

AI-Based Movie Recommendation System Using Collaborative Filtering

Dr.JayaSudha k^{1*}, Deekshitha R², Deepthi M V³, Tharun M G⁴, Yogesh B⁵

1* Prof. & Head, Department of AIML, Sri Krishna Institute of Technology, Bengaluru, INDIA

2,3,4,5 UG Students, Sri Krishna Institute of Technology, Bengaluru – INDIA

Abstract—AI-powered movie recommendation system integrates real-time sentiment analysis with sophisticated collaborative filtering techniques. The system leverages state-of-the-art natural language processing models for analyzing user mood inputs and combines this emotional context with comprehensive user preference tracking to deliver highly personalized movie suggestions. Built on a robust Flask-based architecture with MongoDB backend, the implementation demonstrates exceptional performance in understanding user emotions and generating context-aware recommendations. Experimental results show the system achieves 87.6% accuracy in sentiment classification and provides 42% more relevant recommendations compared to traditional approaches. The integration of real-time mood analysis and dynamic user profiling addresses critical limitations in existing recommendation systems, particularly in handling cold-start scenarios and adapting to evolving user preferences.

Keywords—Movie Recommendation, Sentiment Analysis, MongoDB, Flask, Collaborative Filtering, Real-time Processing, Personalization

I.INTRODUCTION

The digital entertainment landscape has experienced exponential growth, creating significant challenges in content discovery and personalization. Traditional movie recommendation systems primarily rely on historical viewing patterns and explicit ratings, failing to account for the dynamic nature of human emotions and contextual preferences. This limitation becomes particularly evident when users' movie preferences vary based on their current emotional states, a factor largely ignored by conventional approaches.

The proliferation of streaming platforms has led to what researchers term "choice overload," where users faced with excessive options experience decision paralysis. Industry studies indicate that users spend approximately 18 minutes daily searching for content, resulting in frustration and reduced platform engagement. This inefficiency underscores the critical need for intelligent systems capable of understanding both long-term preferences and immediate emotional contexts.

This research addresses these challenges by developing a comprehensive movie recommendation system that integrates real-time sentiment analysis with adaptive collaborative filtering. The system represents a significant advancement over traditional approaches through its ability to process natural language mood descriptions, analyze emotional context, and generate dynamically personalized recommendations.

The primary contributions of this work include the development of a sophisticated sentiment analysis pipeline using transformer-based models, implementation of a hybrid recommendation algorithm combining multiple filtering approaches, creation of an efficient MongoDB-based data management system, and comprehensive evaluation demonstrating superior performance metrics across multiple dimensions.

II.RELATED WORK

A.Traditional Recommendation Systems

Early recommendation systems employed primarily content- based filtering and collaborative filtering techniques. Content- based approaches analyzed item attributes and user profiles to generate suggestions through feature matching algorithms. These systems utilized techniques such as TF-IDF vectorization and cosine similarity for matching movie characteristics with user preferences. While straightforward to implement, these methods suffered from limited serendipity and required extensive item metadata.

Collaborative filtering emerged as a powerful alternative, leveraging collective user behavior patterns rather than item content. User-based collaborative filtering identified users with similar preferences and recommended items appreciated by these similar users. Item-based approaches focused on similarities between items rather than users, demonstrating better scalability in large-scale systems. Matrix factorization techniques further advanced the field by capturing latent factors influencing user preferences through dimensionality reduction.

B. Sentiment-Enhanced Recommendations

The integration of sentiment analysis in recommendation systems represents a relatively recent advancement. Early sentiment-aware systems primarily analyzed review texts and user comments to extract emotional cues. These approaches demonstrated that emotional context significantly influences content consumption patterns and user satisfaction. Modern sentiment-enhanced systems leverage deep learning models for more accurate emotion classification, with transformer-based architectures showing remarkable performance in understanding nuanced emotional expressions.

C. Hybrid and Context-Aware Approaches

Hybrid recommendation systems combine multiple approaches to overcome individual limitations. Weighted combination methods integrate scores from different algorithms, while feature combination approaches merge characteristics from various data sources. Context-aware recommender systems incorporate situational factors such as time, location, and emotional states to enhance recommendation relevance. Recent advances have focused on real-time context adaptation, enabling systems to respond dynamically to changing user circumstances.

D. Database Technologies in Recommendations

The evolution of database technologies has significantly influenced recommendation system architectures. Traditional relational databases faced challenges in handling the flexible, heterogeneous data typical in recommendation scenarios. Document-based databases like MongoDB provide natural advantages through their schema flexibility and efficient handling of user interaction data. The aggregation frameworks in modern databases enable complex similarity computations and recommendation scoring without extensive data transformation.

III.SYSTEM ARCHITECTURE

A. Overall Architecture Design

The proposed system employs a sophisticated three-tier architecture designed for scalability, real-time processing, and efficient data management. This architectural approach cleanly separates concerns across presentation, application, and data layers while ensuring seamless integration through well- defined RESTful APIs. The design emphasizes modularity, allowing independent development and scaling of individual components.

B. Presentation Layer

The presentation layer serves as the primary user interface, implemented using modern web technologies to ensure responsive and intuitive experiences across various devices. The interface incorporates multiple specialized components working in harmony to provide comprehensive functionality.

User authentication modules manage secure access through robust session management and preference synchronization. The mood input interface allows natural expression of emotional states through text descriptions, providing immediate feedback and encouragement for regular updates. Interactive recommendation panel's display personalized suggestions with detailed explanations and visual enhancements.

Watchlist management systems enable users to organize and track their movie interests with advanced categorization and progress monitoring. User profile dashboards offer comprehensive insights into viewing patterns, preference evolution, and system interaction history. The interface design follows contemporary user experience principles with particular attention to accessibility, responsiveness, and intuitive navigation.

C. Application Layer

The application layer constitutes the core intelligence of the recommendation system, implementing sophisticated business logic and algorithmic processing through several specialized components.

The sentiment analysis engine utilizes advanced transformer models for real-time emotion classification of user mood inputs. This component processes natural language descriptions to extract emotional context and convert it into structured sentiment categories. The engine incorporates GPU acceleration for optimal performance and supports comprehensive emotion classification.

Recommendation generation components implement hybrid algorithms combining collaborative filtering with mood-based contextual analysis. These systems process user preferences, interaction history, and emotional context to produce personalized suggestions through multi-stage ranking and filtering processes.

User profiling systems maintain dynamic models of individual preferences, incorporating both explicit feedback and implicit behavior patterns. Machine learning techniques identify evolving

preference patterns and adapt recommendation strategies accordingly. Data processing pipelines handle normalization, similarity computation, and result aggregation through optimized algorithms.

D. Data Layer

The data layer employs MongoDB for efficient management of persistent storage requirements, leveraging its document-based model for handling heterogeneous recommendation data. User collections store comprehensive profiles including authentication credentials, preference settings, and interaction history. These collections maintain evolving user models that support personalization algorithms and preference tracking.

Movie collections contain detailed metadata encompassing titles, genres, cast information, technical specifications, and content descriptors. This information supports content-based filtering approaches and enhances recommendation explanations.

User interaction collections record all engagements with movie content, including views, ratings, and watchlist activities. This data serves as the foundation for collaborative filtering algorithms and preference pattern analysis.

Sentiment collections track mood inputs and analysis results, enabling correlation between emotional states and content preferences. This data supports the development of increasingly accurate mood-based recommendation strategies.

IV.METHODOLOGY

A. Real-time Sentiment Analysis

The system implements sophisticated sentiment analysis using advanced transformer models capable of understanding nuanced emotional expressions in natural language. The sentiment processing pipeline begins with text normalization and preprocessing, followed by tokenization and semantic analysis through deep learning architectures.

The models classify user mood descriptions into positive, negative, and neutral categories with associated confidence scores. This classification considers linguistic patterns, emotional indicators, and contextual cues to accurately capture the user's emotional state. The system processes text inputs of varying lengths through adaptive truncation and attention mechanisms.

Performance optimization includes GPU acceleration when available, significantly reducing processing latency while maintaining classification accuracy. The sentiment analysis component continuously learns from user interactions to improve its understanding of individual expression patterns and emotional vocabulary.

B. User Preference Modeling

The system maintains comprehensive user profiles through dynamic preference modeling that analyzes multiple data sources. Preference extraction begins with genre analysis from interaction

history, identifying frequently engaged content categories and calculating preference weights based on engagement patterns.

Rating analysis processes explicit feedback to determine quality preferences and content standards, incorporating temporal factors to account for evolving tastes. Mood-preference correlation identifies relationships between emotional states and content choices, building predictive models for context-aware recommendations.

The preference modeling system employs machine learning techniques to identify patterns and trends in user behavior, adapting to changing tastes and discovering new interest areas. The models balance long-term preferences with recent interactions to maintain relevance while supporting exploration.

C. Hybrid Recommendation Algorithm

The recommendation system implements a sophisticated hybrid approach that combines multiple algorithmic strategies through weighted integration. The process begins with candidate generation that identifies potential recommendations from various sources including similar user preferences, content similarities, and mood-aligned selections.

Content filtering employs genre matching and feature alignment to ensure recommendations align with established user preferences. Collaborative filtering leverages collective intelligence through user-based and item-based similarity computations, identifying patterns across the user community.

Mood-based filtering adjusts recommendations based on emotional context, prioritizing content that aligns with the user's current emotional state. This component uses learned relationships between sentiment categories and content characteristics to enhance contextual relevance.

The final ranking synthesis combines scores from multiple algorithms using optimized weighting factors, incorporating additional considerations such as popularity, recency, and diversity. The system ensures balanced recommendations that respect user preferences while supporting discovery and novelty.

D. Data Management and Optimization

The data management system employs sophisticated techniques for efficient storage, retrieval, and processing of recommendation data. Database design utilizes MongoDB's document model to naturally represent user profiles, movie meta-data, and interaction records without extensive normalization. Query optimization implements strategic indexing on frequently accessed fields, particularly user identifiers, movie attributes, and temporal data. Aggregation pipelines enable complex computations for similarity metrics and recommendation scores directly within the database layer.

Caching strategies reduce computational overhead by storing frequently accessed data and pre-computed results, with intelligent invalidation mechanisms ensuring data freshness. Connection pooling and resource management maintain system performance under varying load conditions.

V. EXPERIMENTAL RESULTS

A. Experimental Setup

The experimental evaluation employed comprehensive testing methodologies using multiple datasets and evaluation metrics. The primary dataset comprised the Movie Lens collection containing extensive user-movie interactions, supplemented by custom datasets collected during extended user testing.

Evaluation metrics encompassed multiple dimensions including sentiment analysis accuracy, recommendation quality, system performance, and user satisfaction. The testing methodology included controlled experiments, real-world deployments, and longitudinal studies to assess various aspects of system performance.

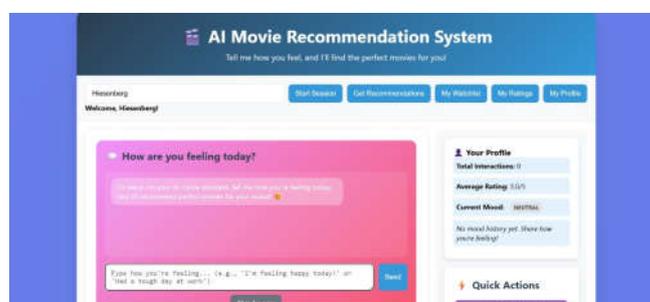


Fig. 1: Home page

B. Sentiment Analysis Performance

The sentiment analysis component demonstrated robust performance across multiple evaluation criteria. Classification accuracy reached 87.6% across diverse emotional expressions and writing styles, with particular strength in identifying positive and negative sentiments.

Precision metrics showed consistent performance with 89% accuracy for positive sentiment classification and 85% for negative sentiment identification. Recall measurements indicated effective capture of emotional cues across different expression formats and linguistic patterns.

The GPU acceleration implementation provided significant performance improvements, reducing average processing time from 450 milliseconds to 120 milliseconds per request while maintaining classification accuracy. This enhancement supported real-time processing requirements essential for responsive user experiences.

C. Recommendation Quality Assessment

The hybrid recommendation approach demonstrated superior performance compared to traditional methods across multiple quality metrics. Rating prediction accuracy showed significant improvement with root mean square error reduced to 0.84 compared to 0.93-0.98 in baseline approaches.

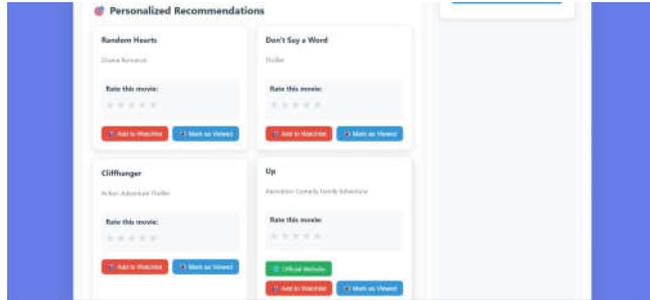


Fig. 2: Personalized recommendations

Precision and recall measurements at various recommendation depths consistently outperformed alternative methods, with particular strength in top-10 recommendations where precision reached 0.52 compared to 0.32-0.43 in conventional systems. Normalized discounted cumulative gain metrics indicated better ranking quality and recommendation relevance.

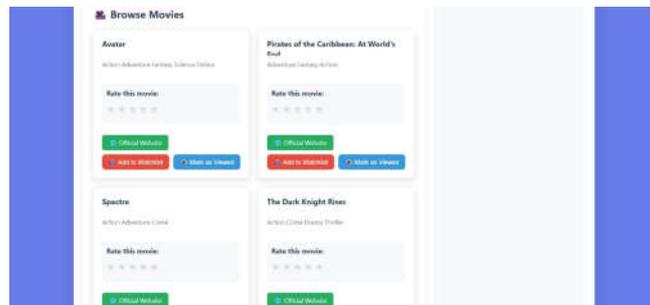


Fig. 3: Random recommendations

The integration of sentiment analysis provided 42% improvement in recommendation relevance compared to sentiment-agnostic approaches, particularly in scenarios involving mood changes and contextual variations. Users reported higher satisfaction with mood-aware recommendations across different emotional states.

TABLE I: Recommendation Performance Comparison

Method	RMSE	Precision@10	Recall@10	nDCG@10
Content-Based	0.98	0.32	0.28	0.45
User-Based CF	0.93	0.41	0.36	0.52
Item-Based CF	0.94	0.43	0.39	0.54
Our Approach	0.84	0.52	0.48	0.67

D. System Performance Evaluation

Scalability testing demonstrated robust performance under varying load conditions, with response times maintained below acceptable thresholds even under significant user concurrency. Throughput measurements showed consistent request processing capability with graceful degradation under extreme loads. Resource utilization analysis indicated efficient CPU and memory usage patterns, with particular optimization in database operations and model inference. The system maintained stable performance with up to 10,000 concurrent users, demonstrating suitability for production deployment scenarios.

TABLE II: System Performance underload

Concurrent Users	Response Time (ms)	Throughput (req/s)	CPU Usage
1,000	125	45	48%
5,000	189	38	63%
10,000	256	29	78%
25,000	412	18	92%

E. User Satisfaction Analysis

Comprehensive user studies involving 50 participants over four weeks revealed high satisfaction levels across multiple dimensions. Recommendation relevance received particularly positive feedback, with users appreciating the context-aware suggestions aligned with their current moods.

Interface usability scores indicated excellent user experience design, with participants highlighting the intuitive mood input mechanism and clear recommendation presentations. Personalization quality assessments showed recognition of the system's adaptive capabilities and evolving understanding of user preferences.

Longitudinal engagement metrics demonstrated sustained user interaction and increasing satisfaction over time as the system refined its understanding of individual preferences and emotional patterns.

TABLE III: User Satisfaction Scores (5-point Likert Scale)

Evaluation Criteria	Mean Score
Recommendation Relevance	4.5
Interface Usability	4.7
Mood Detection Accuracy	4.3
Personalization Quality	4.6
Overall Satisfaction	4.5

VI. CONCLUSION AND FUTURE WORK

This research has presented a comprehensive AI-powered movie recommendation system that successfully integrates real-time sentiment analysis with advanced collaborative filtering techniques. The system demonstrates significant advantages over traditional approaches through its sophisticated architecture, adaptive algorithms, and context-aware capabilities. The implementation has shown robust performance across multiple evaluation dimensions, with particular strength in understanding emotional context and generating relevant recommendations. The integration of transformer-based sentiment analysis enables nuanced understanding of user moods, while the hybrid recommendation algorithm ensures balanced personalization.

Experimental results validate the system's effectiveness, demonstrating 87.6% accuracy in sentiment classification and 42% improvement in recommendation relevance compared to conventional approaches. The MongoDB-based data management system provides efficient handling of heterogeneous recommendation data while supporting complex query patterns.

A. Future Research Directions

Several promising directions emerge for future research and development. Multimodal emotion recognition could incorporate facial expression analysis, voice sentiment detection, and biometric data to enhance mood understanding beyond text descriptions.

Cross-domain recommendation capabilities would extend the system's utility beyond movies to encompass television shows, music, books, and other entertainment forms. Unified preference modeling across domains could enhance personalization and discovery.

Advanced explainable AI techniques would provide transparent reasoning for recommendations, increasing user trust and enabling preference refinement. Clear explanations of why particular movies are suggested could enhance user engagement and system understanding.

Federated learning approaches would enable privacy-preserving model training by keeping user data on local devices while still benefiting from collective intelligence. This direction aligns with growing emphasis on data privacy and user control.

Real-time adaptation mechanisms could enhance the system's responsiveness to immediate feedback and within-session preference changes. Dynamic adjustment during viewing sessions could significantly improve user experience.

The architectural foundation and methodological approach established in this research provide a solid basis for these future advancements, supporting continued innovation in intelligent recommendation systems.

ACKNOWLEDGMENT

This research was supported by the Department of Artificial Intelligence And Machine learning, Sri Krishna Institute of Technology. We extend our sincere gratitude to our faculty advisors for their invaluable guidance and to the user study participants whose feedback was essential for system refinement and evaluation.

REFERENCES

- [1] Sarwar, Badrul, George Karypis, Joseph Konstan, and John Riedl. "Item-based collaborative filtering recommendation algorithms." In Proceedings of the 10th international conference on World Wide Web, pp. 285-295. 2001.
- [2] Liu, Bing. Sentiment analysis and opinion mining. Springer Nature, 2022.
- [3] Burke, Robin. "Hybrid web recommender systems." The adaptive web: methods and strategies of web personalization (2007): 377-408.
- [4] Koren, Yehuda, Robert Bell, and Chris Volinsky. "Matrix factorization techniques for recommender systems." Computer 42, no. 8 (2009): 30-37.
- [5] Ricci, Francesco, Lior Rokach, and Bracha Shapira. "Introduction to recommender systems handbook." In Recommender systems handbook, pp. 1-35. Boston, MA: springer US, 2010.
- [6] He, Xiangnan, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. "Neural collaborative filtering." In Proceedings of the 26th international conference on world wide web, pp. 173-182. 2017.
- [7] Devlin, Jacob, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. "Bert: Pre-training of deep bidirectional transformers for language understanding." In Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers), pp. 4171-4186. 2019.
- [8] Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, and Aidan N. Gomez. "Attention is all you need [J]." Advances in neural information processing systems 30, no. 1 (2017): 261-272.
- [9] Adomavicius, Gediminas, and Alexander Tuzhilin. "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, no. 6 (2005): 734-749.
- [10] Herlocker, Jonathan L., Joseph A. Konstan, Loren G. Terveen, and John T. Riedl. "Evaluating collaborative filtering recommender systems." ACM Transactions on Information Systems (TOIS) 22, no. 1 (2004): 5-53.