “economic policies,” while avoiding redundancy.

4 Experimental results

Experimental results were obtained in a high-performance computing environment to ensure efficient model training and validation. This environment uses multicore processors such as Intel i7, at least 16 GiB of random-access memory (RAM), and a GPU such as NVIDIA GTX 1080 to perform the DL training. In particular, the software environment includes Python and DL libraries such as TensorFlow and PyTorch. The RL-KG results in an optimized and highly relevant input set that enhances classification accuracy, reduces redundancy, and ensures interpretability for science and technology finance evaluations. The evaluation metrics utilized are formulated in Equations (8)-(11), as follows:

$$Accuracy = \frac{TP+TN}{TP+TN+FP+FN} \quad (8)$$

$$Precision = \frac{TP}{TP+FP} \quad (9)$$

$$Recall = \frac{TP}{TP+FN} \quad (10)$$

$$F1-score = \frac{2\times pr\times re}{pr+re} \quad (11)$$

Here, TN and TP denote true negatives and positives, respectively, FN, FP denote false negatives and positives.

4.1 Performance analysis

The performance of the proposed model was compared with that of various existing methods, such as the Deep Knowledge-Aware Network (DKN), news attention multilevel model (NATML), and Long- and Short-Term User Representation (LSTUR) using evaluation metrics. Existing approaches are listed in Table 1.

Table 1. Performance analysis of proposed RL-KG.

Performance Models	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)
DKN	82.55	84.67	83.41	81.66
NATML	88.44	87.67	85.33	88.67
LSTUR	92.56	94.88	92.47	93.32
Proposed RL-KG	96.77	95.55	96.45	93.76

Table 1 presents the performance of the proposed RL-KG framework compared with the existing methods. RL-KG achieved better results than the traditional approaches, including accuracy (96.77%), precision (95.55%), recall (96.45%), and F1-score (93.76%), outperforming the DKN model in terms of accuracy (82.55%), precision (84.67%), recall (83.41%), and F1-score (81.66%). Moreover, the NATML model exhibited accuracies of (88.44%), precision (87.67%), recall (85.33%), and F1-score (88.67%), respectively. The LSTUR model in terms of accuracy (92.56%), precision (94.88%), recall (92.47%), and F1-score (93.32%).

4.2 Comparative analysis

The proposed RL-KG framework was compared with the existing Rec4Mit [10] model by using comprehensive evaluation metrics. An evaluation of the proposed RL-KG method is presented in Table 2.

Table 2. Comparative analysis of proposed RL-KG.

Comparative Models	Accuracy (%)	Precision (%)	Recall (%)	F1-score (%)
Rec4Mit [10]	83.00	81.00	80.00	81.00
Proposed RL-KG	96.77	95.55	96.45	93.76

Table 2 shows that the proposed RL-KG framework obtained better results than the existing Rec4Mit [10] model in terms of accuracy (96.77%), precision (96.16%), recall (96.32%), and F1-score (96.19%) which represent significant improvements in model performance.

4.3 Discussion

The primary objective of this research is to propose a model that progresses personalized news recommendations by joining RL with KG embedding. Existing models such as the Deep Knowledge-Aware Network (DKN), news attention multilevel model (NAML), and Long- and Short-Term User Representation (LSTUR) have progressive the field, but

they face prominent limits. DKN participates semantic and entity-level information; however, it ruins static and suffers to adapt to developing user preferences. Moreover, the NAML captures multilevel semantic features through attention but lacks explicit reasoning from structured knowledge and does not address credibility. Additionally, LSTUR models both long- and short-term user preferences but struggles with limited explainability and redundancy issues. More recently, Rec4Mit [10] presented a RL-based approach that integrates event and veracity information for fake news mitigation by obtaining the robust performance but narrowly focusing on misinformation. To addresses these challenges, the proposed RL-KG framework fuses BERT semantic embeddings with KG-based entity reasoning, whereas RL dynamically optimizes recommendation policies by user feedback. This simultaneously confirms personalization, credibility, and diversity. Overall, RL-KG spreads beyond existing models by as long as an adaptive, knowledge-aware, and trustworthy solution and carrying reliable recommendations suited for dynamic digital environments.

5 Conclusion

The proposed RL-KG framework is a controlling, adaptive, and knowledge-aware system for personalized news recommendations in advanced digital environments. By incorporating the RL with KG embeddings, the framework effectively overcome the limitations of evolving user preferences, information overload, redundancy, and misinformation. The hybrid feature-extraction process joins BERT-based semantic embeddings with structured KG reasoning, which confirms that both contextual meaning and factual relationships are preserved. The RL component frames recommendations as a sequential decision-making task by allowing the system to incessantly adapt to user behavior while balancing engagement and credibility. Finally, the experimental estimation establishes that the RL-KG outperforms existing models like DKN, NAML, LSTUR, and Rec4Mit by obtaining the enhanced results in terms of accuracy (96.77%), precision (95.55%), recall (96.45%), and F1-score (93.76%). Moreover, the inclusion of KG embeddings delivers explainable recommendations to progress the transparency and user trust. Unlike static approaches, the proposed system evolves dynamically, thereby confirming relevance and reliability in a constantly changing news ecosystem. In future work, the proposed framework will be prolonged with transformer-based sequence modeling to progress the temporal understanding of user behavior and privacy-preserving Federated Learning (FL) to strengthen the security and scalability in real-world deployment.

References

1. W. Deng, P. Zhu, H. Chen, T. Yuan, J. Wu, Knowledge-aware sequence modelling with deep learning for online course recommendation. Inf. Process. Manag. **60**, 103377 (2023). https://doi.org/10.1016/j.ipm.2023.103377
2. S. Bansal, K. Gowda, N. Kumar, Multilingual personalized hashtag recommendation for low resource Indic languages using graph-based deep neural network. Expert Syst. Appl. **236**, 121188 (2024). https://doi.org/10.1016/j.eswa.2023.121188
3. D. Cai, S. Qian, Q. Fang, J. Hu, W. Ding, C. Xu, Heterogeneous graph contrastive learning network for personalized micro-video recommendation. IEEE Trans. Multimed. **25**, 2761-2773 (2022). https://doi.org/10.1109/TMM.2022.3151026
4. Z. Shokrzadeh, M.R. Feizi-Derakhshi, M.A. Balafar, J.B. Mohasefi, Knowledge graph-based recommendation system enhanced by neural collaborative filtering and knowledge graph embedding. Ain Shams Eng. J. **15**, 102263 (2024). https://doi.org/10.1016/j.asej.2023.102263
5. J. Zhou, G. Jiang, W. Du, C. Han, Profiling temporal learning interests with time-aware transformers and knowledge graph for online course recommendation. Electron. Commer. Res. **23**, 2357-2377 (2023). https://doi.org/10.1007/s10660-022-09541-z
6. D.H. Tran, Q.Z. Sheng, W.E. Zhang, N.H. Tran, N.L. Khoa, CupMar: A deep learning model for personalized news recommendation based on contextual user-profile and multi-aspect article representation. World Wide Web **26**, 713-732 (2023). https://doi.org/10.1007/s11280-022-01059-6
7. R. Alabduljabbar, H. Almazrou, A. Aldawod, Context-aware news recommendation system: Incorporating contextual information and collaborative filtering techniques. Int. J. Comput. Intell. Syst. **16**, 137 (2023). https://doi.org/10.1007/s44196-023-00315-5
8. M. Wang, Y. Hu, S. Wu, W. Li, Q. Bai, Z. Yuan, C. Jiang, Nudging toward responsible recommendations: A graph-based approach to mitigate belief filter bubbles. IEEE Trans. Artif. Intell. **6**, 378-392 (2024). https://doi.org/10.1109/TAI.2024.3373392
9. T. Anwar, V. Uma, CD-SPM: Cross-domain book recommendation using sequential pattern mining and rule mining. J. King Saud Univ. Comput. Inf. Sci. **34**, 793-800 (2022). https://doi.org/10.1016/j.jksuci.2019.01.012
10. S. Wang, X. Xu, X. Zhang, Y. Wang, W. Song, Veracity-aware and event-driven personalized news recommendation for fake news mitigation, In Proceedings of the ACM web conference 2022, The ACM Web Conference, Virtual Event, Lyon, France, April 25 (2022), 3673-3684
11. W. Lu, R. Wang, S. Wang, X. Peng, H. Wu, Q. Zhang, Aspect-driven user preference and news representation learning for news recommendation. IEEE Trans. Intell. Transp. Syst. **23**, 25297-25307

(2022). https://doi.org/10.1109/TITS.2022.3182568

12. X. Chen, C. Huang, L. Yao, X. Wang, W. Zhang, Knowledge-guided deep reinforcement learning for interactive recommendation. In 2020 International Joint Conference on Neural Networks (IJCNN), IEEE, Glasgow, UK, September 20 (2020), 1-8
13. D. Wu, M. Tang, S. Zhang, A. You, W. Gao, KPRLN: deep knowledge preference-aware reinforcement learning network for recommendation. Complex Intell. Syst. **9**, 6645-6659 (2023). https://doi.org/10.1007/s40747-023-01083-7
14. M. Zhang, Y. Li, S. Li, Y. Wang, J. Yan, Knowledge graph-enhanced hierarchical reinforcement learning for interactive and explainable recommendation. IEEE Access **12**, 137345-137359 (2024). https://doi.org/10.1109/ACCESS.2024.3464733
15. U. Aslam, F. Dine, Unifying NLP and Knowledge Graphs: AI-Powered Techniques for Intelligent Information Extraction. (2025)