

IEEE YESIST12 IEngage Track Problem Statement

AI-Assisted Ticket Classification, Urgency Estimation, KB Retrieval, and Grounded Response Drafting

2. Abstract

Customer support teams handle high volumes of tickets across categories such as Refund, Login, Delivery, Billing, Account, and Other. Manual triage delays response times, creates inconsistent prioritization, and increases workload for support agents. This problem focuses on building a system that classifies ticket category, estimates urgency, retrieves relevant Knowledge Base snippets, summarizes the issue, and drafts a short reply grounded in the KB. The system must not guess when the KB lacks sufficient information and should ask for clarification instead.

3. Keywords

Ticket Triage, Customer Support Automation, Knowledge Base Retrieval, Classification, Urgency Estimation, Draft Reply, Grounded AI, Semantic Search, Helpdesk AI

4. Introduction

Customer support is a critical customer-facing function where timely, accurate, and consistent responses directly affect customer satisfaction. Support tickets often arrive with unstructured subject lines, descriptions, and optional metadata such as channel or timestamp. An AI-assisted triage system can reduce the burden on support teams by routing tickets, estimating urgency, retrieving relevant KB content, and preparing a grounded draft reply for agent review.

5. Background and Motivation

The motivation is to augment support agents rather than replace human judgment. Many support workflows suffer from repetitive classification, delayed prioritization, and inconsistent replies. A grounded triage system can standardize first-level processing while ensuring that draft replies are based on approved Knowledge Base content.

6. Problem Statement

Build a system to classify support tickets, estimate urgency, retrieve KB snippets, and draft short replies using a Knowledge Base. The system must classify ticket category as Refund, Login, Delivery, Billing, Account, or Other; estimate urgency as High, Medium, or Low; retrieve relevant KB snippets; and draft a concise reply grounded in KB evidence. If KB information is insufficient, the system must not guess and should ask for clarification.

Scope & Constraints

In Scope:

- Ticket preprocessing using subject, description, and optional metadata.
- Category classification using rules, machine learning, or hybrid logic.
- Urgency estimation based on ticket content, severity indicators, and optional metadata.
- KB ingestion from .txt, .md, and .pdf files.
- Embedding generation, vector search, and top-k retrieval of KB snippets.
- Draft reply generation grounded in retrieved KB content.
- Structured output containing category, urgency, summary, draft_reply, sources, and confidence.

Out of Scope:

- Sending replies directly to customers without review, unless explicitly configured in a controlled workflow.
- Making policy exceptions or refunds without business approval.
- Using unsupported information outside the provided KB.
- Replacing CRM/helpdesk systems end-to-end; integration may be considered optional.

7. Scope of the Problem

The problem domain includes the technical and operational activities required to design, implement, evaluate, and demonstrate the proposed solution.

- Ticket preprocessing using subject, description, and optional metadata.
- Category classification using rules, machine learning, or hybrid logic.
- Urgency estimation based on ticket content, severity indicators, and optional metadata.
- KB ingestion from .txt, .md, and .pdf files.
- Embedding generation, vector search, and top-k retrieval of KB snippets.

8. Objectives

- Reduce manual effort in ticket categorization and prioritization.
- Improve consistency of urgency assessment across support queues.
- Generate concise summaries to help agents understand issues quickly.
- Retrieve relevant KB snippets and cite them in the response.
- Produce grounded draft replies that reduce agent drafting time.
- Ask clarification instead of guessing when KB content is missing.

9. Constraints and Assumptions

- Ticket text may be incomplete, emotional, multilingual, noisy, or ambiguous.
- KB coverage may be limited or outdated, affecting answerability.
- Urgency signals may differ by business domain and require calibration.
- Draft replies must be professional, concise, and aligned with KB policy.
- Classification confidence must be visible to support agents.
- The system must prevent unsupported commitments such as refunds, delivery promises, or account changes.

Assumptions:

- Participants may use open-source frameworks, public datasets, or synthetically prepared sample data where appropriate.
- The final solution should include a working prototype, documented architecture, sample inputs/outputs, and measurable evaluation results.
- All generated outputs should remain explainable, traceable, and safe for human review.

10. Significance of the Problem

Academic Relevance

- Supports applied research in text classification, intent detection, urgency scoring, and retrieval-augmented response generation.
- Provides opportunities to compare rule-based, ML, and hybrid triage approaches.
- Encourages evaluation of grounded response generation and confidence estimation in service workflows.

Industry and Societal Impact

- Improves support team productivity by automating repetitive triage steps.
- Reduces first-response time and improves customer experience.
- Standardizes support responses using approved KB content.
- Helps prioritize high-urgency tickets and reduce SLA risk.

11. Expected Outcomes

Technical Outputs

- Structured ticket analysis containing category, urgency, and summary.
- Relevant KB sources with citations and confidence.
- Short draft reply grounded in retrieved KB snippets.
- Clarification request when information is insufficient.
- Optional analytics on ticket category distribution and urgency trends.

Suggested KPIs to Track

Metric / KPI	Expected Measurement / Target Direction
Category Classification Accuracy	Higher percentage of tickets assigned to the correct category.
Urgency Estimation Accuracy	Higher agreement with human-labeled urgency.
Draft Reply Grounding	All reply claims supported by retrieved KB snippets.
Average Triage Time Reduction	Lower time spent by agents on first-level triage.
Clarification Quality	Appropriate clarification asked when KB evidence is insufficient.
Agent Acceptance Rate	Percentage of draft replies accepted or minimally edited by agents.