INTERNATIONAL JOURNAL OF RESEARCH, EXPLORATION & KNOWLEDGE (IJREK)
Official Scholarly Journal Portal: **https://www.ijrek.org** | DOI: **https://doi.org/10.5281/zenodo.23777596**

OFFICIAL TRACKING SYSTEM IN AI/ML

Ms. Gawand R. R.[1], Bhumika Barge[2]

[1] Faculty Guide & Supervisor | [2] Dattakala group of institution, India

Corresponding author: **bhumikabarge0202@gmail.com**

ARTICLE ID	ARTICLE TYPE	RESEARCH AREA	PUBLICATION
IJREK-2026-00181	Research Article	Artificial Intelligence	Volume 1, Issue Issue 1 (2026)
RECEIVED	**REVISED**	**ACCEPTED**	**PUBLISHED**
29 October 2026	30 October 2026	29 October 2026	02 November 2026

Abstract

Online businesses must answer product questions, explain policies, guide purchases and resolve order-related concerns across multiple channels. Human-only support can be costly and slow during peak demand, while rigid rule-based chatbots struggle with varied wording and multi-step questions. This paper presents a conceptual design for an Artificial Intelligence (AI) Chat Assistant for Online Business that combines natural language understanding, retrieval from approved business knowledge, controlled access to commerce services and escalation to human staff. The proposed assistant interprets customer intent, retrieves relevant product and policy information, generates grounded responses, and invokes permitted actions such as checking order status only after appropriate authentication. The architecture emphasizes factual consistency, privacy, security, accessibility and measurable service quality. It also describes a development methodology, evaluation metrics, risk controls and a staged deployment plan. Since this is a design-oriented study, no original performance results are claimed; the proposed evaluation must be conducted using representative customer queries and verified business data.

Keywords: *Artificial Intelligence; Conversational Agent; Online Business; E-Commerce; Natural Language Processing; Retrieval-Augmented Generation; Customer Support; Personalization*

I. INTRODUCTION

Digital commerce depends not only on product availability and checkout usability but also on the quality of information provided before and after a purchase. Customers may ask about product specifications, stock, delivery estimates, returns, payment methods, discounts, warranty conditions and order status. When the same questions arrive through a website, social channel and email, small businesses can struggle to provide timely and consistent answers. An AI chat assistant can provide a conversational entry point to approved information and reduce repetitive work, while human employees remain responsible for exceptions, disputes and sensitive decisions. Modern language models can interpret flexible language and formulate natural responses, but a fluent answer is not necessarily a correct answer. For a business setting, responses should be grounded in current product records and policy documents rather than relying only on a model's internal knowledge. Tool access must be constrained so that a chat interface cannot make unauthorized changes to orders, expose another customer's data or promise unsupported discounts. The goal is therefore not simply to add a chatbot, but to design a reliable customer-service workflow. Problem Statement and Objectives The problem addressed is the gap between customer expectations for immediate assistance and the limited capacity of many online businesses to provide continuous, consistent support. Objectives are to: (1) design a modular conversational architecture; (2) connect answers to trusted business sources; (3) support product discovery and common service questions; (4) establish safe human handoff; and (5) define measurable quality, privacy and operational controls.

II. BACKGROUND AND RELATED WORK

Conversational agents have evolved from fixed decision trees to systems that combine intent classification, retrieval, language generation and external tools. Rule-based systems are predictable for narrow flows such as collecting a return request, but they require frequent manual updates and can fail when users phrase questions unexpectedly. Intent-based machine-learning systems generalize across wording but may struggle with unseen intents or context that spans multiple turns. Generative language models produce flexible responses, while retrievalaugmented generation (RAG) can connect these responses to a curated knowledge base and provide source snippets for verification. Research on AI-enabled commerce chatbots examines adoption, trust, service quality, customer satisfaction and operational efficiency. Studies in IEEE conference proceedings discuss AI chatbots for e-commerce customer experience and support, while broader reviews identify the importance of context, usability, reliability and user trust [1]–[5]. Findings across this field should not be treated as proof that every deployment improves sales: results depend on task design, business domain, data quality and customer expectations. A business assistant should therefore be evaluated against a baseline and its own service objectives. Limitations in Existing Approaches Common weaknesses include stale product information, unsupported claims about delivery or refunds, weak handling of ambiguous requests, poor memory boundaries, and failure to recognize when a human should take over. A model may also misread a request containing several products or confuse general advice with account-specific information. These risks motivate retrieval grounding, explicit tool permissions, confidence thresholds, conversation-state management and continuous review of unanswered questions.

III. PROPOSED SYSTEM

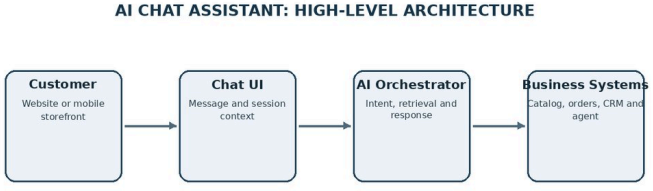


Fig. 1. Proposed architecture connecting the customer interface, AI orchestration and approved business systems.

IV. SYSTEM ARCHITECTURE AND MODULES

A. Customer Interface and Session Manager The interface provides a visible chat launcher, readable message history, suggested prompts and clear controls to end a conversation or request a person. The session manager stores only the context needed to resolve the current request and applies expiration rules. It should avoid placing passwords, payment-card data or unnecessary personal details in prompts and logs. B. Intent and Context Processing The intent component distinguishes product information, product comparison, availability, shipping, returns, order status, payment help and human-agent requests. Entity extraction identifies product names, order references and other parameters. Context tracking supports followup questions such as “Does it come in blue?” without silently carrying unrelated information from an earlier session. If the intent or required parameters are unclear, the assistant asks a concise clarifying question. C. Knowledge Retrieval and Response Generation Business-approved sources may include the product catalogue, frequently asked questions, shipping rules, return policy, warranty documentation and current promotional terms. Content is normalized, versioned and indexed. A retrieval step searches for relevant passages and passes only the needed evidence to the response generator. Answers should be short, actionable and traceable to their sources where appropriate. Product prices and stock status should be read from current systems rather than copied into a static prompt. D. Tool and Integration Layer A tool layer can expose narrowly scoped functions such as search_products, get_public_shipping_policy, check_order_status and create_support_ticket. Each tool validates input, checks authorization, applies rate limits and records a privacy-aware audit event. Read operations should be separated from write operations. Actions such as cancelling an order, issuing a refund or changing an address require explicit confirmation and must follow the business’s normal permission and approval rules. E. Human Handoff and Administration The handoff module activates when the customer asks for a person, retrieval fails, confidence is low, a policy exception is requested or the case involves a dispute. It passes the user’s issue, relevant source references and attempted steps to the agent, subject to consent and retention policy. Administrators need tools to update knowledge, inspect quality metrics, disable problematic responses and manage escalation rules. Module Main responsibility Control Chat interface Collect and display messages Accessible UI; clear AI disclosure Retriever Find relevant approved content Source freshness and relevance checks Response engine Compose grounded answer No unsupported claims; fallback Business API Read permitted records or act Authentication, authorization, validation Human support Resolve complex cases Context summary and audit trail

V. CONVERSATION WORKFLOW

A typical interaction begins with a greeting and a customer question. The system classifies intent, extracts entities and determines whether the request is public or account-specific. It retrieves approved information or calls a permitted API, generates a concise response and checks whether the answer is sufficiently supported. If not, it asks for clarification or escalates. The system then records only necessary operational metadata and updates evaluation logs without retaining sensitive information longer than required.

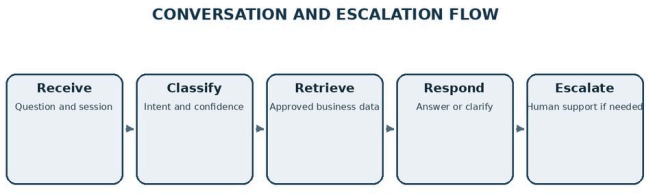


Fig. 2. Conversation processing and escalation workflow.

VI. METHODOLOGY AND IMPLEMENTATION PLAN

The proposed methodology follows an iterative development lifecycle: requirements analysis, knowledge preparation, prototype construction, integration, testing, limited deployment and monitoring. Each phase has measurable acceptance criteria so the system is not released solely because its responses sound natural. A. Requirement Analysis Stakeholders identify customer groups, supported languages, expected query volume, business hours, products and policy boundaries. Queries are grouped by frequency, risk and required system access. The first release should prioritize low-risk, high-frequency questions such as product features, store policies and public shipping information. Sensitive or irreversible actions remain outside the initial autonomous scope. B. Knowledge Preparation Source documents are collected from the business owner and checked for accuracy, ownership, effective date and duplication. Product records should include identifiers, descriptions, variants, price and availability source. Policies should have explicit effective dates and version identifiers. Documents are segmented into meaningful passages and indexed for semantic and keyword retrieval. A content owner must approve updates, and obsolete material must be removed or clearly marked. C. Model and Retrieval Selection A practical implementation may use an NLP model or hosted language-model API for intent understanding and response generation, a vector or hybrid search index for document retrieval, and a backend service to control access to business data. Model choice should consider quality, latency, operating cost, supported languages, privacy terms and deployment constraints. Retrieval quality should be tested independently from answer generation to locate errors precisely. D. Integration and Testing Integration begins with read-only endpoints and test accounts. Functional tests cover product search, policy questions, order lookups, missing parameters and handoff. Security tests attempt unauthorized access, prompt injection, data leakage and unsafe tool invocation. Load tests measure response time under expected concurrent sessions. Human reviewers inspect a representative set of questions, including difficult, ambiguous and out-of-scope examples. Implementation Pipeline

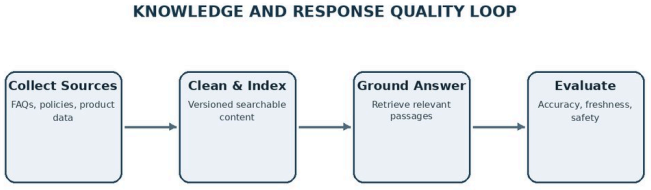


Fig. 3. Knowledge preparation and quality improvement cycle.

VII. FUNCTIONAL AND NON-FUNCTIONAL REQUIREMENTS

TABLE II

Requirement	Example acceptance test	Evidence
Accuracy	Answer matches current approved source	Reviewed query set
Latency	Meet a target agreed before pilot	P50/P95 timing logs
Availability	Fallback works during dependency failure	Failure-injection test
Privacy	No cross-user order disclosure	Authorization tests
Accessibility	Keyboard and screen-reader flows work	Accessibility review
Maintainability	Knowledge update can be rolled back	Release checklist

VIII. EVALUATION FRAMEWORK

Evaluation should combine offline testing, expert review and controlled live monitoring. A curated benchmark set should include common questions, paraphrases, multi-turn conversations, adversarial inputs, policy edge cases and queries with no answer in the knowledge base. The benchmark should be separated from examples used to tune prompts or retrieval, reducing the risk of reporting overly optimistic results. A. Core Metrics Intent classification can be assessed with precision, recall and F1-score. Retrieval can be assessed using Recall@k, mean reciprocal rank or nDCG. Answer quality should be reviewed for factual correctness, completeness, relevance, source support and clarity. Operational measures include first-response latency, resolution rate, escalation rate, API failure rate and cost per conversation. Business outcomes such as conversion or repeat contact should be measured with a controlled comparison and interpreted carefully because seasonality, promotions and traffic mix can affect them. B. Human Evaluation and Error Analysis Reviewers label whether an answer is supported, whether it answers the question, whether it follows the relevant policy and whether the handoff decision was appropriate. Errors are grouped into retrieval miss, stale source, reasoning failure, tool failure, ambiguous request, policy conflict and unsafe disclosure. The team prioritizes fixes by severity

and frequency rather than optimizing only average response time. C. Suggested Pilot Protocol Before launch, define a test set, a baseline support process and acceptance thresholds. Run the assistant in shadow mode or a limited pilot, compare its answers with approved responses, and inspect all high-severity failures. Report sample size, test conditions, confidence intervals where relevant, and limitations. No numerical improvement is asserted in this paper because an actual pilot dataset has not been supplied.

IX. SECURITY, PRIVACY AND RESPONSIBLE AI

An online-business assistant may process names, email addresses, order identifiers, delivery details and customer messages. The design should collect the minimum information required for a task, clearly explain data use, apply access controls and define retention and deletion procedures. Logs should avoid storing secrets and payment details. Where external AI providers are used, the business must review dataprocessing terms, retention settings, regional requirements and whether customer content may be used for provider model improvement. A. Authentication and Authorization A customer must not retrieve account-specific information merely by providing another person's order number. The system should rely on the storefront's authenticated session or a verified process and check authorization at the backend for every request. Tool permissions should be narrowly scoped, credentials stored in a secret manager and service-to-service traffic protected. The language model must never receive unrestricted database credentials or direct production SQL access. B. Prompt Injection and Untrusted Content Customer messages, uploaded documents, product reviews and retrieved web content are untrusted input. Instructions found in those sources must not override system policy or tool permissions. The assistant should treat retrieved content as evidence, not executable instructions. Input validation, output checks, allow-listed tools, restricted parameters and confirmation steps reduce the impact of prompt injection. These measures do not eliminate risk and must be tested continuously. C. Accuracy, Fairness and Transparency The assistant should disclose that it is automated and avoid pretending to be a human employee. It should not infer sensitive personal characteristics or use customer profiling beyond a clearly defined, lawful purpose. Product recommendations should be explainable in terms of stated needs and available product attributes. If a user asks for a guarantee that the business cannot make, the assistant must communicate uncertainty and refer to the authoritative policy or a human agent.

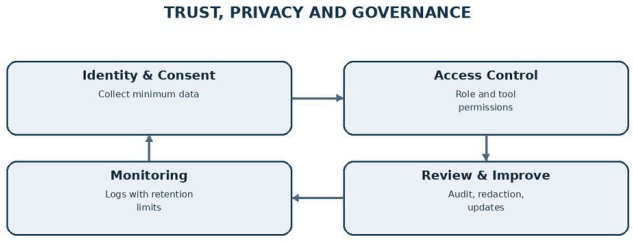


Fig. 4. Governance controls across identity, access, monitoring and review.

TABLE I

Module	Main responsibility	Control
Chat interface	Collect and display messages	Accessible UI; clear AI disclosure
Retriever	Find relevant approved content	Source freshness and relevance checks
Response engine	Compose grounded answer	No unsupported claims; fallback
Business API	Read permitted records or act	Authentication, authorization, validation
Human support	Resolve complex cases	Context summary and audit trail

X. EXPECTED BENEFITS, CHALLENGES AND DISCUSSION

TABLE III

Risk	Impact	Mitigation
Wrong policy answer	Customer dispute or loss of trust	Ground answers in versioned policies; escalate uncertainty
Unauthorized data access	Privacy breach	Backend authorization on every tool call
Stale catalogue	Incorrect price or availability	Live lookup for volatile fields; scheduled sync
Prompt injection	Unsafe action or disclosure	Treat content as untrusted; allow-list tools
API outage	Incomplete support	Timeouts, retries, graceful fallback
Over-automation	Poor handling of exceptions	Easy human handoff and clear scope

XI. FUTURE ENHANCEMENTS

Future versions could support voice interaction, carefully evaluated multilingual conversations, image-based product queries and integration with inventory forecasting. Personalization should be opt-in and based on transparent signals, such as a customer's stated preferences or current browsing context, rather than hidden sensitive inference. Analytics could help identify common unanswered questions, but reporting should aggregate data and respect retention limits. Agent-assist features may be safer initial extensions than allowing the model to autonomously perform high-impact transactions.

XII. CONCLUSION

This paper proposed an AI chat assistant for online business that combines natural-language interaction, retrieval from approved knowledge, restricted backend tools and human escalation. Its value depends on more than conversational fluency: accurate sources, identity checks, safe action boundaries, clear disclosure, measurable evaluation and operational governance are essential. A staged pilot with representative test cases and explicit acceptance criteria is recommended before broad deployment. Future work should implement the architecture in a selected business domain and publish reproducible results from a real evaluation, including limitations and failure analysis.

REFERENCES

[1] “AI-Driven Chatbots for E-Commerce Customer Support,” in Proc. 2024 International Conference on Advances in Computing, Communication and Applied Informatics (ACCAI), 2024, doi: 10.1109/ACCAI61061.2024.10602261. Available: https://ieeexplore.ieee.org/abstract/document/10602261

[2] “Transforming E-Commerce: AI Chatbots for Supercharged Customer Experiences,” in Proc. 2024 International Conference on Information Technology Research and Innovation (ICITRI), 2024, doi: 10.1109/ICITRI62858.2024.10698874. Available: https://ieeexplore.ieee.org/document/10698874/

[3] “Implementation of Retrieval-Augmented Generation (RAG) in Chatbot Systems for Enhanced Real-Time Customer Support in E- Commerce,” in Proc. 2024 3rd International Conference on Automation, Computing and Renewable Systems (ICACRS), 2024, doi: 10.1109/ICACRS62842.2024.10841586. Available: https://ieeexplore.ieee.org/document/10841586/

[4] “Optimizing Artificial Intelligence (AI) Chatbot Customer Service in Small and Medium Enterprises (SMEs) in E-Marketplace,” in Proc. 19th International Conference on Ubiquitous Information Management and Communication (IMCOM), 2025, doi: 10.1109/IMCOM64595.2025.10857522. Available: https://doi.org/10.1109/IMCOM64595.2025.10857522

[5] “The Contribution of Chatbot to Enhanced Customer Satisfaction: A Systematic Review,” in Proc. ASU International Conference on Emerging Technologies for Sustainability and Intelligent Systems (ICETSIS), 2024, doi: 10.1109/ICETSIS61505.2024.10459497. Available: https://doi.org/10.1109/ICETSIS61505.2024.10459497

[6] M. Adam, M. Wessel, and A. Benlian, “AI-based chatbots in customer service and their effects on user compliance,” Electronic Markets, vol. 31, pp. 427–445, 2021, doi: 10.1007/s12525-020-00414-7. Available: https://link.springer.com/article/10.1007/s12525-020-00414-7

[7] M. M. Mariani, N. Hashemi, and J. Wirtz, “Artificial intelligence empowered conversational agents: A systematic literature review and research agenda,” Journal of Business Research, vol. 161, 2023, article 113838, doi: 10.1016/j.jbusres.2023.113838. Available: https://doi.org/10.1016/j.jbusres.2023.113838

[8] J. Balakrishnan and Y. K. Dwivedi, “Conversational commerce: entering the next stage of AI-powered digital assistants,” Annals of Operations Research, vol. 333, pp. 653–687, 2024, doi: 10.1007/s10479-021-04049-5. Available: https://link.springer.com/article/10.1007/s10479-021-04049-5

[9] “The Role of AI Chatbots on E-Commerce Platforms: Understanding its Influence on Customer Trust and Dependability,” in Proc. 9th International Conference on Information Technology and Digital Applications (ICITDA), 2024, doi: 10.1109/ICITDA64560.2024.10809919. Available: https://doi.org/10.1109/ICITDA64560.2024.10809919

[10] “Customer’s Intention to Adopt AI Chatbots in E-Commerce Framework: Using Structural Equation Modeling,” in Proc. 8th International Conference on Business and Information Management (ICBIM), 2024, doi: 10.1109/ICBIM63313.2024.10823613. Available: https://ieeexplore.ieee.org/abstract/document/10823613

Publication archival note: Official scholarly contribution published under CC BY 4.0 International License. Permanent record on **https://www.ijrek.org** | Article ID: IJREK-2026-00181