International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

# RESOLVE X: AI-POWERED AUTONOMOUS CUSTOMER SUPPORT AGENT

[1] K Sunanda, [2] Md Saziya, [3] S Adharsh, [4] A Tejas
[1]AssistantProfessor, [234]Students
Department of IOT
Siddhartha Institute of Technology & Sciences, Narapally
kondapallisunanda.cse@siddhartha.co.in, 23tq1a6925@siddhartha.co.in,
23tq1a6918@siddhartha.co.in, 23tq1a6945@siddhartha.co.in

## Abstract

Customer support is an essential component of modern digital businesses, but traditional support systems often depend heavily on human agents, predefined rules, and static knowledge bases. As the number and complexity of customer queries increase, organizations face challenges such as long response times, inconsistent answers, repeated queries, and high operational costs. **RESOLVE X** is proposed as an AI-powered autonomous customer support agent designed to provide intelligent, context-aware, and automated assistance to customers. The system combines Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), natural language processing, and agent-based decision-making to understand customer requests and generate relevant responses.

The primary objective of RESOLVE X is to move beyond conventional chatbot behavior by enabling the system to understand the intent of a customer, retrieve relevant information, reason over the available context, and determine the appropriate action. Instead of depending only on predefined question-answer pairs, the system can retrieve information from organizational documents, FAQs, product manuals, policies, and knowledge bases. RAG helps ground generated responses in available business information, reducing the possibility of unsupported or hallucinated answers. Recent research has demonstrated the practical value of RAG-based assistants for customer-support applications.

RESOLVE X also introduces autonomous agent capabilities that allow the system to perform multi-step support workflows. The agent can classify customer intent, retrieve relevant information, formulate an answer, use authorized tools or APIs when required, and decide whether a query can be resolved automatically or should be escalated to a human representative. Agentic-RAG research has shown that routing, retrieval validation, and response generation can be combined to improve the quality and relevance of customer-service interactions.

The proposed system is designed with continuous evaluation, monitoring, security, and human escalation as important components. Customer conversations can be analyzed to identify unresolved issues, frequently asked questions, and areas where the knowledge base requires improvement.

# I. Introduction

Customer support has traditionally been implemented through call centers, email communication, FAQs, ticketing systems, and rule-based chatbots. Although these approaches remain useful, they can become inefficient when organizations receive a large volume of customer requests. Customers often expect immediate answers, while support teams need to process multiple conversations simultaneously. Repetitive questions related to products, orders, account services, policies, troubleshooting, and general information consume significant human effort. This creates a need for intelligent systems that can automate routine interactions without sacrificing response quality.

The development of Large Language Models has significantly changed conversational artificial intelligence. Modern language models can understand natural-language questions, maintain conversational context, summarize information, and generate fluent responses. However, relying only on an LLM can introduce problems because the model may not possess the organization's latest information and can sometimes generate inaccurate or unsupported responses. Retrieval-Augmented Generation addresses this limitation by retrieving relevant information from external knowledge sources and providing that information as context during response generation. Production research on RAG-based virtual assistants has demonstrated its usefulness for flexible customer-support interactions.

RESOLVE X extends this concept by incorporating autonomous agent behavior into the customer-support workflow. Rather than simply receiving a question and producing an answer, the system can determine what needs to be done to resolve the customer's problem. For example, it can identify the user's intent, search the knowledge base, verify retrieved information, determine whether a tool or business API is required, generate an appropriate response, and escalate complex cases. This approach transforms a conventional chatbot into a goal-oriented customer-support agent.

Another important aspect of RESOLVE X is contextual personalization. During a conversation, the system can maintain relevant information about the current interaction and use previous messages to understand follow-up questions. For example, if a customer initially asks about an order and later asks about its delivery date, the system can use the conversation context instead of requiring the customer to repeat the order information. Context-aware retrieval and improved grounding are important because factual reliability is particularly significant in customer-service environments.

The proposed system therefore focuses on developing an intelligent customer-support platform capable of handling routine queries autonomously while maintaining appropriate human intervention for sensitive or unresolved situations. RESOLVE X combines natural language understanding, RAG, LLM-based response generation, autonomous decision-making, tool integration, conversation management, and human escalation. The overall goal is to improve response speed, consistency, scalability, and customer experience while reducing unnecessary workload on support teams.

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

## II. Literature Survey

Research in conversational artificial intelligence has progressed from rule-based chatbots toward machine-learning-based dialogue systems and, more recently, LLM-powered agents. Traditional chatbots generally depend on predefined intents, keywords, decision trees, and fixed responses. Although these systems can perform well for predictable questions, they have difficulty handling ambiguous, multi-step, or previously unseen requests. The emergence of LLMs has improved natural-language understanding and generation, allowing customer-support systems to handle more flexible conversations.

Retrieval-Augmented Generation has become an important approach for improving the factual grounding of LLM-based applications. Instead of depending exclusively on information stored within model parameters, RAG retrieves relevant documents from an external knowledge base and provides them to the language model during generation. Research on RAG-based virtual assistants has shown how this architecture can replace rigid rule-based assistants with systems capable of handling a wider range of customer interactions. Engineering studies also emphasize document preparation, retrieval, response generation, and automated evaluation as important components of production RAG systems.

Recent work has investigated RAG specifically for commercial customer support. An industry study on e-commerce support examined document chunking, contextualization, query refinement, and automated LLM-based evaluation as factors affecting RAG quality. Such findings demonstrate that simply connecting an LLM to a vector database is not sufficient; retrieval quality and evaluation must also be carefully designed.

Research has also moved toward **agentic-RAG**, where multiple autonomous components participate in the reasoning and retrieval process. Instead of using a single retrieval-and-generation pipeline, an agentic architecture can perform routing, retrieval validation, response generation, and other specialized tasks. A 2026 study on customer-service interactions reported that agentic-RAG improved response quality, relevance, and accuracy compared with a naïve RAG baseline, particularly for unstructured conversations and ambiguous scenarios.

Another important research direction is the development of self-improving customer-support agents. Recent work on an enterprise support system combines RAG with evolutionary prompt optimization and evaluation workflows so that the system can improve without repeatedly retraining the underlying foundation model. The study emphasizes that enterprise knowledge and policies change continuously, making static support assistants difficult to maintain.

Recent research also highlights the importance of evaluation-driven development for large-scale AI customer-support systems. A 2026 study describing deployments at very large user scale emphasizes context engineering, human-in-the-loop prompt iteration, LLM-based evaluation, and production validation. This indicates that successful customer-support agents require more than a capable language model; they require systematic testing, monitoring, safety controls, and continuous improvement.

## III. System Analysis

RESOLVE X is analyzed as an intelligent customer-support system that receives customer queries through a conversational interface and processes them using multiple AI components. The system first accepts the customer's natural-language request through a web or application interface. A natural language processing layer analyzes the message and identifies the customer's intent, entities, sentiment, and contextual requirements. The query is then passed to an orchestration layer responsible for deciding the appropriate processing workflow. If external organizational information is required, the system performs semantic retrieval from a vector-based knowledge repository. Retrieved information is passed to the LLM together with the relevant conversation context and system instructions. The agent can additionally determine whether an authorized business tool or API needs to be invoked to complete the requested task. The generated response is checked against defined rules and available evidence before being presented to the customer. Complex, sensitive, or unresolved queries can be transferred to a human support representative. Conversation data and feedback can subsequently be used for evaluation and knowledge-base improvement. This architecture enables RESOLVE X to function as an autonomous support system rather than a simple predefined chatbot.

### Existing System

Existing customer-support systems generally consist of human-operated help desks, FAQ pages, ticketing platforms, rule-based chatbots, and basic AI assistants. Human support agents can understand complex situations but require significant time and operational resources, particularly when customer volumes are high. FAQ-based systems provide quick access to predefined information but require customers to search manually and may not answer context-specific questions. Rule-based chatbots generally identify keywords or predefined intents and return fixed responses. These systems work effectively for simple and predictable requests but become less effective when customers use different wording or combine multiple issues in one conversation. Traditional ticketing systems can record unresolved issues but often require human intervention for investigation and resolution. Some modern systems use LLMs, but systems without reliable retrieval mechanisms may produce responses that are not grounded in current organizational information. Existing approaches also commonly separate conversation handling from business tools, limiting their ability to complete tasks autonomously. Consequently, there is a need for an integrated system capable of understanding, retrieving, reasoning, acting, and escalating within one workflow.

### Disadvantages of Existing System

- High dependence on human support representatives.
- Long response and resolution times during peak workloads.
- Limited ability to understand complex natural-language queries.
- Rule-based chatbots depend heavily on predefined intents and responses.
- Difficulty handling multi-step customer problems.
- Repetitive questions increase the workload of support employees.
- Basic LLM chatbots may generate inaccurate or unsupported responses.
- Inadequate automated evaluation and continuous improvement mechanisms.

**Proposed System**

RESOLVE X is proposed as an AI-powered autonomous customer-support agent that integrates LLMs, RAG, natural language processing, autonomous decision-making, and tool integration into a unified platform. The system begins by receiving a customer's query through an interactive conversational interface. An intent-analysis component determines what the customer wants and identifies important entities and contextual information. The orchestration agent then selects the appropriate workflow based on the nature of the request. For knowledge-based questions, the system performs semantic retrieval from a domain-specific vector database containing FAQs, policies, manuals, product information, and organizational documents. The retrieved content is provided to the LLM so that responses are generated using relevant organizational knowledge rather than unsupported assumptions. For action-oriented requests, the agent can invoke authorized tools or APIs, such as order-status services, account systems, ticketing platforms, or scheduling services. The system maintains conversational context to provide coherent follow-up responses. Confidence checks and business rules can be applied before final responses are delivered. When the system cannot safely or confidently resolve a problem, it can escalate the conversation to a human representative along with the relevant conversation history and retrieved information. This combination enables RESOLVE X to provide intelligent, scalable, and controlled customer support.

**Advantages of Proposed System**

- Provides faster automated responses.
- Supports natural-language conversations.
- Uses RAG to ground answers in organizational knowledge.
- Reduces dependence on predefined chatbot rules.
- Maintains conversational context.
- Supports autonomous multi-step workflows.
- Can integrate with business APIs and tools.
- Reduces repetitive workload for human support teams.
- Provides human escalation for complex cases.
- Improves consistency of customer responses.
- Supports centralized knowledge management.
- Enables continuous monitoring and system evaluation.
- Scales more effectively for large customer volumes.
- Can provide personalized responses based on available context.

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

## IV. Methodology

The methodology of RESOLVE X consists of a sequence of interconnected stages designed to convert a customer query into an accurate and actionable response. The first stage is data collection, where organizational documents such as FAQs, product manuals, policies, troubleshooting guides, and support information are gathered. In the second stage, the collected documents are cleaned, structured, divided into meaningful chunks, and converted into vector embeddings. These embeddings are stored in a vector database to enable efficient semantic retrieval. When a customer submits a query, the natural language processing layer analyzes the request and identifies its intent and important entities. The orchestration agent then determines whether the request requires knowledge retrieval, conversation-based reasoning, an external tool, or human escalation. Relevant information is retrieved from the vector database and combined with the customer's current conversational context. The LLM processes this information and generates a response according to system instructions and business rules. For action-oriented queries, the agent can securely invoke authorized APIs or tools and incorporate their results into the response. A validation and safety layer checks the response before it reaches the customer. Finally, conversation outcomes and evaluation metrics can be analyzed to improve prompts, retrieval strategies, workflows, and knowledge resources over time. This methodology follows the broader direction of modern customer-support systems that combine retrieval, agentic workflows, and systematic evaluation.

### System Architecture

The architecture of RESOLVE X follows a modular AI-agent architecture consisting of the Customer Interface, API Gateway, Conversation Manager, AI Orchestrator, Intent and Context Analyzer, RAG Pipeline, Knowledge Base, LLM Engine, Tool/API Layer, Validation Layer, Human Escalation Module, and Monitoring System. The customer interacts with the system through a web or application-based chat interface. The API gateway receives the request and forwards it to the conversation manager, which maintains the relevant session context. The AI orchestrator analyzes the request and determines the appropriate workflow. For knowledge-based questions, the RAG pipeline converts the query into an embedding and retrieves relevant information from the vector database. The retrieved documents are combined with the conversation context and passed to the LLM for grounded response generation. For transactional or action-oriented requests, the orchestrator can communicate with authorized external APIs and business services through a controlled tool layer. A validation layer checks generated responses against business rules, retrieved evidence, and safety requirements. If the system lacks sufficient confidence or encounters a complex case, the human escalation module transfers the conversation to a support representative with the required context. Finally, the monitoring and analytics layer records system performance, retrieval quality, resolution rates, escalation patterns, and user feedback. This modular architecture allows RESOLVE X to combine autonomous reasoning with controlled execution and human oversight.

International Journal of
DATA SCIENCE AND IOT MANAGEMENT SYSTEM
Peer Reviewed, Referred & Indexed Journal
ISSN: 3068-272X www.ijdim.com Original Research Paper

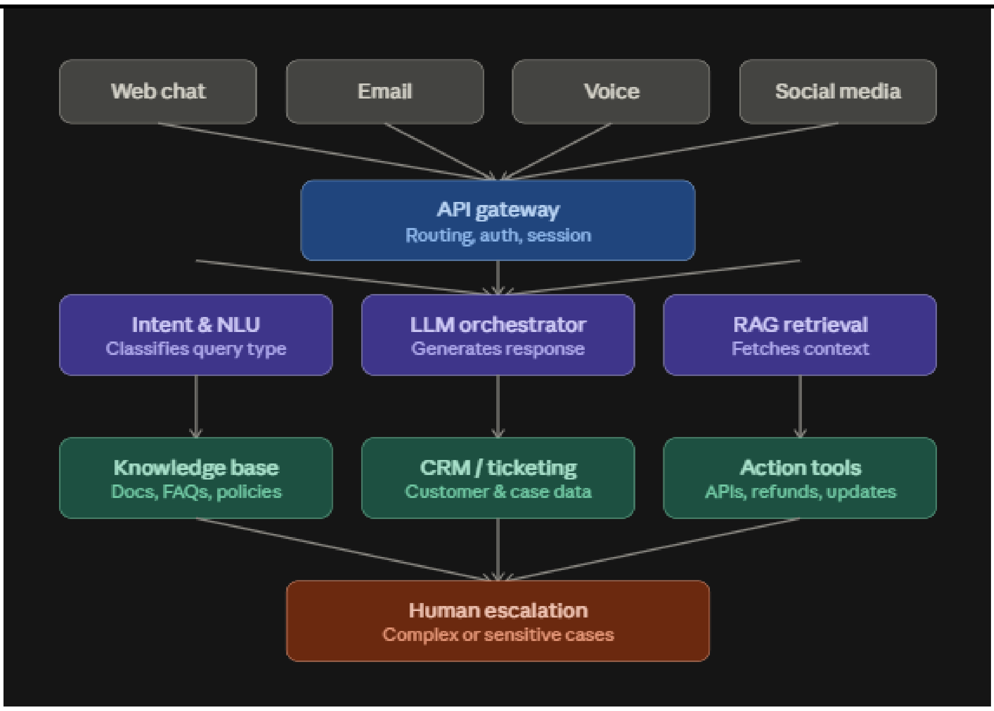


## V. Result and Output

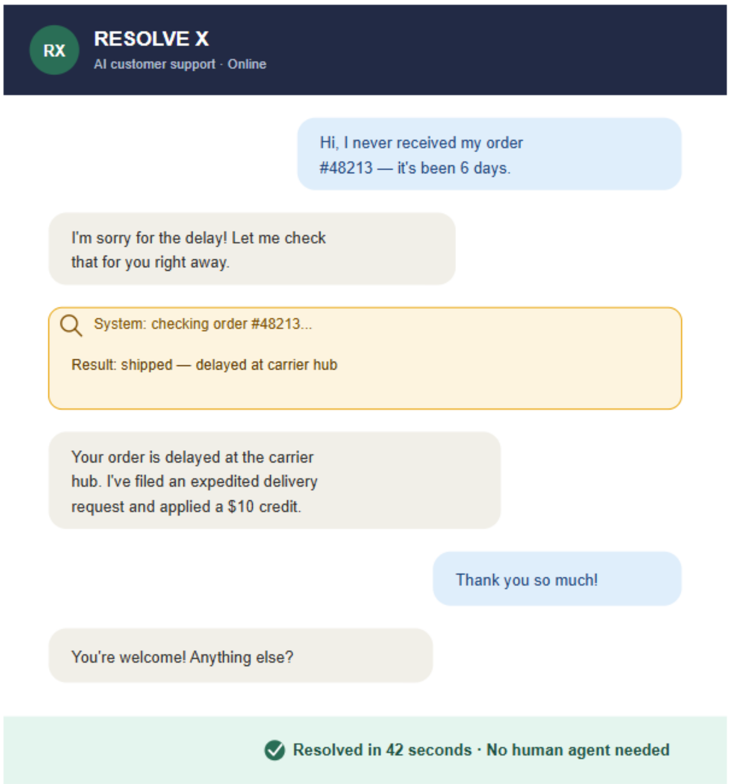

## VI. Conclusion

RESOLVE X: AI-POWERED AUTONOMOUS CUSTOMER SUPPORT AGENT provides an intelligent and automated solution for handling customer queries efficiently. The system uses Artificial Intelligence, Natural Language Processing, Large Language Models, sentiment analysis, and knowledge retrieval techniques to understand customer requirements and generate relevant responses. By maintaining conversation context, the agent can provide personalized and continuous support without requiring manual intervention for every request.

The proposed system can automatically identify customer intent, retrieve appropriate information from the knowledge base, execute supported actions, and provide real-time responses. It can also recognize complex or sensitive queries and escalate them to human support agents when necessary. This combination of automation and human intervention helps improve customer satisfaction while reducing the workload of support teams.

RESOLVE X can significantly reduce response time, improve consistency in customer communication, and provide 24/7 support. Integration with CRM systems, ticketing platforms, databases, and external business APIs further increases its practical usefulness. The system can also collect conversation data and performance metrics, allowing organizations to continuously evaluate and improve their customer-support operations.

Overall, RESOLVE X demonstrates how autonomous AI agents can transform traditional customer-support processes into intelligent, scalable, and responsive services. With future enhancements such as advanced multi-agent collaboration, multilingual support, voice-based interaction, predictive customer assistance, and improved learning mechanisms, the system can become a powerful platform for next-generation automated customer service.

## References

[1] Kumar, R. D., Prudhviraj, G., Vijay, K., Kumar, P. S., & Plugmann, P. (2024). Exploring COVID-19 through intensive investigation with supervised machine learning algorithm. In Handbook of Artificial Intelligence and Wearables (pp. 145-158). CRC Press.
[2] Swathi, B., Vijay, K., Sushanth Babu, M., & Dinesh Kumar, R. (2024, November). Machine Learning Techniques in Cloud Based Intrusion Detection. In The International Conference on Artificial Intelligence and Smart Environment (pp. 557-564). Cham: Springer Nature Switzerland.
[3] Sv satyakrishna, shirisha rangu ,bhargavi nalacheruve.(2024) Prospective investigation on colorectal cancer with SMOTE on machine learning Algorithm
[4] Dr.G.Vishnu Murthy, BhargaviNalacheruve 1Professor, Department of computer Science & engineering, Anurag University, TS, India. 2Student, Department of computer Science & engineering, Anurag University, TS, India.
[5] V. N. S. Manaswini, K. K, C. Nigam, S. S. Ali, R. Niranjana, and Suman, "Real-Time Object Detection in Drone Surveillance Using YOLOv5," in Proc. 2025 3rd Int.

Conf. IoT, Communication and Automation Technology (ICICAT), Gorakhpur, India, 2025, pp. 1–6, doi: 10.1109/ICICAT68430.2025.11414670.

[6] B. Soundarya, V. N. S. Manaswini, M. Ayyakrishnan, R. D. Kumar, "Contextual Analysis of Big Data Analytics in Intelligent Transportation Frameworks," in Intersection of Artificial Intelligence, Data Science, and Cutting-Edge Technologies: From Concepts to Applications in Smart Environment, Lecture Notes in Networks and Systems, vol. 1353, Cham: Springer, 2025, doi: 10.1007/978-3-031-88304-0_79.

[7] R. D. Kumar, V. N. S.Manaswini, "Applications of blockchain in smart cities: detecting fake documents from land records using blockchain technology," in Blockchain for Smart Cities, Elsevier, 2021, pp. 105–117, doi: 10.1016/B978-0-12-824446-3.00017-X.

[8] Tejavath Veeramma, Badarla Anil, Guguloth Ravinder, "An advanced movie recommender using collaborative filtering and sentiment analysis," International Research Journal of Modernization in Engineering Technology and Science, vol. 7, no. 7, July 2025, doi: 10.56726/IRJMETS81618.

[9] Ravi Kumar Banoth, Ramana Murthy B V, "Automatic crop recommendation system using LightGBM and decision tree machine learning models," Journal of Machine and Computing, vol. 5, no. 1, pp. 343, Jan. 2025, doi: 10.53759/7669/jmc202505026.

[10] Ravi Kumar Banoth, Dr. B.V. Ramana Murthy, "Smart agriculture through IoT and machine learning for analyzing carbon footprints," in Proc. Int. Conf. Computer Science and Communication Engineering (ICCSCE), Apr. 2025.

[11] Ravi Kumar Banoth, B. V. Ramana Murthy, "Soil image classification using transfer learning approach: MobileNetV2 with CNN," SN Computer Science, vol. 5, art. no. 199, 2024, doi: 10.1007/s42979-023-02500-x.