The European Journal of Technical and Natural Sciences 2025, No 4–5
ISSN 2414-2352

PREMIER Publishing
ppublishing.org

DOI:10.29013/EJHSS-25-4.5-18-29

CC BY

PERSONALIZATION OF NEWS CONTENT USING AI: ARCHITECTURE AND EFFICIENCY OF NEURAL NETWORK SYSTEMS

I. S. Shalamov [1]

[1] Independent Researcher, Russia, Moscow

Cite: *Shalamov I.S. (2025). Personalization of News Content Using Ai: Architecture and Efficiency of Neural Network Systems. European Journal of Technical and Natural Sciences 2025, No 1*. https://doi.org/10.29013/EJHSS-25-4.5-18-29

Abstract

The rapid growth of digital information has made news content personalization a critical tool for addressing information overload and enhancing user engagement. This study investigates the architecture and efficiency of neural network systems for personalizing news content, focusing on a proposed hybrid model combining BERT for semantic text analysis and LSTM for modeling temporal user interactions. Using the MIND dataset, we compared the hybrid model against baseline architectures (TextCNN, LSTM, BERT-base) and traditional matrix factorization. The hybrid model achieved superior performance, with Precision@5 of 0.78, Recall@5 of 0.75, F1-score of 0.76, and a click-through rate (CTR) of 6.3%, outperforming baselines by up to 43% in F1-score and 66% in CTR. The model demonstrated robustness to noisy data and improved recommendations for cold-start scenarios. However, challenges such as computational complexity and applicability to multilingual datasets remain. The findings highlight the potential of hybrid neural architectures for real-world news personalization while underscoring the need to address ethical concerns, such as filter bubbles and privacy. Future research should explore multimodal data integration and lightweight models to enhance scalability and diversity in recommendations.

Keywords: *News personalization, Neural networks, Recommender systems, BERT, LSTM, Hybrid models, Content-based filtering, Collaborative filtering, Filter bubbles, Artificial intelligence*

Introduction

The growing digitization of the information landscape is characterized by an exponential increase in available content volume, necessitating the development of effective personalization mechanisms to overcome user information overload. In the context of news platforms, content personalization represents the adaptation of information flows to individual audience characteristics, including preferences, behavioral patterns, and contextual factors, thereby enhancing user engagement and overall platform interaction satisfaction.

Contemporary personalization methods are based on three primary approaches: collaborative filtering, grounded in the analysis of user or item similarity; content-based methods utilizing content characteristics

themselves; and hybrid systems integrating the advantages of both paradigms. Traditional algorithms, including matrix factorization methods, demonstrate limitations in processing multidimensional textual data and adapting to dynamically changing user preferences.

The application of neural network architectures opens qualitatively new possibilities for solving personalization tasks through their capacity to identify latent patterns in large data arrays. Convolutional neural networks (CNNs) effectively extract local features from textual data, recurrent architectures (RNNs) model temporal dependencies in user behavior, while transformer models such as BERT provide deep understanding of contextual semantics in news. Despite the significant potential of neural network approaches, their practical implementation is associated with a range of technical and ethical challenges. Critical issues include: the cold start problem for new users and content, systematic biases in training data, the formation of informational filter bubbles, and potential violations of user privacy.

The present study aims to provide a systematic analysis of architectural solutions for neural network-based news content personalization systems and comprehensive evaluation of their effectiveness based on quantitative metrics (precision, recall, click-through rate) and qualitative indicators of user experience. The work encompasses both technical aspects of comparative analysis of various neural network models and critical examination of practical and ethical considerations in their deployment in production systems.

Materials and methods

Data. Two primary types of datasets were utilized for conducting a comprehensive study: open-source news datasets and synthetic user data, enabling both experimental realism and controlled testing conditions.

The primary source of news data was the MIND (Microsoft News Dataset), representing a large-scale collection containing over 1 million English-language news articles. The dataset includes rich metadata: article headlines, abstracts, categorical markup (politics, sports, technology, entertainment, business, etc.), as well as detailed user interaction logs including clicks, views, interaction timestamps, and action sequences. The temporal coverage spans from October to November 2019, providing sufficient data volume for model training and validation. Of particular value are the anonymized behavioral data from over 1 million users, encompassing interaction histories with news content.

For modeling diverse user profiles, synthetic data were generated, incorporating detailed demographic characteristics (age ranging 18–65 years, gender distribution, geographic location by regions), professional affiliation, education level, and individual preferences across 17 primary news categories. Synthetic data were created using an extended Monte Carlo method accounting for realistic statistical distributions of demographic characteristics and correlations between various user attributes. To enhance realism, temporal user activity patterns, seasonal preferences, and evolution of interests over time were modeled.

Data preprocessing included a multistage procedure: text tokenization using BPE (Byte Pair Encoding), stop word removal based on an extended NLTK vocabulary, lemmatization and stemming, numerical value normalization, and categorical variable encoding. Additionally, anomalous value filtering and missing data handling were performed using interpolation and imputation methods.

Neural Network Architectures. The study conducted a systematic comparison of three primary types of neural network architectures adapted for news content personalization tasks:

- **Convolutional Neural Networks (CNN):** The base architecture was founded on the TextCNN model with substantial modifications for recommender system tasks. The model included three parallel convolutional layers with filters of varying sizes (3, 4, and 5), enabling capture of n-grams of different lengths and identification of local patterns in texts. The convolutional layers were followed by max-pooling layers for extracting the most significant features, batch normalization for training stabilization, and dropout layers (coefficient 0.3) for overfitting prevention. The output

layer dimensionality comprised 256 neurons;

- **Recurrent Neural Networks (RNN):** An enhanced bidirectional LSTM architecture with 128 hidden units in each direction was employed, specifically adapted for processing user interaction sequences and analyzing temporal dependencies in preferences. The model incorporated attention mechanisms for focusing on the most relevant sequence elements, as well as residual connections for improved gradient flow. Gradient clipping techniques were additionally applied to prevent vanishing gradient problems;
- **Transformers:** The foundation consisted of a pre-trained BERT-base model with 12 transformer layers and 768-dimensional embeddings, fine-tuned on the MIND dataset for specific news content relevance classification tasks. The fine-tuning procedure included freezing lower layers and training the upper 4 layers with adaptive learning rates. For user data processing, a simplified transformer with 4 encoder layers, 8 attention heads, and an intermediate layer of dimensionality 1024 was developed, optimized for working with user interaction sequences;
- **Hybrid Model:** An innovative architecture was proposed that combines the advantages of BERT for semantic text content analysis and LSTM for modeling dynamic temporal dependencies in user behavior. The architecture includes two primary components: a BERT-based text encoder for extracting contextual representations of news articles and an LSTM-based user encoder for modeling user preference evolution. Integration is achieved through a multi-layer perceptron with three hidden layers (512, 256, 128 neurons), ReLU activation functions, and regularization mechanisms. Cross-attention mechanisms were additionally implemented for efficient combination of textual and user representations.

Personalization Methods. News content personalization was implemented through comprehensive integration of several complementary approaches:

- **Content-based Filtering:** High-dimensional news text embeddings (768-dimensional vectors) extracted from the fine-tuned BERT model were used for precise computation of semantic similarity between articles and user preferences. Both cosine similarity metrics and Euclidean distance with weight coefficients optimized for each content category were applied. TF-IDF characteristics and semantic clusters obtained through k-means clustering were additionally considered;
- **Collaborative Filtering:** An enhanced neural matrix factorization algorithm with deep architecture was implemented for modeling complex nonlinear interactions between users and content. The model includes embedding layers for users and news articles (dimensionality 64), several fully connected layers with batch normalization and dropout for latent factor extraction, and an output layer with sigmoid activation for predicting interaction probability;
- **Hybrid Approach:** A comprehensive system was developed that integrates content-based embeddings, collaborative interaction data, and a rich set of contextual factors through the proposed hybrid neural architecture. Contextual variables included: user geographic location, time of day and day of week, seasonality, access device, previous session activity, and dynamic article popularity characteristics. All contextual factors were transformed into vector representations and integrated into the model as additional input features through specialized embedding layers.

Effectiveness Evaluation. Comprehensive evaluation of the proposed models' effectiveness was conducted using an extended set of quantitative and qualitative metrics:

- **Precision@K:** The proportion of relevant recommendations among K suggested articles was computed for various K values (5, 10, 20, 50), enabling

quality assessment of recommendations at different positions in the list;

- **Recall@K:** Defined as the proportion of all relevant articles successfully found in top-K recommendations, considering various relevance definitions (clicks, reading time, expert evaluation);
- **F1-Score and Modifications:** Classical F1-score was calculated as the harmonic mean between precision and recall, along with weighted F-score versions to account for the importance of different error types;
- **CTR (Click-Through Rate):** The percentage of clicks on recommended articles was measured across various time windows (1 hour, 6 hours, 24 hours post-recommendation) to evaluate both short-term and long-term effectiveness;
- **Diversity Metrics:** Additional indicators were introduced for assessing recommendation diversity, including intra-list diversity (ILD), catalog coverage, and recommendation novelty to prevent filter bubble effects;
- **User Satisfaction:** Comprehensive evaluation obtained through simulated user surveys based on synthetic data with modeling of various user behavior types and preferences.

The experimental methodology included rigorous A/B testing with randomized user distribution between control and experimental groups. The proposed hybrid model was systematically compared against baseline architectures (TextCNN, LSTM, BERT) and classical methods (collaborative filtering, popularity). To ensure statistical reliability of results, 5-fold stratified cross-validation with temporal data order preservation was applied. All models were tested on a standard split of 80% training, 10% validation, and 10% test data from the MIND dataset with additional early stopping procedures to prevent overfitting.

Tools and Technologies. Technical implementation of models was performed using a modern machine learning technology stack:

- **Core Frameworks:** PyTorch version 1.9 for neural network construction, training, and optimization with automatic differentiation and GPU acceleration support; TensorFlow 2.8 for separate experiments and comparative analysis;
- **Specialized Libraries:** Hugging Face Transformers for efficient work with pre-trained transformer models, including BERT, with fine-tuning and tokenization support; NLTK and spaCy for comprehensive text preprocessing, including tokenization, lemmatization, POS-tagging, and named entity extraction;
- **Data Processing:** Pandas and NumPy for efficient processing and manipulation of large volumes of structured data; Dask for parallel processing of data exceeding RAM capacity; Apache Spark for distributed cluster processing;
- **Machine Learning:** Scikit-learn for quality metric computation, baseline algorithm implementation, and data preprocessing; XGBoost for comparative experiments with gradient boosting;
- **Computational Infrastructure:** Experiments were conducted on high-performance servers equipped with NVIDIA Tesla V100 graphics accelerators (32 GB video memory) for neural network training acceleration, as well as on the scalable Amazon Web Services cloud platform (EC2 P3 instances) for distributed computing and hyperparameter optimization. Docker containers were used to ensure experiment reproducibility and MLflow for experiment tracking and model versioning.

The presented methodology provides a comprehensive approach to analyzing neural network architectures for news content personalization. The utilization of the large-scale MIND dataset combined with synthetic user data creates a robust empirical foundation for evaluating the effectiveness of various personalization approaches. The proposed hybrid architecture, integrating transformer models for semantic analysis and recurrent networks for modeling temporal dependencies, represents an innovative approach to personalization tasks. The combination of content-based and collaborative methods with contextual

factors enables the creation of a more accurate recommendation system. The multi-criteria evaluation system, incorporating both classical information retrieval metrics and recommender system-specific indicators, provides comprehensive assessment of personalization quality. The application of rigorous statistical validation methods ensures reliability and reproducibility of results. Technical implementation based on modern machine learning frameworks and high-performance infrastructure creates conditions for efficient training of complex models. The presented methodology establishes the foundation for obtaining objective results from comparative analysis of neural network architectures and formulating evidence-based recommendations for optimal approaches to news content personalization.

Results

Experiment Description. The experiments aimed to compare the effectiveness of four approaches to news content personalization: three neural network architectures (TextCNN, bidirectional LSTM, BERT-base) and the proposed hybrid model combining BERT for text processing and LSTM for analyzing temporal dependencies in user interactions. All models were trained and tested on the MIND (Microsoft News Dataset) containing over 1 million news articles and user interaction logs. The dataset was split into training (80%) and test (20%) sets using stratified division to preserve news category proportions.

Training was conducted using the Adam optimizer with an initial learning rate of 0.001 and L2 regularization (weight = 0.01) to prevent overfitting. For TextCNN, filter sizes (3, 4, 5) and number of filters (100, 200) were varied; for LSTM – number of hidden units (64, 128, 256) and layers (1, 2); for BERT-base – number of fine-tuning epochs (3, 5, 10). The hybrid model combined BERT embeddings (dimensionality 768) with LSTM outputs (128 hidden units) through a fully connected layer with ReLU activation and subsequent softmax for relevance prediction. To enhance result reliability, 5-fold cross-validation was applied, with each experiment repeated 5 times using different initial weight initializations. Additionally, A/B testing was conducted in simulated scenarios where model recommendations were compared against a baseline system based on classical matrix factorization (MF).

Quantitative Results

Model effectiveness was evaluated using four metrics: Precision@5 (proportion of relevant recommendations in top-5), Recall@5 (proportion of relevant articles found in top-5), F1-score (harmonic mean of Precision and Recall), and CTR (click-through rate, percentage of clicks on recommendations). Results are presented in Table 1.

Table 1. *Model Performance Comparison on MIND Test Dataset*

Model	Precision@5	Recall@5	F1-Score	CTR (%)
MF (Baseline)	0.55	0.52	0.53	3.8
TextCNN	0.62	0.58	0.60	4.2
LSTM	0.65	0.61	0.63	4.8
BERT-base	0.71	0.68	0.70	5.5
Hybrid (BERT+LSTM)	0.78	0.75	0.76	6.3

The hybrid model demonstrated the best results across all metrics, achieving Precision@5 = 0.78, Recall@5 = 0.75, F1-score = 0.76, and CTR = 6.3%. Compared to the baseline matrix factorization (MF) model, the hybrid model improved F1-score by 43% and CTR by 66%. BERT-base showed significant advantage over TextCNN and LSTM due to deep semantic analysis of news texts, while LSTM performed better with temporal dependencies in user interactions compared to TextCN N. TextCNN, despite its simplicity, proved least effective among neural network models due to limited ability to capture complex contextual dependencies in texts.

Hyperparameter analysis showed that increasing BERT fine-tuning epochs from 3

to 10 improved the hybrid model's F1-score from 0.68 to 0.76, indicating the importance of fine-tuning for adaptation to specific news domain data. For LSTM, optimal performance was achieved using 128 hidden units and two layers, providing balance between computational complexity and modeling quality. A/B testing in simulated scenarios showed that the hybrid model increased CTR by 15% compared to BERT-base, 31% compared to LSTM, and 50% compared to TextCN N. Compared to baseline MF, CTR improvement was 66%, highlighting the effectiveness of neural network approaches in personalization tasks.

Additionally, model robustness to data noise was investigated. 10% random interactions (noise) were artificially added to the test set, leading to decreased Precision@5 for all models, but the hybrid model maintained its advantage, losing only 5% accuracy, while TextCNN and MF showed decreases of 12% and 15% respectively.

Qualitative Results

To evaluate recommendation quality, examples of personalized news were analyzed for synthetic user profiles with different interests and contexts. For instance, for a user with profile "interests: technology, geolocation: USA, time: evening," the hybrid model suggested articles about latest AI and cybersecurity achievements, such as "New NVIDIA AI Chips" or "Cloud Computing Trends." Meanwhile, TextCNN more frequently recommended general technology news like "New Smartphone Reviews," ignoring contextual features, while MF suggested less relevant articles from adjacent categories (e.g., "Electronics News"). BERT-base showed good results in semantic relevance but was inferior to the hybrid model in accounting for user temporal preferences.

User satisfaction was evaluated through simulated surveys based on synthetic data. Users were asked to rate recommendation relevance on a scale from 1 to 5. The average score for the hybrid model was 4.1 (82%), for BERT-base – 3.6 (72%), for LSTM – 3.4 (68%), for TextCNN – 3.25 (65%), and for MF – 3.0 (60%). These results confirm that integrating semantic text analysis (BERT) with temporal dependency modeling (LSTM) enables creation of more relevant and personalized recommendations. Additionally, it was noted that the hybrid model better handled cold start scenarios, using content-based embeddings for new users, improving recommendation quality by 10% compared to LSTM in such scenarios.

Additional Observations

During experiments, it was revealed that the hybrid model demonstrated lower sensitivity to filter bubble problems. For example, for users with narrow interests (e.g., only "politics"), the model suggested diverse articles from related topics (e.g., "economics" or "international relations"), achieved through inclusion of contextual features in the architecture. It was also noted that inference time for the hybrid model (average 0.12 seconds per recommendation) remained acceptable for real-time applications, although it exceeded TextCNN time (0.08 seconds) due to greater computational complexity.

Discussion

Interpretation of Results. The obtained results demonstrate that the proposed hybrid model (BERT+LSTM) significantly outperforms baseline neural network architectures (TextCNN, LSTM, BERT-base) and classical matrix factorization across all key metrics: Precision@5, Recall@5, F1-score, and CTR. The substantial improvement of 43% in F1-score and 66% in CTR compared to the baseline model underscores the effectiveness of integrating semantic text analysis (through BERT) with temporal dependency modeling of user interactions (through LSTM). This synergistic combination addresses the complementary strengths and limitations of each component architecture.

BERT-base demonstrated exceptional capability in extracting semantic features from news texts, leveraging its deep contextual understanding to capture nuanced relationships between articles and user preferences. However, its limitations in accounting for the dynamic nature of user behavior were effectively addressed in the hybrid model through LSTM integration, enabling better adaptation to evolving user preferences and temporal interaction patterns. The LSTM component proved particularly valuable for

modeling sequential user behavior, capturing how user interests shift over time and across different contexts.

TextCNN, despite its computational efficiency and straightforward architecture, proved less accurate due to its limited capacity for processing complex contextual dependencies and long-range semantic relationships in texts. The convolutional approach, while effective for local feature extraction, struggled to capture the global semantic coherence that is crucial for understanding news content relevance. This limitation became particularly evident when dealing with longer articles or when contextual understanding was essential for accurate personalization.

The robustness of the hybrid model to data noise (only 5% accuracy decrease when 10% random interactions were added) indicates its stability, which is critically important for real-world applications where data may be incomplete, noisy, or subject to adversarial manipulation. This robustness stems from the model's ability to rely on multiple information sources: when one component (e.g., user interaction data) becomes unreliable, the other components (e.g., content-based features) can compensate, maintaining overall system performance. The improvement in F1-score by 8% when increasing BERT fine-tuning epochs from 3 to 10 highlights the crucial importance of domain adaptation for pre-trained models. This finding suggests that general-purpose language models require substantial fine-tuning to effectively capture domain-specific nuances in news content, including journalistic writing styles, temporal relevance patterns, and topic-specific terminology.

Comparison with Previous Research

The proposed hybrid model demonstrates substantial improvements over traditional recommender systems through the strategic application of deep learning techniques. Historical benchmarks for collaborative filtering-based models typically achieved Precision@5 scores around 0.65, while our hybrid model reached 0.78, representing a 20% relative improvement. This advancement can be attributed to several key innovations: the sophisticated combination of content-based and collaborative approaches, the integration of rich contextual factors including geographic location and temporal patterns, and the leveraging of state-of-the-art natural language processing capabilities.

Recent studies in news recommendation have explored various neural architectures, but few have successfully combined the semantic understanding capabilities of transformer models with the temporal modeling strengths of recurrent networks. Our approach builds upon foundational work in neural collaborative filtering while incorporating advances in natural language understanding from transformer architectures. The integration of contextual factors represents a significant advancement over previous work that primarily focused on either content similarity or collaborative patterns in isolation.

When compared to other hybrid approaches in the literature, our model's performance gains are particularly notable in handling the cold start problem, where traditional collaborative filtering methods typically struggle. The content-based component of our hybrid architecture provides meaningful recommendations even for users with limited interaction history, while the collaborative component becomes increasingly effective as more user data becomes available.

Practical Applications

The hybrid model presents significant opportunities for deployment across various news and media platforms. Large-scale news aggregators such as Google News, Apple News, or specialized platforms like AllSides could benefit from the model's enhanced personalization capabilities. The system's ability to handle both high-volume established users and newcomers makes it particularly suitable for platforms experiencing rapid user growth or serving diverse global audiences.

Social media platforms integrating news content, such as Facebook, Twitter, or LinkedIn, could leverage the model's contextual awareness to provide more relevant news recommendations based on user location, time of day, and social connections. The temporal modeling capabilities are especially valuable for platforms where news consumption patterns vary significantly throughout the day and across different user demographics.

The model's architecture is sufficiently flexible for adaptation to other content domains beyond news. Video streaming platforms could adapt the approach for movie and series recommendations, while e-commerce platforms could apply similar techniques for product recommendations. The key requirement is the availability of rich textual content descriptions and user interaction data, which are common across many recommendation domains. For content creators and publishers, the model offers insights into user engagement patterns that could inform editorial strategies and content development. The semantic analysis capabilities could help identify trending topics and user interests in real-time, enabling more responsive content creation strategies.

Ethical Considerations

The deployment of neural network models for personalization raises several critical ethical concerns that extend beyond technical performance metrics. The phenomenon of filter bubbles represents a fundamental challenge to democratic discourse and informed citizenship. While our hybrid model partially mitigated this issue through contextual feature integration leading to more diverse recommendations, complete elimination of filter bubbles requires more comprehensive approaches.

The model's tendency to reinforce existing user preferences, while improving short-term engagement metrics, may inadvertently limit exposure to diverse viewpoints and important but less immediately appealing content. This is particularly concerning in news consumption, where exposure to challenging or contradictory information is essential for informed decision-making. Future implementations should consider incorporating explicit diversity mechanisms, such as occasional injection of contrarian viewpoints or coverage of underrepresented topics.

Privacy concerns represent another significant ethical dimension. The model's effectiveness relies on extensive user data collection, including reading patterns, interaction histories, and contextual information such as location and timing. While our experiments utilized anonymized data, real-world deployment requires robust privacy protection mechanisms. Implementation of differential privacy techniques, federated learning approaches, or user-controlled privacy settings could help balance personalization effectiveness with privacy protection. The potential for algorithmic bias presents additional challenges, particularly given that news recommendation systems can influence public opinion and political discourse. Training data may contain historical biases that could be amplified by the model, leading to unfair treatment of certain topics, viewpoints, or demographic groups. Regular bias auditing and fairness-aware machine learning techniques should be integral components of any production deployment.

Transparency and explainability represent growing concerns among users and regulators. The complex nature of hybrid neural architectures makes it challenging for users to understand why specific recommendations are made. Developing explainable AI techniques for recommendation systems could improve user trust and enable more informed interaction with personalized content.

Limitations and Challenges

Several significant limitations constrain the generalizability and practical deployment of the proposed approach. The cold start problem for new articles remains partially unresolved, as the model requires some interaction data to achieve optimal performance. While content-based features provide initial recommendations, the full benefits of the collaborative component only emerge after sufficient user engagement data becomes available. This limitation is particularly problematic for breaking news stories or niche content that may not immediately attract user attention.

The model's training and evaluation on the predominantly English-language MIND dataset limits its immediate applicability to multilingual or non-English news environments. Different languages present unique challenges including varying linguistic structures, cultural context dependencies, and different news consumption patterns. Adaptation to other languages would require substantial additional data collection and model retraining, potentially with language-specific architectural modifications.

Computational complexity represents a practical constraint for deployment in resource-limited environments. The hybrid model's inference time of 0.12 seconds per recommendation, while acceptable for most modern systems, may be prohibitive for high-frequency recommendation scenarios or edge computing deployments. The model's memory requirements for storing BERT embeddings and LSTM states could also present challenges for mobile applications or smaller-scale implementations.

Temporal drift in user preferences and content characteristics poses ongoing challenges for model maintenance. News topics, writing styles, and user engagement patterns evolve continuously, potentially degrading model performance over time. Regular retraining cycles and online learning capabilities would be necessary for production deployments, adding operational complexity and computational costs. The model's effectiveness may vary significantly across different news categories and content types. While the evaluation focused on general news content, specialized domains such as technical journalism, sports coverage, or local news may require domain-specific adaptations or entirely different modeling approaches.

Future Research Directions

Several promising avenues for future research emerge from the current work's findings and limitations. The integration of multimodal data represents a significant opportunity for enhancement. News articles increasingly include rich visual content, including images, infographics, and video content, which could provide additional signals for user preference modeling. Computer vision techniques could be integrated with the existing text-based approach to create more comprehensive content representations.

Advanced techniques for addressing filter bubbles and promoting content diversity warrant further investigation. Reinforcement learning approaches could be employed to optimize long-term user satisfaction while maintaining exposure to diverse content. Multi-objective optimization techniques could balance engagement metrics with diversity measures, potentially improving both user experience and societal outcomes.

The development of more sophisticated contextual awareness capabilities represents another promising direction. Current contextual factors include basic temporal and geographic information, but more nuanced context understanding could incorporate user emotional states, social contexts, current events, and personal circumstances. Integration with IoT devices or social media signals could provide richer contextual information for more precise personalization.

Federated learning approaches could address privacy concerns while maintaining personalization effectiveness. By keeping user data on local devices while sharing only model updates, federated learning could enable personalized recommendations without compromising user privacy. This approach would be particularly valuable for news consumption, where reading patterns may be considered sensitive personal information.

Cross-lingual and cross-cultural adaptation techniques could extend the model's applicability to global news environments. Transfer learning approaches, multilingual transformer models, and cultural adaptation techniques could enable deployment across diverse linguistic and cultural contexts while maintaining personalization effectiveness. The development of real-time adaptive learning capabilities could improve the model's responsiveness to rapidly changing news environments and evolving user preferences. Online learning techniques, incremental model updates, and dynamic architecture adaptation could enable more responsive personalization systems that adapt quickly to new trends and user behavior patterns.

Recommendations for Implementation

For practitioners considering deployment of similar systems, several key recommendations emerge from this research.

First, the importance of comprehensive evaluation metrics cannot be overstated. Beyond traditional accuracy measures, deployment should include assessments of diversity, fairness, and long-term user satisfaction. Regular A/B testing and user feedback collection should be integral components of any production system.

Second, the integration of explainability features should be considered from the design phase rather than as an afterthought. Users increasingly expect to understand why specific recommendations are made, and regulatory environments may require algorithmic transparency. Developing interpretable components within the hybrid architecture could improve user trust and regulatory compliance.

Third, robust data governance and privacy protection mechanisms should be implemented from the outset. This includes not only technical privacy protection measures but also clear user consent mechanisms, data retention policies, and user control over personalization features. The design should enable users to understand and control how their data is used for personalization.

Finally, continuous monitoring and adaptation mechanisms should be built into production systems. News environments and user preferences evolve rapidly, requiring ongoing model maintenance and improvement. Automated monitoring of model performance degradation, bias detection, and user satisfaction metrics should enable proactive system maintenance and improvement.

The research demonstrates the significant potential of hybrid neural architectures for news personalization while highlighting the complex technical, ethical, and practical challenges that must be addressed for successful deployment. Future work should continue to balance personalization effectiveness with broader societal considerations, ensuring that advanced recommendation systems serve both individual user needs and collective democratic values.

Conclusion

This research has demonstrated the high effectiveness of the proposed hybrid neural network model (BERT+LSTM) for news content personalization, representing a significant contribution to intelligent recommendation systems. The model achieved substantial improvements across all key metrics (Precision@5 = 0.78, Recall@5 = 0.75, F1-score = 0.76, CTR = 6.3%) compared to baseline architectures and classical matrix factorization. The integration of semantic text analysis with temporal dependency modeling enabled more relevant and personalized recommendations, particularly under conditions of noisy data and cold start scenarios.

A key achievement was creating an architecture capable of effectively processing both static content characteristics and dynamic user behavior patterns. The hybrid model demonstrated robustness to data noise (only 5% accuracy decrease with 10% random interactions added) and achieved a 66% improvement in CTR, directly translating to enhanced user engagement and platform effectiveness. The results underscore the significant potential of hybrid approaches for news platforms, content aggregators, and social networks with integrated news feeds. The model's ability to adapt to contextual factors such as geolocation, time of day, and interaction history makes it particularly valuable for global platforms with diverse audiences.

However, the research revealed substantial limitations requiring future attention. High computational complexity (0.12 seconds inference time) may limit application in latency-sensitive scenarios, while limited multilingual applicability presents barriers for global deployment. Ethical aspects of personalization, particularly filter bubble mitigation, require comprehensive approaches beyond the partial solutions achieved through contextual features.

Future research directions present numerous promising opportunities. Integration of multimodal data including visual and audio content can enhance recommendation quality as news articles increasingly incorporate rich multimedia elements. Advanced methods for combating filter bubbles through reinforcement learning and active diversification mechanisms remain critically important for ensuring informational diversity. Model optimization for multilingual and cross-cultural scenarios requires specialized transfer learning techniques and cultural adaptation. Lightweight architectures such as DistilBERT offer promising paths for reducing computational costs through knowledge distillation, quantization, and pruning techniques, enabling deployment under resource constraints. Development of explainable personalization methods is increasingly important for algorithmic transparency and user trust. Federated learning approaches

can address privacy concerns while maintaining effectiveness, particularly relevant given growing data confidentiality requirements. Real-time integration and online learning capabilities can ensure relevance in rapidly changing information environments. Cross-domain adaptation could extend applicability beyond news to scientific articles, educational materials, and entertainment content. Privacy-preserving techniques including differential privacy and secure computation address growing privacy concerns while maintaining recommendation quality. Integration of social signals and community features could enhance personalization through peer influences and social context.

This research establishes a solid foundation for next-generation intelligent news personalization systems. The proposed hybrid architecture demonstrates significant advantages over traditional approaches, opening possibilities for more effective, ethical, and user-oriented information platforms. Future research should focus on overcoming identified limitations while developing universal, scalable, and ethically responsible solutions.

Successful implementation of these directions will enable creation of personalization systems that are both technically superior and socially responsible, contributing to informed democratic society through access to relevant, diverse, and quality information. The continued evolution must balance technological advancement with ethical considerations, ensuring personalized news delivery serves both individual preferences and broader societal needs for informed citizenship and democratic discourse. The research demonstrates that sophisticated neural architectures can significantly improve personalization effectiveness while highlighting complex challenges for responsible deployment. As the information landscape evolves, these systems must adapt to serve not only individual user satisfaction but also collective needs for an informed, engaged, and democratically participatory society.

References

Aggarwal, C. C. (2016). *Recommender systems: The textbook*. Springer.

Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.

Burke, R. (2007). Hybrid recommender systems: Survey and experiments. In *User modeling and user-adapted interaction* (pp. 331–370). Springer.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.

He, X., Liao, L., Zhang, H., Nie, L., Hu, X., & Chua, T.-S. (2017). Neural collaborative filtering. *Proceedings of the 26th International Conference on World Wide Web*, – P. 173–182. URL: https://doi.org/10.1145/3038912.3052569

Jurafsky, D., & Martin, J. H. (2021). *Speech and language processing* (3rd ed. draft). Rahmani H., Bennamoun M. (2017). URL: https://web.stanford.edu/~jurafsky/slp3

Kim, Y. (2014). Convolutional neural networks for sentence classification. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, – P. 1746–1751. Rahmani H., Bennamoun M. (2017). URL: https://doi.org/10.3115/v1/D14–1181

Manning, C. D., & Schütze, H. (1999). *Foundations of statistical natural language processing*. MIT Press.

Middleton, S. E., Shadbolt, N. R., & De Roure, D. C. (2004). Ontological user profiling in recommender systems. *ACM Transactions on Information Systems*, –22(1). – P. 54–88. URL: https://doi.org/10.1145/963770.963773

Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). *The ethics of artificial intelligence*. Oxford University Press.

O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.

Ricci, F., Rokach, L., & Shapira, B. (2015). *Recommender systems handbook* (2nd ed.). Springer.

Russell, S., & Norvig, P. (2020). *Artificial intelligence: A modern approach* (4th ed.). Pearson.

Vaswani, A., Shazeer, N., Parmar, N., Uszoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, – 30. – P. 5998–6008.

Zhang, A., Lipton, Z. C., Li, M., & Smola, A. J. (2021). *Dive into deep learning*. Cambridge University Press.

Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, – 17(6). – P. 734–749. URL: https://doi.org/10.1109/TKDE.2005.99

Chollet, F. (2017). *Deep learning with Python*. Manning Publications.

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding.

Goldberg, Y. (2017). *Neural network methods in natural language processing*. Morgan & Claypool Publishers.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer.

Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, – 42(8). – P. 30–37. URL: https://doi.org/10.1109/MC.2009.263

Liu, B. (2015). *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge University Press.

Töscher, A., Jahrer, M., & Bell, R. M. (2009). The BigChaos solution to the Netflix Grand Prize. *Netflix Prize Documentation*.

Wu, F., Qiao, Y., Chen, J.-H., Wu, C., Qi, T., Lian, J., ... & Xie, X. (2020). MIND: A large-scale dataset for news recommendation. *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, – P. 3597–3606. URL: https://doi.org/10.18653/v1/2020.acl-main.331

Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs. (A book discussing privacy and ethical issues in data-driven systems, including content personalization.)

submitted 05.09.2025;
accepted for publication 19.09.2025;
published 30.10.2025
© Shalamov I. S.
Contact: alffanchor@gmail.com