Human Systems Engineering and Design (IHSED2025), Vol. 198, 2025, 329–336
https://doi.org/10.54941/ahfe1006786

Context-Aware Product Recommendations Using Weather Data and AI Models

Smajić Alma, Rovis Mia, and Lorencin Ivan

Faculty of Informatics, Juraj Dobrila University of Pula, Zagrebačka 30, 52100 Pula, Croatia

ABSTRACT

Traditional recommender systems generate product suggestions based on user purchase history or collaborative filtering techniques. Although effective under static conditions, dynamic contextual factors, such as weather conditions and geographic location, are frequently overlooked despite their clear influence on consumer behaviour. To overcome these limitations, context-aware recommender systems integrate real-time situational data, including ambient temperature, precipitation, and demographic attributes, into the recommendation process, therefore providing more precise and relevant suggestions. An author-developed framework employs OpenAI's GPT-3.5 Turbo and GPT-4 variants to produce personalized order recommendations. Upon receiving a user-provided location (e.g., "Pula"), current weather data are retrieved from an external API and combined with user profile information to construct contextualized prompts. These prompts are sent directly to the OpenAI API, which returns context-aware recommendations based on the provided inputs. By merging environmental context and user preferences with advanced generative AI, alongside a given product database, recommendations are demonstrated to be substantially more relevant than those produced using traditional methods.

Keywords: CARS (context-aware recommender systems), Generative AI, Contextualized prompts, OpenAI

INTRODUCTION

Recommender systems represent a core component of digital platforms, tasked with suggesting relevant items based on user behavior and interaction history. While such approaches have demonstrated effectiveness in stable settings, dynamic contextual factors, such as, weather conditions, time of day, or geographical location are often not accommodated, despite their substantial influence on user preferences and decision making in real world scenarios. To address this limitation, context aware recommender systems (CARS) have been proposed. These systems enhance traditional recommendation pipelines by integrating situational data into the modeling process, thereby improving the timelines and contextual relevance of suggestions. (Adomavicius & Tuzhilin, 2011). Furthermore, real-time data improves the accuracy and relevance of recommendations. For instance, Yoon and Choi (2023) demonstrate that harnessing weather and demographic information can significantly improve performance in product recommendation tasks (Yoon & Choi, 2023).

Alongside the growing adoption of contextual data, recent advances in attention-based models have opened new directions for recommender systems. Through self-attention, such architectures can capture long-range dependencies across user histories and simultaneously process diverse signals, including browsing patterns, cart contents, and environmental context (Mehdiabadi et al., 2024).

These models, particularly in the form of large language models (LLMs), offer new opportunities for building adaptive, prompt-driven recommendation workflows capable of reasoning over both user preferences and real-time situational inputs.

In this study, we present an author-developed prototype that integrates real-time contextual signals and user profile information into a generative recommendation framework. Upon receiving a user-provided location (e.g., "Pula"), the system sends a request to the external API (OpenWeatherMap) to retrieve current weather data and combines it with a product database, alongside stored user preferences, to build enriched prompts. These prompts are processed by OpenAI's GPT-3.5 Turbo, GPT-4, GPT-4o, and lightweight variants to generate context-aware product suggestions. Final outputs are evaluated based on their alignment with both the environmental context, product recommendation and user intent. By combining environmental context and personalized user information with generative AI capabilities, the system produces product recommendations that are evidently more relevant than those generated by static, behavior-only methods. This work provides a flexible and extensible approach to real-time, context-aware recommendation, showcasing the potential of LLMs in adaptive retail scenarios.

BACKGROUND AND RELATED WORK

CARS extend recommendation approaches by integrating environmental and situational variables into the user item interaction framework. By incorporating these contextual dimensions, CARS can accurately capture the user's circumstances and thereby enhance the pertinence of suggested items. CARS extend traditional recommendation models by incorporating contextual factors like weather and location directly into the interaction between users and items. Through context-aware mechanisms, such as time, location and weather, CARS achieve a deeper understanding of the user's current situation and deliver suggestions that are both relevant and timely. (Mateos and Bellogin, 2025).

Surveys, such as the one by Sloke et al. (2021), provide a comprehensive taxonomy of context models (e.g., contextual pre filtering, post filtering and modelling), highlighting significant gains in prediction accuracy when environmental features are incorporated. Furthermore, these findings highlight the proven value of incorporating weather awareness into recommendation systems (Casillo et al., 2022). showed that incorporating localized and precipitation data into a hybrid matrix factorization model resulted in 12% improvement in click-through rates for fashion retail platforms operating in Germany. Similarly, in their analysis of Netflix's contextual pipeline, (Gómez-Uribe and Hunt, 2016). underscored how temporal and regional context, including local weather anomalies, can boost user engagement by surfacing seasonally relevant titles. However, a recent study by Duncan et al. (2024), introduces the MAWI Rec framework, which exhibits the utility of severe weather information in product recommendation workflows. In their experiments on diverse retail datasets, the MAWI Rec system leverages localized weather alerts, such as high-wind advisories or precipitation forecasts to adjust the ranking of merchandise. What is more, their findings demonstrate statistically significant gains over a state-of-the-art baseline, underscoring the capacity of weather-aware recommendations to better align with consumer needs.

Another line of work focuses on prompt optimization to adapt LLM outputs for recommendation tasks. Liu et al. (2023) introduce RecPrompt, a framework that iteratively refines prompt templates to align GPT-4's generative tendencies with explicit user preference signals in the news domain. By systematically adjusting prompt parameters, such as, context length, keyword emphasis, and instruction framing, RecPrompt achieves improved relevance and coherence in generated news suggestions, demonstrating the feasibility of fine-tuning LLM interactions for domain-specific recommendations. Nonetheless, challenges regarding controllability, reproducibility, and the alignment of generative outputs with long-term user objectives still remain.

METHODS

The context-aware methodologies represented by MAWI Rec's weather-driven enhancements and LLM augmented recommendation paradigms form the foundation for next generation recommender systems.

System Architecture

The developed context-aware recommender system is designed to generate personalized product recommendations by integrating real-time contextual data alongside user's preferences, as shown in Figure 1.

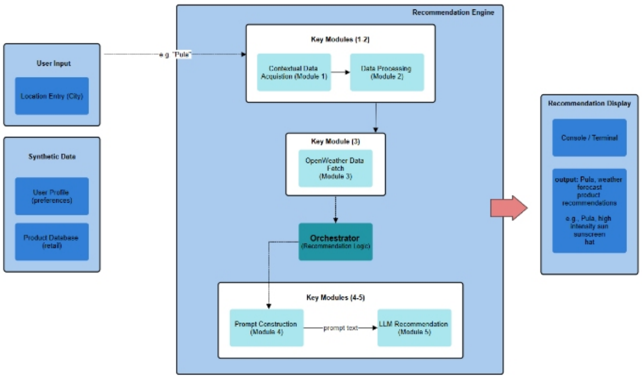

Figure 1: Block diagram based on authors CARS.

The architecture consists of several key modules: data processing, contextual data acquisition and prompt engineering for large language models (LLMs), alongside product recommendation generation.

Data Preparation and Acquisition

Product data is manually created in the form of a dictionary, making it easy to generate and modify synthetic data without external file dependencies. The data is then preprocessed and transformed into dense vector representations (embeddings) using OpenAI's embedding model which enables efficient semantic search and similarity matching. This approach aligns with recent advances in recommender systems that leverage neural embeddings to capture nuanced relationships between items and user needs. Upon receiving a user query, including location, the system sends a request to the *OpenWeatherMap* API to retrieve current weather data for the specified city. The weather data collected includes temperature, humidity, wind speed, atmospheric pressure, and cloud cover. These features are known to significantly influence user preferences and have been shown in literature to improve the accuracy of recommendations when incorporated into the modeling process.

User Profile Integration

Integration of user profile is essential for product recommendation. The profile is constructed to capture preferences, such as dietary restrictions, likes and dislikes in order to generate most credible results. The profile is used to further filter and personalize recommendations, ensuring that the suggestions align with both the user's static preferences and dynamic contextual needs.

Prompt Engineering

The recommendation logic leverages prompt orchestration for a specified LLMs, for the purposes of this study, OpenAI's models. A contextualized prompt is generated combining the user profile, current weather data, and the semantic product database. The prompt explicitly instructs the model to recommend diverse retail products (food, clothes, cosmetics, beverages) relevant to both the weather and user preferences. This method demonstrates recent trends in prompt-driven workflows, where orchestrated prompts guide LLMs through multi-step, context-rich tasks (Mehdiabadi et al., 2024).

Moreover, the system sends the constructed prompt to the OpenAI API and returns the parsed JSON object, which contains weather data for the specified city and three recommended products with descriptions tailored to the weather context and user profile.

Recommendation Generation

Rather than issuing a distinct database query, the system embeds the complete product database directly into the LLM prompt alongside the current weather conditions and user profile. The model then generates its recommendations, leveraging its internal knowledge and the provided context to select and diversify the most appropriate products. This prompt-driven approach simplifies the architecture while still producing highly relevant, context-aware suggestions.

EVALUATION AND RESULTS

Evaluation Strategy

To assess the effectiveness of the developed context-aware recommender system, both qualitative and quantitative evaluation methods were implemented, following best practices in recommender system research (Slokom et al., 2021). The evaluation focused on two main aspects: recommendation relevance and contextual suitability.

Datasets

The evaluation was conducted using a custom-built dataset of retail products, alongside user profiles and real-time weather data collected through the OpenWeatherMap API. User profiles were constructed to reflect diverse preferences, dietary restrictions, and interests, while weather data captured key contextual variables such as temperature, humidity, and precipitation. Moreover, this experimental setup mirrors methodologies found effective in recent context-aware recommendation studies (Casillo et al., 2022).

Comparison

To measure the value of contextual integration, the system's outputs were compared with a conventional recommendation approach that relies solely on user purchase history and collaborative filtering. In addition, such models do not incorporate real-time contextual data or user profile enrichment, which limits their ability to provide personalized and adaptive recommendations. This approach reflects standard industry practices prior to the emergence of context-aware methods, where recommendations were typically based on static historical data without consideration for dynamic user behavior or environmental factors. All in all, while traditional recommender systems rely on static inputs and lack real-time adaptability, the emergence of context-aware and LLM-powered models represents a significant shift towards more intelligent, responsive, and personalized recommendation approaches tailored to continuously customize recommendations according to shifting user inputs and context.

RESULTS

To evaluate the performance and reliability of the developed context-aware recommender system, five OpenAI models GPT-4, GPT-4o, GPT-3.5 Turbo, GPT-4-mini, and GPT-4.1-mini were tested across a series of cities using real-time weather data and a static user profile. The evaluation was based on two primary aspects, structural correctness and contextual alignment. Models were tested one-by-one under identical conditions. In other words, each model was prompted with the same format, and it was expected to return five randomized retail product recommendations formatted as a valid JSON object based on the weather of the specified city. Outputs were manually verified, and the results recorded in a structured Excel table showing binary values of correct (1) and incorrect (0) responses for each model across all tested cities.

As shown in Figure 2, GPT-4o outperformed the other models, returning structurally valid and contextually appropriate responses in 74% of cases. GPT-4o-mini followed with 12%, alongside GPT-4.1-mini with 7%. GPT-3.5-turbo achieved 2%, while GPT-4 failed behind with 0% accuracy. The bar chart in Figure 2 illustrates these results, with the X-axis listing the evaluated models and the Y-axis representing the percentage of accurate responses.

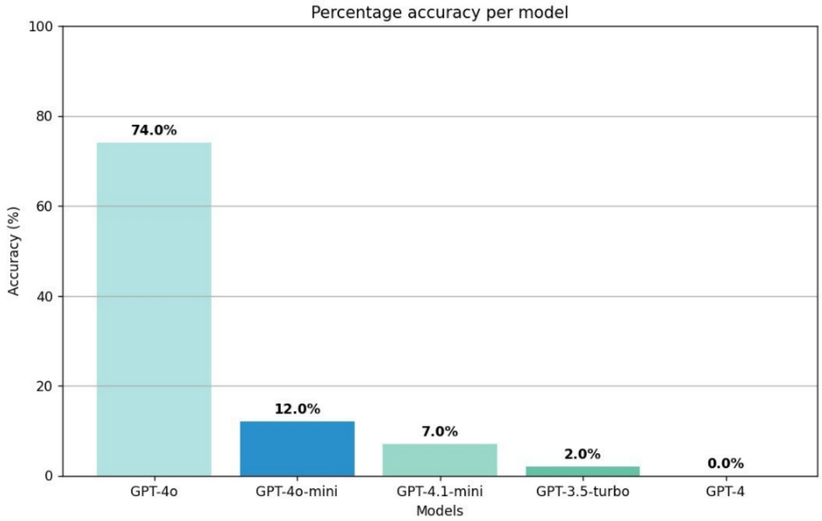

Figure 2: Percentage accuracy per model.

Overall, the results suggest that GPT-4o is significantly more reliable when strict formatting and contextual understanding are required. Its ability to generate semantically relevant retail product suggestions aligned with both the weather data and user preferences marks it as the most suitable model for the developed context-aware recommender system tested in this study.

DISCUSSION

This study highlights the importance of incorporating contextual data into recommendation workflows. The implemented system demonstrated the ability to align product suggestions with environmental conditions and user preferences, confirming findings from prior research that context significantly shapes consumer behavior. In the mentioned study, participants indicated that the recommendations felt more relevant when based on real-time context (Slokom et al., 2021). Furthermore, the use of large language models (LLMs) via prompt engineering proved effective for combining diverse inputs into coherent and personalized outputs. This finding aligns with recent studies emphasizing the generative potential of LLMs in adaptive recommender systems (Mehdiabadi et al., 2024). Nevertheless, despite these strengths, several challenges were identified. The system's performance is highly influenced by the quality of contextual data. Notably, inaccurate or delayed weather information can reduce recommendation relevance leading to inaccurate results.

Future research could be focused on expanding evaluation across diverse demographic segments, incorporating additional contextual signals, such as time of day or special events and investigating hybrid approaches that combine large language models with rule-based reasoning to improve interpretability and system robustness.

ACKNOWLEDGMENT

This research is supported by SPIN project "INFOBIP Konverzacijski Order Management (IP.1.1.03.0120), "Projektiranje i razvoj nove generacije laboratorijskog informacijskog sustava (iLIS)" (IP.1.1.03.0158), "Istraživanje i razvoj inovativnog sustava preporuka za napredno gostoprimstvo u turizmu (InnovateStay)" (IP.1.1.03.0039), "European Digital Innovation Hub Adriatic Croatia (EDIH Adria) (project no. 101083838)" under the European Commission's Digital Europe Programme and the FIPU project "Sustav za modeliranje i provedbu poslovnih procesa u heterogenom i decentraliziranom računalnom sustavu".

REFERENCES

Adomavicius, G., & Tuzhilin, A. (2011). *Context-aware recommender systems*. U F. Ricci, L. Rokach, B. Shapira i P. B. Kantor (Eds.), *Recommender Systems Handbook* (2. izd., str. 217–253).

Braunhofer, M., Elahi, M., & Ricci, F. (2016). Understanding the impact of weather for POI recommendations. *CEUR Workshop Proceedings, 1685*, 19–26.

Casillo, M., Gupta, B. B., Lombardi, M., Lorusso, A., Santaniello, D., & Valentino, C. (2022). Context Aware Recommender Systems: A Novel Approach Based on Matrix Factorization and Contextual Bias. *Electronics, 11*(7), 1003.

Duncan, B. A., Kallumadi, S., & Berg-Kirkpatrick, T. (2024). MAWI Rec: Leveraging severe weather data in recommendation. U *Proceedings of the 18th ACM Conference on Recommender Systems (RecSys '24)*.

Gómez-Uribe, C. A., & Hunt, N. (2015). The Netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems, 6*(4), 1–19.

Liu, Y., & Liu, X. (2023). Is ChatGPT a good recommender? A preliminary study. *Preprint*.

Mateos, P., Bellogín, A. (2025). A systematic literature review of recent advances on context-aware recommender systems. *Artif Intell Rev* 58(20).

Mehdiabadi, A., Akbari, M. K., & Saadatfar, H. (2024). A survey on context-aware recommender systems using deep learning and reinforcement learning. *Artificial Intelligence Review*.

Slokom, M., Hanjalic, A., & Larson, M. (2021). Towards user-oriented privacy for recommender system data: A personalization-based approach to gender obfuscation for user profiles. *Information Processing & Management, 58*(6), 1–24.

Wu, L., Zheng, Z., Qiu, Z., Wang, H., Gu, H., Shen, T., Qin, C., Zhu, C., Zhu, H., Liu, Q., Xiong, H., & Chen, E. (2023). A survey on large language models for recommendation.

Yoon, J., & Choi, J. (2023). Context-aware product recommendation using weather and demographic data. *Journal of Intelligent Information Systems, 62*(2), 145–162.

Zdravevski, E., Lameski, P., & Trajkovik, V. (2023). A survey of context-aware recommender systems for smart health. *Sensors, 23*(7), 3679.