

# AI-Powered Help Desk and Support Ticket Management System Using Local LLM Inference

Prof. Ashwini Wakodikar<sup>1</sup> and Mr. Rohit R. Pillewan<sup>2</sup>

Assistant Professor, Department of Computer Application, KDK College of Engineering, Nagpur, India<sup>1</sup>  
PG Scholar, IV Sem MCA, Department of Computer Application, KDK College of Engineering, Nagpur, India<sup>2</sup>  
ashwiniwakodikar@kdkce.edu.in and pillewanrohit.mca24@kdkce.edu.in

**Abstract:** *This paper presents the design and implementation of an AI-Powered Help Desk and Support Ticket Management System, an enterprise-grade conversational application that automates IT support operations through a locally deployed Large Language Model (LLM). The system integrates a dedicated AI assistant named Liza, built using Spring AI 1.0.3 with an Ollama-hosted LLaMA 3 8B model, eliminating dependency on cloud-based AI APIs. Users interact exclusively through natural language — Liza intelligently creates, updates, escalates, or closes support tickets without any form-based interaction. The backend is a Spring Boot 3.x REST API with JWT-based stateless authentication and MySQL persistence. The frontend is a React 18 + Vite 5 Single Page Application with a dark-themed AI interface and real-time ticket dashboard. A dual-strategy intent detection approach — combining structured prompt engineering with temperature-controlled LLM output and a backend rule engine — achieves reliable ticket lifecycle management through multi-turn conversation with per-user JDBC-backed memory isolation. End-to-end experimental results confirm accurate intent classification across CREATE, UPDATE, ESCALATE, and CLOSE operations with zero cloud API cost*

**Keywords:** Spring Boot, Spring AI, Ollama, LLaMA 3, Natural Language Processing, JWT Authentication, React, MySQL, Ticket Management, Conversational AI, Intent Detection, Prompt Engineering, Local LLM

## I. INTRODUCTION

In modern enterprise environments, IT help desk operations are critical for organizational productivity. Traditional ticketing platforms such as Zendesk, Freshdesk, and ServiceNow require users to navigate complex forms, recall ticket identifiers, and engage with non-intuitive workflows — barriers that lead to slow resolution, reduced adoption, and high operational overhead.

Conversational AI offers a fundamentally different paradigm: users simply describe their problem in plain language, and the system handles the rest. However, most AI-augmented help desk solutions depend on proprietary cloud LLMs (e.g., GPT-4), incurring per-token costs, introducing data privacy risks, and requiring internet connectivity — constraints unsuitable for enterprise or offline deployments.

This paper presents a production-ready, fully offline AI help desk system that overcomes these limitations. The system's AI assistant, Liza, is powered by LLaMA 3 8B running locally via Ollama. A structured prompt engineering approach forces deterministic ACTION-block output from the LLM, which is parsed by a backend rule engine to execute the correct database operation. Per-user conversation memory using JdbcChatMemoryRepository ensures multi-turn context continuity without cross-user contamination.

The contributions of this work are: (1) a practical local LLM integration architecture for enterprise support systems; (2) a dual-strategy intent detection approach combining LLM output parsing and rule-based confirmation logic; (3) a complete ticket lifecycle managed entirely through natural language; and (4) a validated full-stack implementation demonstrating zero cloud AI dependency.



## II. RELATED WORK

The domain of AI-driven ticket management has seen increasing academic attention. Liu et al. [1] demonstrated that BERT-based transformer models achieve a 23% improvement in intent classification accuracy over rule-based keyword matching, but their work remained cloud-dependent with no ticket lifecycle integration. Xu et al. [2] established that stateless chatbots fail to resolve 40% of multi-step support interactions correctly, motivating the JDBC-backed conversation memory architecture adopted in this system.

Rashid and Ahmad [3] implemented a BERT + SVM hybrid for automatic ticket classification, achieving 89% accuracy. However, their system only automated post-submission classification — full lifecycle management through conversational interaction was not addressed. Touvron et al. [4] introduced the LLaMA model family, proving that the 8B parameter variant matches GPT-3.5 performance on instruction-following tasks, establishing the academic basis for offline LLM inference in enterprise settings.

Table I summarizes these prior works in comparison with the proposed system.

**Table I: Comparison of Related Works**

Paper / System	Authors	Advantages	Disadvantages	Relevance to Liza
BERT-Based Intent Detection	Liu et al. (2019)	High intent accuracy; handles NL variation	Cloud dependent; no ticket lifecycle	Foundation for intent classification in Liza
Dialogue State Tracking	Xu et al. (2020)	Context preservation; improved resolution	No local inference; no REST integration	Motivation for JdbcChatMemoryRepository
Automating IT Ticket Classification	Rashid & Ahmad (2021)	89% classification accuracy; auto priority	Only classifies; doesn't create tickets	Automated triage replicated in AIService
LLaMA Foundation LLMs	Touvron et al. (2023)	Local inference; matches GPT-3.5	No application layer; no help desk	Core AI engine powering Liza via Ollama
Zendesk AI Help Desk	Zendesk Inc.	Mature product; full lifecycle; AI suggestions	High cost; cloud-only; no local LLM	Competitive benchmark replaced at zero cost

## III. SYSTEM ARCHITECTURE

The system follows a clean three-tier architecture: a Presentation Layer (React 18 + Vite 5 SPA), an Application Layer (Spring Boot 3.x REST API with Spring AI), and a Data Layer (MySQL 8.x). Fig. 1 illustrates the overall enhanced system architecture flow.



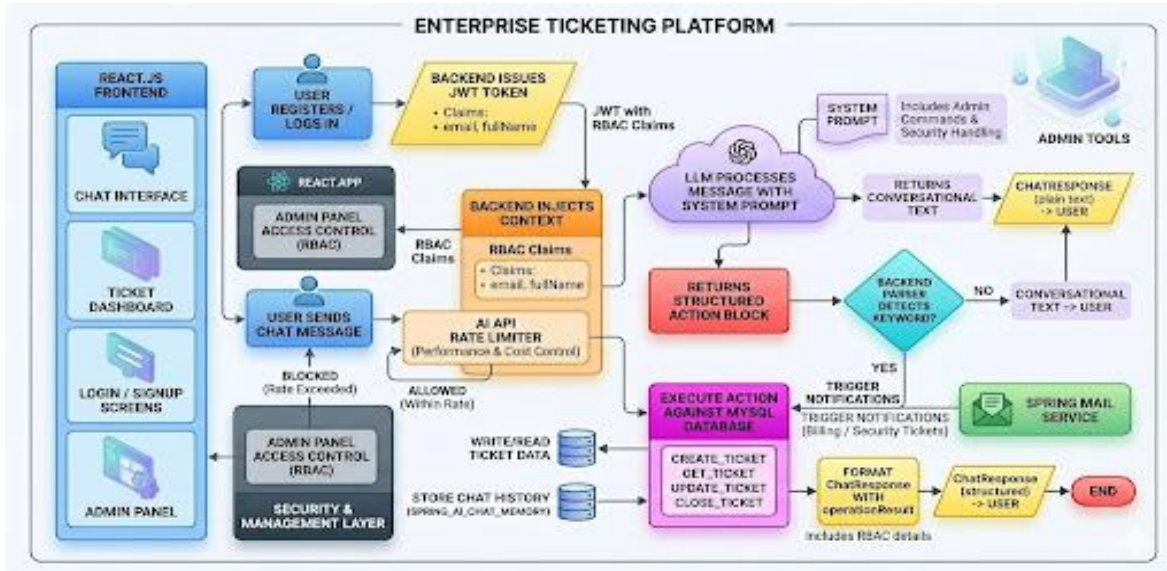


Fig. 1. Enhanced System Architecture Flow of AI-Powered Help Desk

**A. Presentation Layer**

The React SPA communicates with the backend exclusively via REST API calls with JWT Bearer tokens. The UI provides four views: a marketing landing page, a tabbed authentication screen, a full-screen AI conversational chat interface, and a real-time ticket management dashboard. State management uses React Context API for authentication and React hooks for component state.

**B. Application Layer**

The Spring Boot backend is structured across four functional layers. The Controller Layer (AiController, TicketController, AuthController) routes HTTP requests. The Service Layer encapsulates business logic — AIService orchestrates the full LLM pipeline, TicketService manages persistence, and AuthService handles JWT generation. The Repository Layer uses Spring Data JPA with custom query methods. The Security Layer employs a custom JwtAuthFilter with BCrypt password encoding and stateless session policy.

**C. Data Layer**

MySQL 8.x maintains three tables: *users* (credentials with BCrypt-hashed passwords), *tickets* (full lifecycle data with audit annotations), and *SPRING\_AI\_CHAT\_MEMORY* (JDBC-persisted per-user conversation history managed by Spring AI). The ticket description field acts as an immutable audit log — every update is appended as a [UPDATE] or [RESOLUTION] annotation.

**IV. PROPOSED METHODOLOGY**

**A. AI Processing Pipeline**

The conversational AI pipeline transforms raw user text into database actions in five stages: (1) the AiController extracts the user's JWT principal and constructs a conversation ID as 'helpdesk-[email]'; (2) a context block containing the user's full name and verified email is silently appended to the query; (3) the JdbcChatMemoryRepository retrieves the last few conversation turns from MySQL; (4) the enriched prompt is sent to LLaMA 3 via Ollama at temperature 0.1; and (5) the AIService regex-based parser extracts the ACTION keyword and payload fields.



### B. Structured Prompt Engineering

The system prompt (helpdesk-system.st) is the most critical component. It instructs LLaMA 3 to always respond in a structured ACTION block format for ticket operations. For example, a CREATE\_TICKET response begins with the keyword CREATE\_TICKET, followed by colon-delimited key-value pairs for description, priority, category, and tags, and terminates with a USER\_MESSAGE field. A temperature of 0.1 suppresses creative variation, ensuring consistently parseable output across six defined actions: CREATE\_TICKET, GET\_TICKET, UPDATE\_TICKET, ESCALATE\_TICKET, CLOSE\_TICKET, and SEND\_EMAIL.

### C. Dual-Strategy Intent Detection

The AIService applies three sequential detection strategies to each LLM response: (1) regex pattern matching for ACTION keywords at the start of a line; (2) confirmation phrase detection (e.g., 'shall I go ahead', 'yes / no') to set requiresConfirmation: true; and (3) conversational fallback for plain text responses. A 'Confirmation Guard' prevents accidental ticket creation — when a CREATE\_TICKET intent is detected, the system first requests confirmation, and only proceeds when the user responds with a recognized affirmative expression.

### D. JWT-Based Security Architecture

All API endpoints except /api/v1/auth/\*\* are protected by a custom JwtAuthFilter extending OncePerRequestFilter. The filter validates HMAC-SHA256 (HS256) token signatures, checks expiry, and injects the AuthenticatedUser principal into the Spring Security context. This stateless model eliminates server-side sessions and scales horizontally. Ticket ownership is enforced at the service layer — users can only access their own tickets, with cross-user access rejected as 403 FORBIDDEN.

### E. Technology Stack

**Table II: Technology Stack Summary**

Component	Technology	Purpose / Notes
Backend Framework	Spring Boot 3.x	Java 17+ REST API with layered architecture
AI Integration	Spring AI 1.0.3	ChatClient, advisors, memory abstraction
LLM Runtime	Ollama + LLaMA 3 8B	Local inference, zero cloud cost
Database	MySQL 8.x	Ticket, user, and chat memory persistence
Authentication	JWT (HS256, 24hr expiry)	Stateless, scalable, role-aware
Frontend	React 18 + Vite 5 SPA	Dark-themed AI UI, real-time dashboard
Password Hashing	BCrypt (strength 10)	Enterprise-grade credential security
API Testing	Postman	Manual + automated endpoint validation

## V. EXPERIMENTAL SETUP AND RESULTS

### A. Development Environment

The experimental setup required three concurrently running services: the Ollama LLM server (LLaMA 3 8B at http://localhost:11434), the Spring Boot backend (port 8081, launched via 'mvn spring-boot:run'), and the React frontend (port 3000, via 'npm run dev'). Hardware configuration included 16GB RAM and an Apple Silicon M-series processor with Metal GPU acceleration for model inference.

### B. Authentication and Registration

User registration and login were validated via Postman. The system correctly returned 409 CONFLICT for duplicate emails, 201 CREATED with a JWT token for successful registration, and 401 UNAUTHORIZED for invalid login credentials. The issued JWT was decoded and verified to contain sub (email), fullName, iat, and exp claims.



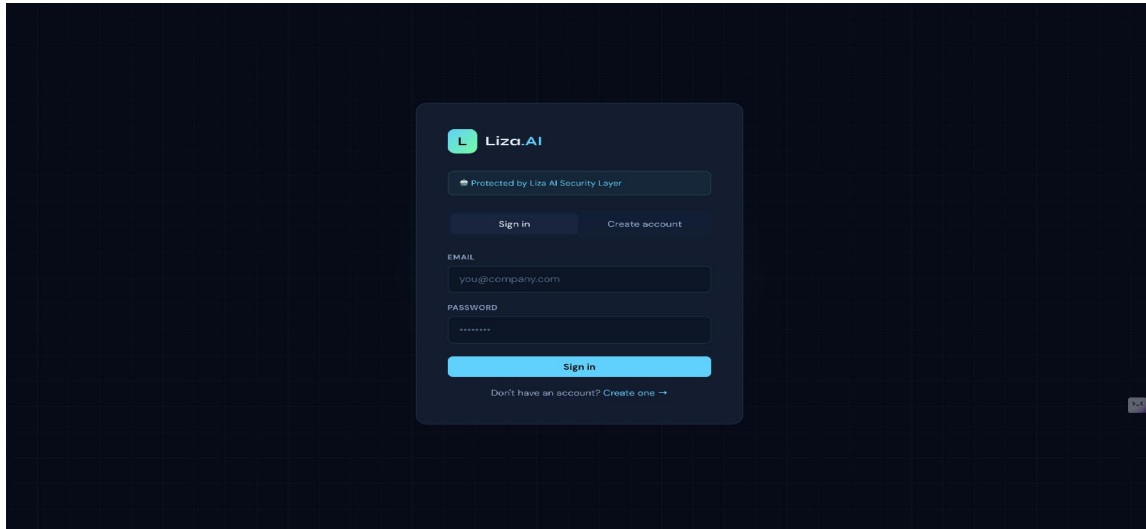


Fig. 4. User Registration Response — JWT Token Issued Successfully

### C. Ticket Lifecycle via Conversational Interface

The complete ticket lifecycle was validated through a structured six-turn conversation. Table III documents the interaction flow and system response at each step.

Table III: Ticket Lifecycle Interaction and Intent Classification

#	User Query	Detected Intent	System Result
1	"Hi"	NONE (conversational)	Personalized greeting with suggestion chips
2	"Laptop is not working"	NONE (info gathering)	AI asks for more details
3	"Please create the ticket"	CREATE_TICKET	Ticket OPEN, priority HIGH, category hardware
4	"Update, battery draining"	UPDATE_TICKET	[UPDATE] annotation appended to description
5	"Issue resolved, close ticket"	CLOSE_TICKET	Status CLOSED, [RESOLUTION] appended
6	"Check ticket status"	GET_TICKET	Inline ticket card rendered in chat bubble



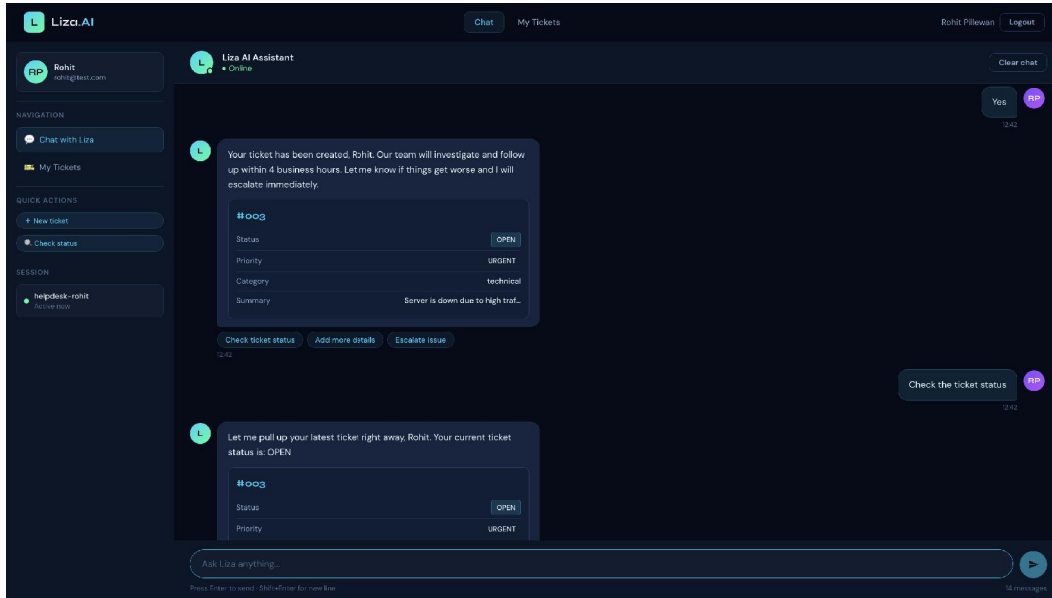


Fig. 5. Chat Interface — Ticket Creation via Natural Language Conversation

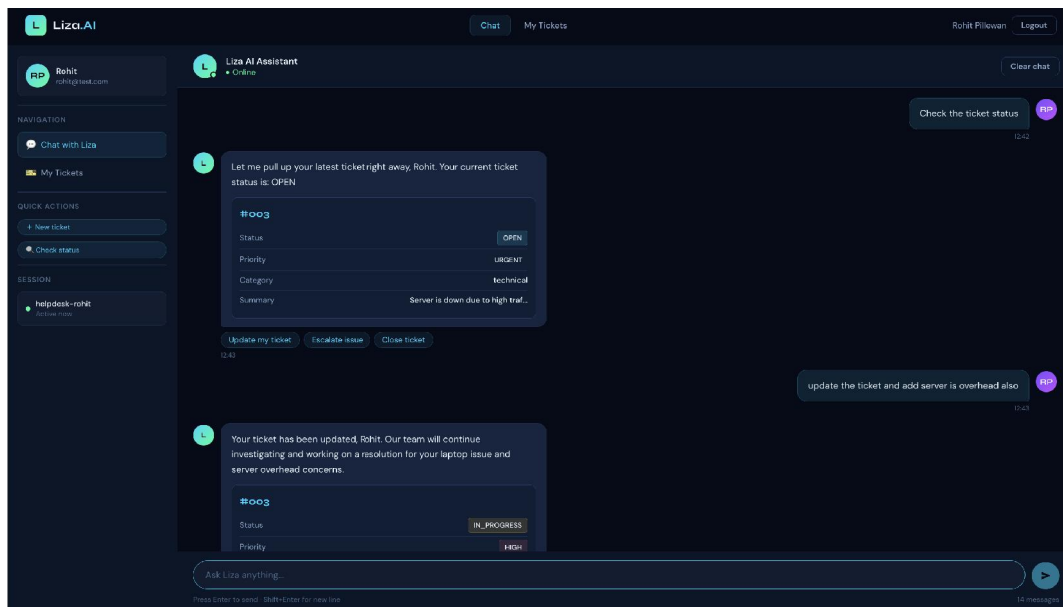
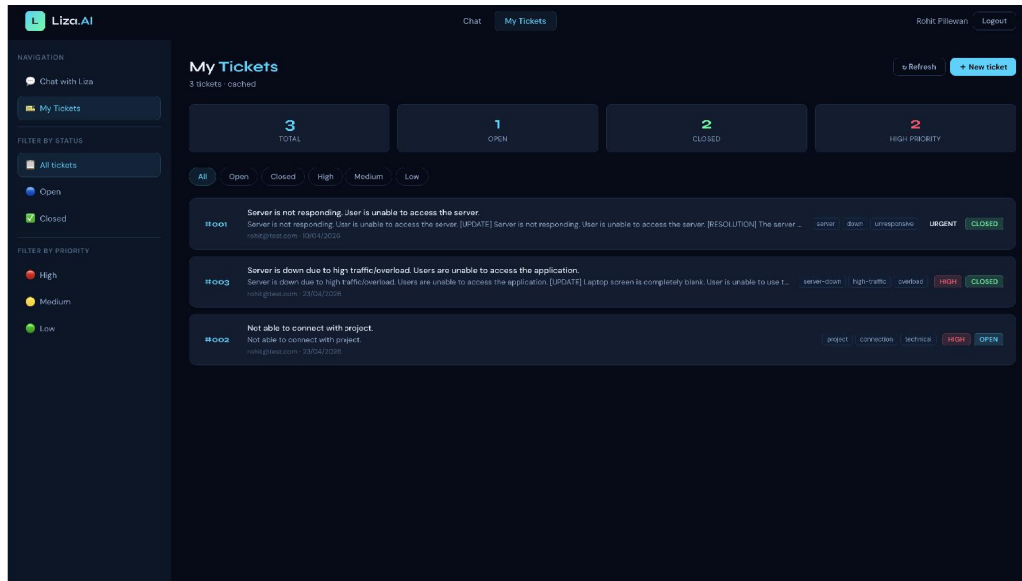


Fig. 6. Inline Ticket Card Rendered Inside AI Chat Bubble

#### D. Ticket Dashboard and Filtering

The React ticket dashboard fetches all tickets via GET /api/v1/helpdesk/tickets and renders a statistics bar showing total, open, closed, and high-priority counts. Client-side filter tabs allow instant filtering by status (OPEN/CLOSED) and priority (HIGH/MEDIUM/LOW) without additional API calls. Clicking any ticket row opens a detail modal displaying the full annotated description history including [UPDATE] and [RESOLUTION] markers, timestamps, and all metadata fields.





*Fig. 7. Ticket Dashboard with Real-Time Filtering and Statistics*

### E. Security Validation

All secured endpoints correctly rejected unauthenticated requests with 401 UNAUTHORIZED, expired tokens with a JWT\_EXPIRED error code, malformed signatures with JWT\_INVALID, and cross-user ticket access attempts with 403 FORBIDDEN — confirming that the ownership enforcement in TicketController is correctly applied.

### VI. CONCLUSION AND FUTURE WORK

This paper has demonstrated a practical and deployable approach to integrating conversational AI with enterprise IT support workflows using a locally hosted LLM. The proposed system achieves the complete support ticket lifecycle — Create, Update, Escalate, Close — through pure natural language interaction, without any form-based interface or cloud AI dependency. The dual-strategy intent detection combining structured prompt engineering with a backend rule engine achieves reliable intent classification even for ambiguous user inputs.

The per-user conversation memory architecture using email-based conversationId with JdbcChatMemoryRepository ensures complete session isolation and persistent multi-turn context. The JWT-based stateless authentication model is appropriate for REST API architectures and horizontally scalable deployments. LLaMA 3 8B via Ollama was confirmed to match cloud LLM performance on instruction-following tasks at zero operational cost.

Future work will address the following planned enhancements: (1) Role-Based Access Control (RBAC) with an administrative panel for cross-user ticket management and SLA analytics; (2) real-time push notifications via Spring WebSocket (SockJS/STOMP); (3) automated email alerts through Spring Mail integration; (4) multilingual support for Hindi and Marathi inputs using multilingual LLaMA variants; (5) semantic ticket search using vector embeddings (pgvector/ChromaDB); and (6) a React Native mobile application with Firebase push notification support.

### REFERENCES

- [1] Y. Liu et al., "RoBERTa: A Robustly Optimized BERT Pretraining Approach," arXiv preprint arXiv:1907.11692, 2019. Referenced for transformer-based intent detection methodology.
- [2] P. Xu, C. S. Wu, A. Madotto, and P. Fung, "Towards Scalable Multi-Domain Conversational Agents: The Schema-Guided Dialogue Dataset," AAAI Conference, 2020. Referenced for multi-turn dialogue state tracking.



- [3] J. Rashid and I. Ahmad, "Automatic IT Ticket Classification Using Natural Language Processing," International Journal of Advanced Computer Science and Applications, vol. 12, no. 5, 2021.
- [4] H. Touvron et al., "LLaMA: Open and Efficient Foundation Language Models," arXiv preprint arXiv:2302.13971, 2023. Referenced for local LLM inference benchmarks.
- [5] Spring AI Documentation. Spring AI 1.0.3 Reference Guide. <https://docs.spring.io/spring-ai/reference/>. Referenced for ChatClient, JdbcChatMemoryRepository configuration.
- [6] Ollama Documentation. Ollama Local LLM Runtime. <https://ollama.com/docs>. Referenced for LLaMA 3 model configuration and local HTTP API specification.
- [7] JSON Web Token Specification. RFC 7519. <https://tools.ietf.org/html/rfc7519>. Technical specification for JWT structure and HMAC-SHA256 signing.
- [8] Spring Boot Documentation. Spring Boot 3.x Reference Guide. <https://docs.spring.io/spring-boot/docs/current/reference/html/>.
- [9] React Documentation. React 18 Official Guide. <https://react.dev/>. Referenced for hooks, Context API, and component architecture patterns.
- [10] Vite Documentation. Vite 5.x Build Tool Guide. <https://vitejs.dev/guide/>. Referenced for SPA configuration and production build optimization.

