



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** X **Month of publication:** October 2025

DOI: <https://doi.org/10.22214/ijraset.2025.74734>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

Real-Time AI Powered StackOverflow

T.P. Yagna Narayanan¹, S. Vishnu Priyan², S. Manikavasagam³, Mrs. Mutheewari⁴

^{1,2,3}Department of Computer Science Engineering, ⁴Associate Professor / CSE, K.L.N College of Engineering, Pottapalayam, Sivagangai

Abstract: *This study proposes an intelligent AI-powered platform for software question-and-answer (Q&A) management, aimed at enhancing developer productivity and collaboration. Traditional platforms, such as Stack Overflow, often face slow responses, duplicate queries, and inconsistent quality of answers due to reliance on manual input. The proposed system leverages advanced AI models, including GPT and BERT, to automatically generate, validate, and rank answers in real-time, while detecting duplicate questions through semantic similarity analysis. Additionally, the system recommends bug fixes and improvements by analyzing submitted code snippets. Built using a Python Flask backend, React.js frontend with Tailwind CSS, and a MySQL/MongoDB database, the platform ensures a responsive, scalable, and user-friendly experience. By integrating AI-driven forecasting of trending technical queries, real-time feedback loops, and automated code validation, the system fosters efficient knowledge sharing and proactive developer assistance, addressing challenges of traditional Q&A platforms.*

Keywords: AI Q&A Platform, Automated Answer Generation, Code Validation, Bug Detection, Semantic Similarity, GPT, BERT, Real-Time Feedback, Developer Productivity

I. INTRODUCTION

Online software developer communities have become essential for knowledge sharing, problem-solving, and collaborative learning. Platforms such as Stack Overflow rely heavily on human contributions to answer technical questions. However, this dependence results in slow response times, duplicate questions, and variable answer quality. In addition, developers often struggle with outdated or incomplete answers, inconsistent coding practices, and a lack of contextual understanding of their queries.

Recent advancements in artificial intelligence, natural language processing (NLP), and deep learning provide opportunities to automate and enhance these processes [1,2]. By leveraging transformer-based models like GPT and BERT, AI systems can automatically generate answers, validate code snippets, and suggest bug fixes. Semantic similarity detection reduces question duplication, while AI-driven ranking ensures answers are relevant, accurate, and clear. Furthermore, AI models can provide contextual explanations, highlight best practices, and suggest performance optimizations for submitted code, enhancing learning outcomes for both novice and experienced developers. The proposed system integrates these capabilities into a single web-based platform, enabling developers to efficiently obtain validated solutions, track trending questions, and proactively address technical issues. Additionally, the system supports real-time feedback loops, which allow AI models to continuously learn from user interactions, improving answer quality and relevance over time. By providing structured summaries, actionable insights, and predictive analytics for emerging topics, the platform not only addresses immediate developer queries but also fosters long-term knowledge retention and skill development. Existing research on AI-assisted Q&A systems highlights the potential for improved efficiency and knowledge management but often lacks real-time code validation, semantic deduplication, predictive insights for trending queries, and user-adaptive learning mechanisms [3,4,5]. This platform aims to fill these gaps, providing an intelligent, automated solution for modern software development communities that reduces cognitive load, minimizes redundancy, and enhances collaborative problem-solving.

II. METHODOLOGY

The system architecture of AI-Powered Stack Overflow (Figure 1) implements a structured pipeline from question submission to answer recommendation, summarization, and contextual highlighting. The main components are detailed below:

- 1) **Question Capture:** Users input queries via a web interface developed in React.js. Queries can include natural language text, embedded code snippets, or technical logs. The interface supports multi-line input, syntax highlighting for code, and error message detection to preserve context.
- 2) **Text Preprocessing:** The system cleans the input text, removes stop words, normalizes terms, and tokenizes sentences. Advanced preprocessing identifies technical entities using NER, regex, and embedding-based classification. For example, if a user asks, "Why is NullPointerException thrown in Java streams?" the system recognizes NullPointerException, Java, and streams as entities.

- 3) **Semantic Analysis:** Transformer-based models (GPT, BERT) encode the question into a semantic vector and compare it against a precomputed embedding database of previous questions and answers. This allows the system to identify contextually similar queries, even if they use different wording. Ranking algorithms prioritize answers based on semantic similarity, user votes, and answer recency.
- 4) **Answer Generation and Summarization:** If a similar question exists, the system retrieves relevant answers; otherwise, GPT generates new answers using learned patterns and technical knowledge. Multi-answer threads are summarized into concise paragraphs, emphasizing critical code snippets, troubleshooting steps, and key concepts.
- 5) **Display and Interaction:** The web interface dynamically displays suggested answers, summaries, and contextual highlights. Users can vote on answers, save summaries, or continue discussion threads.

All data, including questions, answers, embeddings, and interaction logs, are stored in **SQL**, indexed for fast retrieval. Node.js and Express.js handle API requests, while React provides a real-time front-end update system.




Fig. 1 Flow diagram

III. TEXT PREPROCESSING AND ENTITY RECOGNITION

To achieve high-quality semantic understanding, rigorous preprocessing is necessary:

- 1) **Tokenization & Lemmatization:** Converts text into normalized tokens, which reduces variations due to tense, plural forms, or synonyms.
- 2) **Named Entity Recognition (NER):** Identifies programming languages, APIs, libraries, functions, error types, and variable names.
- 3) **Code Detection:** Recognizes inline and block code snippets, preserving syntax and formatting for analysis.
- 4) **Noise Filtering:** Removes irrelevant symbols, extraneous punctuation, or non-informative words without affecting technical meaning.

For instance, the question

“How to fix `IndexError` in Python when using a list comprehension?”

is preprocessed to isolate `IndexError`, Python, and list comprehension as key entities. This structured input ensures precise semantic matching and answer generation.

IV. PROCESS FLOW

The AI-Driven Q&A Platform processes user queries into validated, ranked answers through a structured and methodical pipeline. The first step, User Query Submission, allows developers to post questions through the web interface. The system captures the query, associated tags, and user information to ensure contextual understanding and proper routing.

Next, the Backend Processing stage receives the query and prepares it for semantic analysis. At this stage, natural language processing techniques, including tokenization, entity recognition, and contextual embeddings, are applied to accurately understand the intent of the question. This ensures that subsequent AI modules have a clear and precise representation of the query.

In the Semantic Search & Duplicate Detection stage, the system compares the incoming query against the existing question repository using transformer-based embeddings. Duplicate or semantically similar questions are identified to reduce redundancy, while relevant existing answers are retrieved to enhance response quality.

The AI Answer Generation stage uses GPT and BERT models to produce context-aware, human-like answers. The system generates explanations, code snippets, and best practices, ensuring that responses are technically accurate and easily understandable.

During the Code Validation & Bug Detection stage, any submitted or generated code snippets undergo static and dynamic analysis. Errors, potential bugs, and optimization suggestions are provided in real-time, allowing users to quickly identify and fix issues.




Fig. 2 Process flow

This structured workflow ensures that the system provides developers with fast, accurate, and contextually relevant solutions while minimizing duplicate queries, improving knowledge retention, and enhancing collaborative problem-solving.

V. SEMANTIC MATCHING AND ANSWER RECOMMENDATION

After preprocessing, the system computes embeddings for the question and compares them against the existing database. Answers exceeding a similarity threshold are ranked and presented. For unmatched queries, GPT generates responses using both knowledge from existing answers and its general programming knowledge. The ranking mechanism incorporates:

- 1) Semantic similarity score
- 2) Answer quality indicators (votes, recency, author reputation)
- 3) Contextual relevance to entities detected in the question

This ensures the recommendations are accurate, timely, and relevant.

VI. SUMMARIZATION AND CONTEXTUAL HIGHLIGHTING

Summarization reduces cognitive load and improves usability by highlighting:

- 1) Key functions, classes, or methods
- 2) Code snippets with syntax preservation
- 3) Actionable instructions or troubleshooting steps
- 4) External references or documentation links

Contextual highlighting draws attention to entities and steps critical for problem-solving, allowing developers to scan threads efficiently without reading the full discussion.

VII. SYSTEM OPTIMIZATION

The system optimizes performance via:

- 1) Precomputed embeddings: Reduces similarity search time.
- 2) Caching frequently asked questions: Minimizes repeated processing.
- 3) Efficient MongoDB indexing: Speeds up data retrieval.
- 4) Asynchronous server operations: Ensures low latency for real-time suggestions.
- 5) Resource management: Handles multiple simultaneous users with minimal memory and CPU overhead.

VIII. RESULTS AND DISCUSSION

The AI-Powered Stack Overflow system delivers:

- 1) Real-time answer suggestions for user queries
- 2) High accuracy in entity recognition for code, libraries, and error types
- 3) Concise summaries of multi-answer threads
- 4) Reduced search and decision time for developers

Fig. 3 Question Input Interface




Fig. 4 Answer Recommendation Page.




Fig. 5 Question Posting Page.




Fig. 6 Duplicate detection while upload question.





Fig. 7 Trending Question Page.



Testing showed significant improvements in response relevance, reduced duplicate query creation, and increased user engagement. Feedback indicates developers saved up to 40–50% of time typically spent searching for answers.

IX. CONCLUSION

This study demonstrates an AI-powered, context-aware Stack Overflow assistant that enhances developer productivity through real-time answer recommendations, entity recognition, and automated summarization. By integrating transformer-based NLP models with a MERN stack architecture, the system improves knowledge retrieval, reduces duplication, and highlights key information effectively. Future work includes:

- 1) Multilingual support for international developer communities
- 2) Improved code validation and execution environment integration
- 3) Enhanced IDE plugins for seamless workflow integration
- 4) Advanced learning from user interactions to improve recommendation accuracy

The proposed system provides a scalable, intelligent solution for knowledge management in software development communities, bridging the gap between unstructured forum content and actionable insights.

REFERENCES

- [1] Vaswani, Ashish, et al. "Attention is All You Need." *NeurIPS*, 2017.
- [2] Devlin, Jacob, et al. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." *NAACL*, 2019.
- [3] Brown, Tom, et al. "Language Models are Few-Shot Learners." *NeurIPS*, 2020.
- [4] Jiang, Mi, et al. "NexaNota: An AI-Powered Smart Linked Lecture Note-Taking System Leveraging Large Language Models." *ICBDE*, 2025.
- [5] Wisoff, Josh, et al. "NoteBar: An AI-Assisted Note-Taking System for Personal Knowledge Management." *arXiv preprint*, 2025.
- [6] Radford, Alec, et al. "GPT-3: Language Models are Few-Shot Learners." *OpenAI Technical Report*, 2020.
- [7] Zhou, YunYu, et al. "Extracting Learning Data From Handwritten Notes Using AI." *IEEE Access*, 2025.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)