

# EcoAssist: An AI-Powered Knowledge-Based E-Commerce Customer Support Chatbot

Shivansh Mudgal, Ravi Raj Kumar, Ritik Kumar,  
*Department of Computer Science and Engineering (AI&ML)*

*Under the Guidance of Prof. Priyanka Asthana*

*Dronacharya Group of Institutions, Greater Noida, India*

*Affiliated to Dr. A.P.J. Abdul Kalam Technical University (AKTU), Lucknow, India*

Date of Submission: 25-04-2026

Date of Acceptance: 04-05-2026

## Abstract

The rapid expansion of e-commerce platforms has significantly increased the demand for responsive, scalable, and intelligent customer support systems. Traditional customer service models rely heavily on human agents and rule-based automated responders, which often fail to provide contextual, multilingual, and personalized assistance at scale. Recent academic and industrial research demonstrates that AI-driven, knowledge-grounded chatbots can significantly improve response quality, availability, and operational efficiency in digital commerce environments.

This paper presents EcoAssist, an AI-powered e-commerce customer support chatbot that integrates a structured domain knowledge base with a locally deployed Large Language Model (LLM) to generate accurate and context-aware responses. The system is implemented using a Flask-based backend, asynchronous model communication, prompt-engineered knowledge grounding, and a responsive web chat interface. The chatbot architecture incorporates multilingual response capability, response caching, and modular deployment using a local LLM server via Ollama to ensure privacy and low-latency inference.

The proposed design draws from multilingual and semantic chatbot frameworks, context-aware conversational systems, and knowledge-based customer service bots. Experimental evaluation shows that EcoAssist effectively handles frequently asked e-commerce queries, reduces repeated support workload, and improves user interaction quality. The system provides a lightweight yet extensible foundation for AI-driven customer support automation in small and medium-scale e-commerce platforms.

## Keywords

AI Chatbot, E-Commerce Support, Knowledge-Based System, Large Language Model, Multilingual NLP, Prompt Engineering, Customer Automation

## I. Introduction

The global e-commerce ecosystem has experienced rapid growth over the past decade, leading to a proportional increase in customer service demands. Customers expect instant responses regarding order tracking, payments, cancellations, refunds, delivery timelines, and account management. Providing continuous, high-quality customer support using only human agents is expensive, time-consuming, and difficult to scale.

To address these challenges, organizations have increasingly adopted chatbot systems for automated customer interaction. Early chatbot implementations were primarily rule-based and decision-tree driven. While useful for simple queries, such systems lack contextual understanding and fail when user phrasing deviates from predefined patterns. Advances in Natural Language Processing (NLP) and Large Language Models (LLMs) have enabled a new generation of conversational agents capable of generating contextually meaningful responses.

Research shows that AI-powered chatbots improve service availability and reduce operational costs in e-commerce environments. Multilingual and semantic-aware chatbot systems further increase accessibility across diverse user populations. Knowledge-grounded chatbots improve factual reliability by constraining responses using domain data.

The objective of this work is to design and implement EcoAssist, a domain-focused, knowledge-grounded, multilingual-capable chatbot that leverages a local LLM and structured knowledge base to deliver

accurate and context-aware customer support responses through a web interface.

## II. Literature Review

Wibowo et al. demonstrated the practical deployment of chatbots in e-commerce customer service and reported measurable efficiency gains and reduced response times. Their work shows that chatbot-based support systems can successfully handle repetitive queries and reduce human workload. Vinay et al. proposed a multilingual semantic search chatbot framework combining retrieval mechanisms with multilingual NLP, highlighting the importance of semantic grounding and cross-language support.

Wołk et al. focused on multilingual chatbot construction using machine translation and synthetic dialogue data generation, showing that multilingual capability can be achieved even with limited native datasets. Ben Hamida et al. introduced a context-aware multilingual chatbot framework with personalization and adaptive learning, emphasizing contextual memory and dynamic response adaptation.

Knowledge-based chatbot architectures described in Electronic Commerce Research and Applications demonstrate that structured domain knowledge improves factual correctness and reduces hallucinated responses. User behaviour studies by Cheng and Bao indicate that consumer trust and satisfaction depend strongly on chatbot clarity and task appropriateness. Language style and tone also affect engagement and brand perception.

Foundational multilingual conversational evaluation, such as XPersona, supports cross-language personalization strategies. Industrial-scale assistants such as AliMe Assist demonstrate that AI assistants can operate successfully

in real-world e-commerce ecosystems at large scale. These works collectively inform the design principles behind EcoAssist: domain grounding, multilingual support, contextual response generation, and practical deployability.

## References

The development of **EcoAssist** is grounded in existing research regarding chatbot efficiency, multilingual capabilities, and domain-specific knowledge integration.

- **Operational Efficiency and Automation:** Research by Wibowo et al. establishes that chatbots in e-

commerce provide measurable efficiency gains and reduce the workload for human agents by handling repetitive queries.

- **Multilingual and Semantic Grounding:** To address diverse user bases, Vinay et al. proposed a framework that combines retrieval mechanisms with multilingual NLP. Furthermore, Wołk et al. demonstrated that multilingual support can be achieved effectively using machine translation and synthetic data.

- **Contextual Awareness and Personalization:** Modern systems must be context-aware; Ben Hamida et al. introduced a framework focusing on adaptive learning and contextual memory to improve interaction quality.

- **Knowledge-Based Accuracy:** Utilizing a structured knowledge base, as described in *Electronic Commerce Research and Applications*, significantly improves factual correctness and reduces "hallucinations" in AI responses.

- **User Trust and Interaction Style:** Consumer satisfaction is heavily influenced by chatbot clarity and task appropriateness. Additionally,

the language style and tone used by the bot directly impact brand perception and user engagement.

- **Scalability and Evaluation:** Personalization strategies can be evaluated using benchmarks like XPersona, while industrial examples like AliMe Assist prove that AI assistants can function at a large scale in real-world e-commerce ecosystems.

### 1. System Architecture and Methodology

Eco Assist follows a modular layered architecture consisting of a frontend chat interface, a Flask backend server, a structured knowledge base, and a Large Language Model inference layer. The architecture is designed to balance flexibility, accuracy, and deployment simplicity.

The frontend layer provides a browser-based chat interface built with HTML, CSS, and JavaScript. It captures user queries and language preferences and sends them to the backend using asynchronous HTTP requests. The backend layer handles routing, prompt construction, caching, and model communication.

A structured JSON knowledge base stores domain-specific FAQ entries, including business hours, shipping policies, returns, and payment methods. This knowledge base is injected into the model prompt through prompt engineering techniques. This method aligns with knowledge-grounded chatbot approaches described in prior research.

The methodology pipeline consists of query capture, validation, prompt construction with knowledge embedding, asynchronous LLM invocation, optional translation, caching, and response delivery. Asynchronous processing using an event loop and non-blocking HTTP calls improves responsiveness under concurrent usage. Multilingual response handling is inspired by multilingual chatbot frameworks.

Caching is introduced as an optimization layer where repeated queries are

served from memory without invoking the model again, reducing inference cost and latency.

### 2. Implementation

The Eco Assist backend is implemented using the Flask web framework. Two primary routes are defined: a home page route that renders

the chat interface template and a chat API endpoint that processes user queries. The chat endpoint accepts POST requests containing the user query and selected language.

The system integrates with a locally hosted LLM using the Ollama API. Model inference requests are sent using an asynchronous HTTP client library, enabling efficient non-blocking communication. The prompt includes system instructions, the serialized knowledge base, and the user query.

The frontend chat interface dynamically renders messages using JavaScript DOM manipulation. The Fetch API is used to send queries and receive JSON responses. The interface supports keyboard-based submission, button submission, avatar icons, and message styling to improve usability.

A translation module is integrated to produce Hindi responses when selected. This follows multilingual chatbot implementation practices discussed in prior work. A response cache implemented as an in-memory dictionary maps repeated queries to stored responses, significantly improving performance for FAQ-heavy workloads.

### 3. Results and Discussion

Functional testing demonstrates that Eco Assist correctly answers a wide range of e-commerce support queries, including order tracking, cancellation, return policies, payment methods, and delivery timelines. Knowledge-grounded prompting reduces incorrect or fabricated responses compared to unconstrained generation, supporting conclusions from knowledge-based chatbot research.

Performance observations show that first-time responses depend on model inference time, while cached responses are delivered almost instantly. This confirms the effectiveness of caching for repeated queries.

Multilingual output improves accessibility and aligns with user expectations from multilingual chatbot frameworks. User experience observations indicate that natural language responses increase clarity and perceived intelligence. These findings are consistent with consumer behaviour research on chatbot interaction style and trust.

Compared with rule-based chatbots, EcoAssist demonstrates stronger contextual understanding, better phrasing flexibility, and improved handling of paraphrased queries.

#### 4. Challenges and Limitations

The system currently relies on prompt-based knowledge injection rather than semantic retrieval, which limits scalability when the knowledge base grows large. Context memory and personalization features described in advanced frameworks are not yet implemented. Local LLM deployment requires sufficient hardware resources. Translation quality depends on external services and may introduce minor semantic variation. The cache is memory-based and not persistent across restarts.

#### 1. Linguistic Complexity and "Hinglish"

While the system is designed to be multilingual, a major limitation is the handling of mixed-language queries.

- **Nuance Loss:** Users often blend languages (e.g., "Mera order kab deliver hoga?"). Current models may struggle with these semantic nuances of "Hinglish" compared to pure English or pure Hindi.

- **Dialect Variation:** Regional variations in Hindi and other Indian languages can lead to minor semantic variations in the bot's response.

#### 2. Static vs. Dynamic Knowledge

- **Inventory Lag:** Currently, the system relies on a structured JSON knowledge base for FAQ-heavy workloads. However, it lacks real-time integration with live e-commerce databases (like SQL or Shopify) to

provide instant updates on stock levels or specific "out of stock" items.

- **Scalability of Injection:** Prompt-based knowledge injection (putting the whole FAQ into the prompt) is effective for small sets but limits scalability as the knowledge base grows.

#### 3. Technical Constraints of Local Deployment

- **Hardware Dependency:** Since EcoAssist uses a local LLM server via Ollama to ensure privacy, the response speed is strictly tied to the host's hardware (GPU/RAM).

- **Inference Latency:** Without high-end hardware, first-time response times (inference) may be higher than cloud-based alternatives, which can impact user experience during peak traffic.

#### 4. Behavioral and Trust Factors

- **Lack of Personalization:** The current framework does not yet implement contextual memory or personalized user history.

- **Handling Ambiguity:** If a user's query is too vague, the bot may generate a contextually "meaningful" but ultimately unhelpful response if the specific answer isn't in the grounded knowledge base.

## 5. Future Scope

Future enhancements include integration of vector databases and semantic retrieval pipelines following multilingual semantic chatbot frameworks.

Context-aware

personalization and adaptive learning modules can be added based on the frameworks proposed in. Multilingual dataset expansion and dialogue augmentation methods can improve response diversity. Large-scale deployment and monitoring strategies, like industrial assistants, such as AliMe Assist, can be explored. User behaviour modelling and language-style adaptation based on and can further improve engagement and trust.

### Transition to RAG Architecture

The current system utilizes prompt-based knowledge injection, which is limited by the LLM's context window.

- **Vector Databases:** Future iterations will integrate a vector database (such as Pinecone or ChromaDB) to implement **Retrieval-Augmented Generation (RAG)**.
- **Semantic Retrieval:** This will allow the bot to search through thousands of documents rather than reading a static JSON file, ensuring scalability as the business grows.

### Real-Time Data Integration

To provide a truly interactive e-commerce experience, the chatbot needs to move beyond static FAQs.

- **Live Database Connectivity:** Integrating the Flask backend with live SQL databases (e-commerce APIs like Shopify or WooCommerce) will enable the bot to provide real-time updates on **order status, refund progress, and inventory levels**.
- **Personalization:** Implementing session-based memory will allow the bot to recognize returning users and offer personalized recommendations based on past purchase history.

### Multimodal and Multilingual Expansion

Expanding the ways users interact with EcoAssist will improve accessibility across diverse demographics.

- **Voice-to-Text Support:** Adding speech recognition will allow users to query the bot via voice, making the platform more inclusive for users with differing literacy levels.
- **Regional Language Depth:** While current support includes English and Hindi, future versions will incorporate regional languages

such as Tamil, Bengali, and Marathi to capture a larger segment of the Indian e-commerce market.

- **"Hinglish" Optimization:** Fine-tuning models specifically on code-switched datasets will improve the bot's ability to handle a natural linguistic blending common in Indian social commerce.

### Advanced Security and Blockchain Integration

Drawing inspiration from industrial-scale assistants, security will be a primary focus.

- **Blockchain Ledgers:** For high-trust transactions, integrating a blockchain-based ledger can provide unalterable records of customer support promises and refund agreements.
- **Encrypted Sessions:** Implementing end-to-end encryption for chat logs will further ensure data privacy, building on the current local deployment strategy.

## 6. Conclusion

EcoAssist demonstrates that combining a structured knowledge base with a locally deployed Large Language Model can produce an effective and scalable e-commerce customer support chatbot. The architecture integrates knowledge grounding, multilingual capability, asynchronous processing, and caching to balance accuracy and performance. The design is supported by prior research in multilingual, knowledge-based, and AI-driven chatbot systems. The system provides a practical automation solution suitable for real-world deployment and extensible toward semantic retrieval and personalization in future work.

### Key Accomplishments:

- **Knowledge-Grounded Accuracy:** The system effectively utilized prompt engineering to anchor AI responses in a domain-specific JSON knowledge base, significantly reducing factual errors or "hallucinations" common in unconstrained LLMs.
- **Multilingual Accessibility:** Through the integration of translation modules, the platform offers seamless support in both English and Hindi, broadening market reach and inclusivity for a diverse user base.
- **Optimized Performance:** The implementation of a multi-tier architecture—combining asynchronous processing with an in-

memorycachinglayer—  
ensuredthatfrequentlyaskedquestionsare  
answeredwithnear-  
zerolatency,whilecomplexqueriesarehandled  
efficiently by the Ollama inference engine.

- **Scalable Framework:** EcoAssist provides a lightweight and extensible foundation suitable for small and medium enterprises (SMEs) looking to automate their customer service operations.

### References

- [1] Wibowo, B., Clarissa, H., & Suhartono, D., The Application of Chatbot for Customer Service in E-Commerce, EMACS Journal, 2(3), 91–95.
- [2] Vinay R., et al., A multilingual semantic search chatbot framework, IAES Int. Journal of AI, 2024.
- [3] Wołk, A., Skowrońska, H., Skubis, I., Multilingual Chatbot for E-Commerce, PACIS 2021.
- [4] BenHamida, A. G., et al., A Multilingual, Context-Aware E-Commerce Chatbot Framework, ICRDICCT 2025.
- [5] An Intelligent Knowledge-Based Chatbot for Customer Service, Electronic Commerce Research and Applications, Vol. 50, 2021.
- [6] Cheng, X., Bao, Y., Exploring Consumers' Response to Text-Based Chatbots in E-Commerce, arXiv, 2024.
- [7] Chatbots in e-commerce: Language style effects, Journal of Retailing and Consumer Services, 2023.
- [8] Lin, Z., Liu, Z., XPersona: Evaluating Multilingual Personalized Chatbot, arXiv, 2020.
- [9] Li, F., L., AliMeAssist: An Intelligent Assistant for E-Commerce, arXiv, 2018.