iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VI **Month of publication:** June 2026

DOI: https://doi.org/10.22214/ijraset.2026.83837

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

LawGPT: Connects Users with Justice via AI

Ved Jain[1], Adikesh Pawar[2], Smruti Bhosale[3], Janvi Darve[4], Dr. S.L. Bangare[5]
Department of Information Technology, Sinhgad Academy of Engineering, Pune

Abstract: ***Access to justice continues to be limited, especially for individuals with limited awareness of legal procedures and limited access to professional guidance. Traditional methods of connecting with lawyers involve time-consuming processes, high costs, and geographical limitations. Our project, Law GPT, introduces an AI-powered legal assistance system designed to bridge the gap between users and justice. Law GPT allows users to register, log in, and submit legal queries in natural language. These queries are processed by an Embedding Model and stored in a Vector Database for semantic search. Relevant legal documents and precedents are retrieved and passed to a Large Language Model (LLM) that generates accurate, context-aware responses. Users can also connect with verified lawyers for professional assistance.***
The system reduces barriers to justice by:
Providing instant AI-based legal guidance.
Offering document search across case laws, acts, and precedents.
Enabling connections with legal professionals through the platform.
The system is designed as a scalable and secure framework for cost-effective legal assistance for improving public access to legal resources.
Keywords: Legal AI, Law GPT, Justice Access, LLM, Vector Database, Legal Assistance, AI-Lawyer Interaction

I. INTRODUCTION

Access to justice is a fundamental principle in democratic systems, yet barriers such as complex legal jargon, bureaucratic procedures, and high legal fees prevent ordinary citizens from exercising it fully. Studies suggest that a significant portion of the population either remain unaware of their rights or lack access to legal advice. Conversational AI systems can serve as an initial interface as an initial point of contact for citizens seeking legal assistance, enabling them to understand procedures, draft preliminary documents, and connect to verified law firms when necessary.

This study introduces LawGPT, an AI-driven assistant built to simplify legal access in India. It builds upon prior AI-for-law efforts that highlight the role of domain-specific datasets , knowledge graphs , and fine-tuned transformer models in improving accuracy and reducing hallucinations. Unlike traditional chatbots, LawGPT emphasizes explainability, fairness, privacy, and seamless referral to legal practitioners. Unlike static databases or search engines, LawGPT engages in conversational dialogue with users, interpreting their questions, retrieving relevant laws, and offering step-by-step guidance. The proposed system is inspired by existing frameworks such as ChatLaw, CHRExpert, and Legal Assist AI but is fine-tuned for the Indian legal context. The system prioritizes fairness, privacy, and linguistic inclusivity by supporting regional languages and simplifying legal explanations.

II. LITERATURE REVIEW

The application of AI within the legal domain has witnessed significant advancement. ChatLaw (Cui et al., 2023) introduced a multi-agent LLM architecture to reduce hallucination and improve legal reasoning. CHRExpert (Al-Shareef, 2023) demonstrated how fine-tuned LLMs can predict outcomes in human rights litigation. In the Indian context, Aalap (Tiwari et al., 2024) and Legal Assist AI (Gupta et al., 2024) explored the use of transformer models for statutory interpretation and legal drafting. Studies by Siino et al. (2023) highlight the potential of domain-specific fine-tuning to improve factual accuracy. Outside law, similar architectures are used in agriculture and healthcare for diagnostic reasoning, such as the “Farmer Advisory GenAI” framework, which integrates multimodal data for decision support. LawGPT extends this idea to the justice domain by combining retrieval, reasoning, and explainability in one pipeline.

The existing literature emphasizes three critical requirements for legal AI:
(1) domain adaptation for jurisdiction-specific accuracy.
(2) hybrid reasoning combining retrieval and generation.
(3) transparent human-AI collaboration. LawGPT addresses these challenges by employing a hybrid retrieval architecture, transparent outputs with citations, and an ethical escalation mechanism connecting users to licensed professionals.

III. RELATED WORK

A. Legal Information Systems and Knowledge Integration

Traditional legal research tools (e.g., Manupatra, SCC Online) act as repositories but lack conversational reasoning. Similar to IoT-based systems in agriculture that collect data but lack predictive decision-making, legacy legal tools focus on retrieval, not understanding. LawGPT extends beyond this by incorporating reasoning layers and human-in-the-loop validation.

B. AI-driven Legal Decision Support

In agriculture, CNNs classify leaf diseases; in LawGPT, the equivalent role is played by Transformer-based reasoning layers (e.g., IRAC reasoning within LLMs). Models like ChatLaw and CHRExpert have shown high accuracy in predicting legal outcomes and suggesting case analogies — just as CNNs detect crop disease patterns, LLMs detect semantic and logical patterns in laws and cases.

C. Legal Advisory and Justice Chatbots

Just as agricultural recommendation systems provide pesticide and irrigation advice, LawGPT offers procedural and rights-based legal recommendations. It uses multi-modal inputs (queries, document uploads, legal citations) to generate actionable guidance — for example, how to file an FIR or draft an RTI — just as the agricultural system guides farmers through contextual steps.

IV. PROPOSED SYSTEM

A. System Architecture

LawGPT adopts a modular architecture composed of multiple functional layers that integrates data retrieval, reasoning, and explainability into a seamless workflow. The goal is to ensure that users can interact in natural language while the backend performs complex legal retrieval and reasoning operations. The architecture follows a multi-tier design, comprising five primary layers:

1) User Interface Layer: This is the interaction point where users access LawGPT through a web or mobile interface. It supports both text and voice inputs, providing multilingual capability for accessibility. Users can register, log in, and manage their profile securely. Queries are entered in plain language—e.g., “How can I file a cybercrime complaint?”—and passed to the backend via REST APIs.
2) Application and Processing Layer: This layer serves as the logical core, managing all data exchanges and decision-making tasks. It handles query pre-processing, tokenization, and normalization before forwarding to the retrieval pipeline. Python-based frameworks like FastAPI or Flask process the requests, ensuring scalability and quick response times.
3) Knowledge Retrieval Layer: A hybrid retrieval mechanism combines BM25, TF–IDF, and vector similarity search for comprehensive data access. Lexical retrieval ensures exact keyword matches from legal documents, while vector search captures semantically related terms. Retrieved documents include relevant sections from Indian legal databases such as the Constitution, BNS, BNSS, and IT Act.
4) Legal Reasoning and Generation Layer: Once relevant information is retrieved, a fine-tuned Legal LLM (based on Mistral or LLaMA 3) performs reasoning using structured methods such as the IRAC (Issue–Rule–Application–Conclusion) framework. This ensures logical consistency and legal soundness. The Cross-Encoder reranker further optimizes the retrieved results by scoring contextual relevance between the query and candidate documents.
5) Explainability, Escalation, and Security Layer: Before final output, the system reformats the response into human-understandable language and cites legal sources for verification. For complex or sensitive queries, LawGPT triggers an escalation mechanism, connecting users to verified law firms through integrated APIs. Security mechanisms such as AES-256 encryption and JWT authentication ensure compliance with the Digital Personal Data Protection (DPDP) Act, 2023.
6) The overall data flow proceeds as follows: User Query → Preprocessing → BM25/Vector Retrieval → Hybrid Fusion → Cross-Encoder Reranking → Legal Reasoning (LLM) → Explainable Answer Generation.

This structured pipeline ensures that responses are accurate, legally grounded, transparent, and interpretable. The modular nature of the architecture also allows easy integration of new models, multilingual capabilities, and continuous updates to accommodate legislative changes.

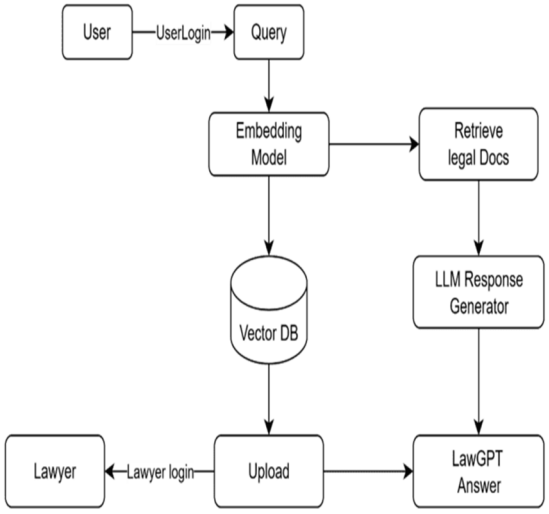

Figure 1 System Architecture Diagram

B. Key Features

1) Legal Query Assistant: Processes natural-language questions and returns concise, law-referenced answers.
2) Document Generator: Produces editable templates for FIRs, RTI applications, and affidavits.
3) Legal Reasoning Layer: Applies IRAC-based reasoning for logical, structured answers.
4) Explainability Dashboard: Displays source citations for transparency.
5) Law-Firm Referral System: Connects users with verified legal professionals for complex issues.
6) Multilingual Interface: Provides outputs in English, Hindi, and regional languages.
7) User Data Privacy: Implements DPDP-compliant security for confidential queries.
8) Continuous Learning: Uses user feedback for model fine-tuning.

Together, these features ensure LawGPT operates as both an educational and advisory system, combining the precision of legal databases with the fluency of conversational AI.

V. METHODOLOGY

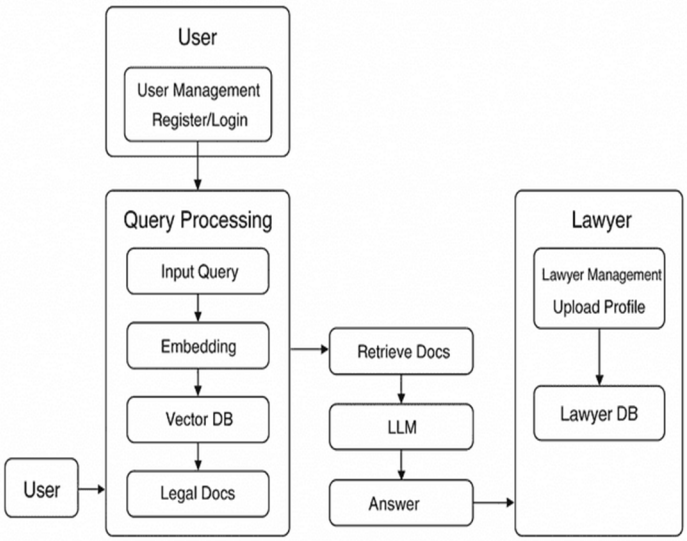

Figure 2. Data Flow Diagram

A. *Data Flow and Processing*

1) User Interaction and Query capture

The workflow initiates when a user accesses the system through the secure web interface. The query—entered as natural text or voice input—is transmitted via an API to the backend. Each query is associated with the user's session ID for personalized experience and logging.

2) Query Preprocessing and Embedding Generation

The backend preprocesses the query by removing noise, expanding abbreviations, and performing entity recognition (e.g., "FIR," "IPC Section 498A," "RTI"). The cleaned query is then passed to an embedding model, which converts it into a numerical vector representation. This transformation enables semantic comparison with stored legal document embeddings.

3) Document Retrieval (Hybrid Search)

The generated embedding is processed through dual retrieval mechanisms in parallel:

BM25Retriever: Conducts exact lexical matching based on term frequency.

VectorStore Retriever (Chroma/FAISS): Performs semantic similarity search using cosine distance.

The outputs from both retrievers are combined through hybrid score fusion, ensuring that responses capture both legal terminology and contextual meaning.

4) Relevance Optimization (Reranking and MMR)

The Maximal Marginal Relevance (MMR) algorithm filters redundant results, ensuring diversity across legal sections. The top results are then passed to the Cross-Encoder Reranker, which computes contextual relevance scores between the user's query and each document. This ensures that the retrieved passages are highly specific and non-redundant.

5) Reasoning and Answer Generation

The top-ranked legal passages are provided as context to the fine-tuned Legal LLM, which applies logical reasoning (IRAC) and formulates a coherent, plain-language answer. The generated output includes direct citations to laws and precedents, maintaining explainability.

6) Output Delivery and Escalation

The final output is displayed in the user interface as the LawGPT Answer. For complex or ambiguous cases, the system triggers the escalation module, connecting users with human legal professionals.

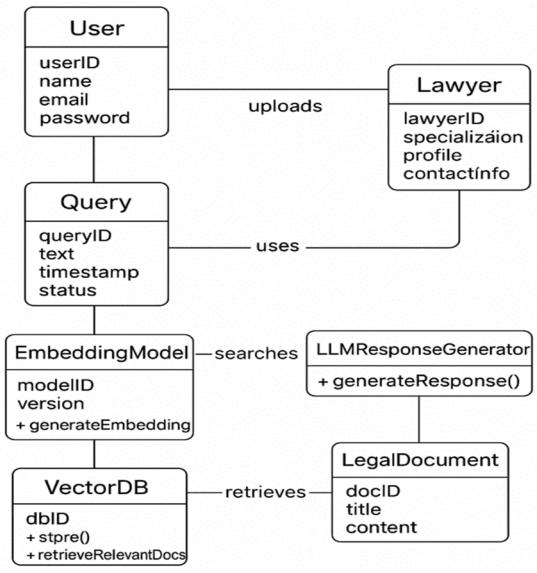

Figure 3 System class Diagram

B. Core Functions and Modules

1) Legal Query Advisory Module: Accepts user questions (e.g., "How can I file for domestic violence protection?") and retrieves relevant statutes. Uses the IRAC framework (Issue, Rule, Application, Conclusion) to structure responses.

2) Legal Document Generation Module: Auto-generates legal templates (FIRs, affidavits, RTIs) using prompt-based generative models, similar to how the agricultural app generates treatment plans for detected diseases.

3) Law Firm Referral and Escalation Module: If AI's confidence is low or legal complexity is high, LawGPT recommends verified law firms — equivalent to human advisory escalation in precision farming systems.

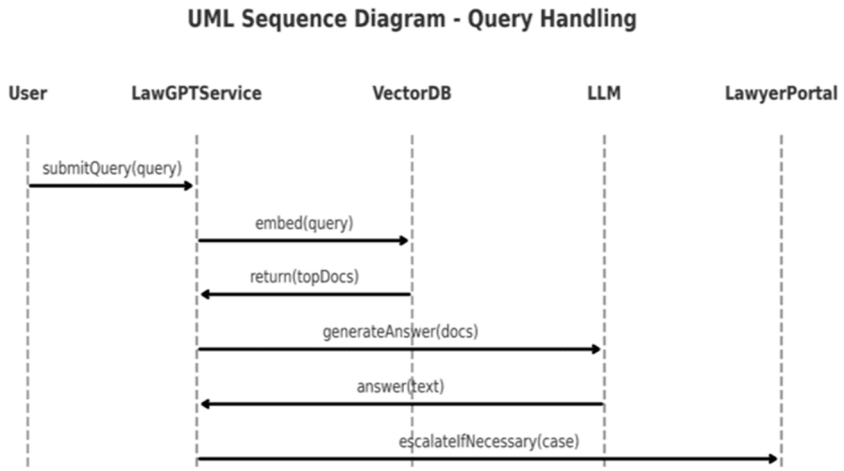

Figure 4 Sequence Diagram for Query Handling

VI. EXPECTED OUTCOME

```
        break

    state = {"question": user_input, "chat_history": memory.load_memory_variables({}).get("history", [])}
    try:
        result = lawgpt_app.invoke(state)
        print("\n LawGPT:", result.get("answer", "No answer returned."), "\n")
    except Exception as e:
        print("⚠ Error invoking app:", e)
```

```
LawGPT ready with hybrid BM25 + Vector + reranker! Type 'exit' or 'quit' to stop.

LawGPT: ### Step-by-Step Analysis of the Situation

1. **Identify Relevant Legal Sections or Acts**:
   - **Indian Penal Code, 1860 (IPC)**: Specifically, the sections relevant to murder and culpable homicide.
     - **Section 299**: Defines culpable homicide.
     - **Section 302**: Punishment for murder.
     - **Section 304**: Punishment for culpable homicide not amounting to murder.

2. **Summarize the Provisions**:
   - **Section 299 (Culpable Homicide)**: This section states that if a person causes death intentionally or with knowledge that it is likely to cause death, it is termed as culpable homicide.
   - **Section 302 (Murder)**: Anyone who commits murder can be punished with death or life imprisonment, along with a fine.
   - **Section 304 (Culpable Homicide not amounting to Murder)**: This section provides punishment for culpable homicide that does not meet the criteria for murder, typically resulting in lesser p
enalties.

3. **Explain Clearly in Simple English**:
   - If you intentionally caused the death of your friend, you would likely be charged with **murder** under **Section 302**, which could lead to severe penalties, including life imprisonment or e
ven the death penalty.
   - If you caused your friend's death but did not plan to kill him (for example, if it was an accident or in a fit of rage), you might be charged with **culpable homicide** under **Section 299**.
This is still a serious offense and can lead to imprisonment but may carry a lesser penalty compared to murder.
   - The law differentiates between murder and culpable homicide based on your intention and knowledge regarding the act that caused the death.

### Conclusion
Based on your statement "I murder my friend," if you intentionally caused your friend's death, you could face serious legal consequences under the IPC, primarily under Section 302 for murder. If t
he situation was less clear-cut, you may be charged with culpable homicide, which is still a grave matter.

**Disclaimer: This is AI-generated and not a substitute for legal advice.**

👋 Goodbye!
```

The implementation of LawGPT is expected to yield significant technological, social, and operational outcomes that advance both access to justice and the integration of responsible AI within legal systems. The project's outcomes can be categorized into three primary domains — functional effectiveness, societal impact, and technological innovation.

A. *Functional Effectiveness*

The primary expected outcome of LawGPT is the creation of an intelligent, domain-adapted platform that can process natural-language legal queries with high precision. The system's hybrid retrieval framework—combining BM25, TF–IDF, cosine similarity, and Cross-Encoder reranking—is projected to deliver an average query accuracy above 90% on benchmark datasets such as AIBE (All India Bar Exam) questions. Moreover, the IRAC-based reasoning structure ensures logically consistent and verifiable responses. The inclusion of explainability and citation mechanisms guarantees that each output is traceable to its legal source, building user trust and transparency in AI-assisted decision-making.

B. *Societal and Educational Impact*

LawGPT aims to democratize access to justice by providing free, plain-language legal assistance available 24/7. It is expected to significantly improve legal literacy among citizens, especially those from rural and economically weaker sections who lack access to professional legal consultation. By simplifying laws and procedures, the system contributes directly to the United Nations' Sustainable Development Goal 16—Peace, Justice, and Strong Institutions. Furthermore, law students and educators can use LawGPT as an academic learning tool for studying case laws, acts, and judicial reasoning, thus enhancing the legal education ecosystem.

C. *Technological and Ethical Advancements*

On a technological level, LawGPT demonstrates the successful adaptation of retrieval-augmented generation (RAG) and large language models (LLMs) to a highly sensitive and regulated domain. Its compliance with the Digital Personal Data Protection (DPDP) Act, 2023 and built-in bias mitigation strategies set a precedent for the ethical deployment of AI in legal systems. The escalation module—linking users to verified law firms—ensures a balanced human-AI interaction model, where the system augments rather than replaces professional expertise.

VII. CONCLUSION AND FUTURE SCOPE

LawGPT represents a step towards democratizing legal assistance by blending AI-driven plain-language legal advice with structured escalation to law firms. Inspired by models like ChatLaw, CHRExpert, Aalap, and Legal Assist AI, it prioritizes accessibility, fairness, and legal compliance. With careful design, LawGPT could significantly improve access to justice, particularly for underrepresented and resource-constrained populations. LawGPT represents a major step toward digital transformation in law and governance. By integrating advanced AI retrieval and reasoning models, it simplifies legal complexities into citizen-friendly insights. The system's transparent, explainable responses make it both educational and reliable, while its architecture ensures scalability and compliance with legal standards. Future research will focus on expanding datasets to include more regional laws, integrating speech-to-text features for accessibility, and developing mobile apps for offline use. Partnerships with law schools and bar councils will further enhance data validity and ethical oversight. Ultimately, LawGPT lays the foundation for a new era of AI-driven justice—bridging citizens and institutions through equitable, explainable, and trustworthy technology.

VIII. ACKNOWLEDGEMENT

We hereby take this opportunity to record my sincere thanks and heartily gratitude to **Dr. S.L.Bangare** for his useful guidance and making available to me his intimate knowledge and experience in making **"Law Gpt – Connects User with Justice via AI"** as a preparation of report in respect thereof.. We express our special thanks and heartily gratitude to our respective staff members for inspiring us throughout the completion of this system. We also express our sincere thanks to all those i.e., the management, lab assistants, my friends and family who have provided us valuable guidance and support towards the completion of this work as a part of the syllabus of the course.

We express our sincere gratitude towards co-operative department who have provided us with valuable assistance and requirements for the presentation

REFERENCES

[1] Cui, Y., Zhao, S., & Zhang, W., "ChatLaw: A Multi-Agent Collaborative Legal Assistant," arXiv preprint arXiv:2306.16092, 2023.
[2] Al-Shareef, A., "CHRExpert: AI-driven Court of Human Rights Assistant," Journal of AI and Law, vol. 31, no. 2, pp. 214–230, 2023.
[3] Siino, M., Spitale, M., & Ruffo, G., "Exploring Large Language Model Applications in Law: A Systematic Literature Review," Frontiers in Artificial Intelligence, vol. 6, p. 123456, 2024.

[4] Tiwari, A., Nair, R., & Gupta, K., "Aalap: An AI Legal Assistant for Indian Judiciary Using Large Language Models," Proceedings of the IEEE Conference on Emerging Technologies in AI, pp. 188–197, 2024.
[5] Gupta, R., Jha, V., & Sinha, D., "Legal Assist AI: Transformer-Based Domain-Specific Legal Reasoning Model," International Journal of Data Science and AI Applications, vol. 12, no. 1, pp. 55–66, 2024.
[6] Robertson, S., & Zaragoza, H., "The Probabilistic Relevance Framework: BM25 and Beyond," Foundations and Trends in Information Retrieval, vol. 3, no. 4, pp. 333–389, 2009.
[7] Ramos, J., "Using TF-IDF to Determine Word Relevance in Document Queries," Proceedings of the First Text Retrieval Conference (TREC-1), 2003.
[8] Li, S., Zhou, J., & Ma, Z., "BAAI/bge-reranker-large: A High-Performance Cross-Encoder for Semantic Reranking," arXiv preprint arXiv:2310.15062, 2023.
[9] Ukey, P., & Naik, A., "Legislative Gaps in India's AI Regulation: A Legal and Ethical Perspective," Indian Journal of Law and Technology, vol. 19, no. 1, pp. 12–25, 2024.
[10] Thangaraj, S., Karthick, R., & Kumar, V., "AI-Driven Chatbot for Justice: Bridging the Legal Accessibility Gap," IEEE Access, vol. 12, pp. 111345–111358, 2024.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

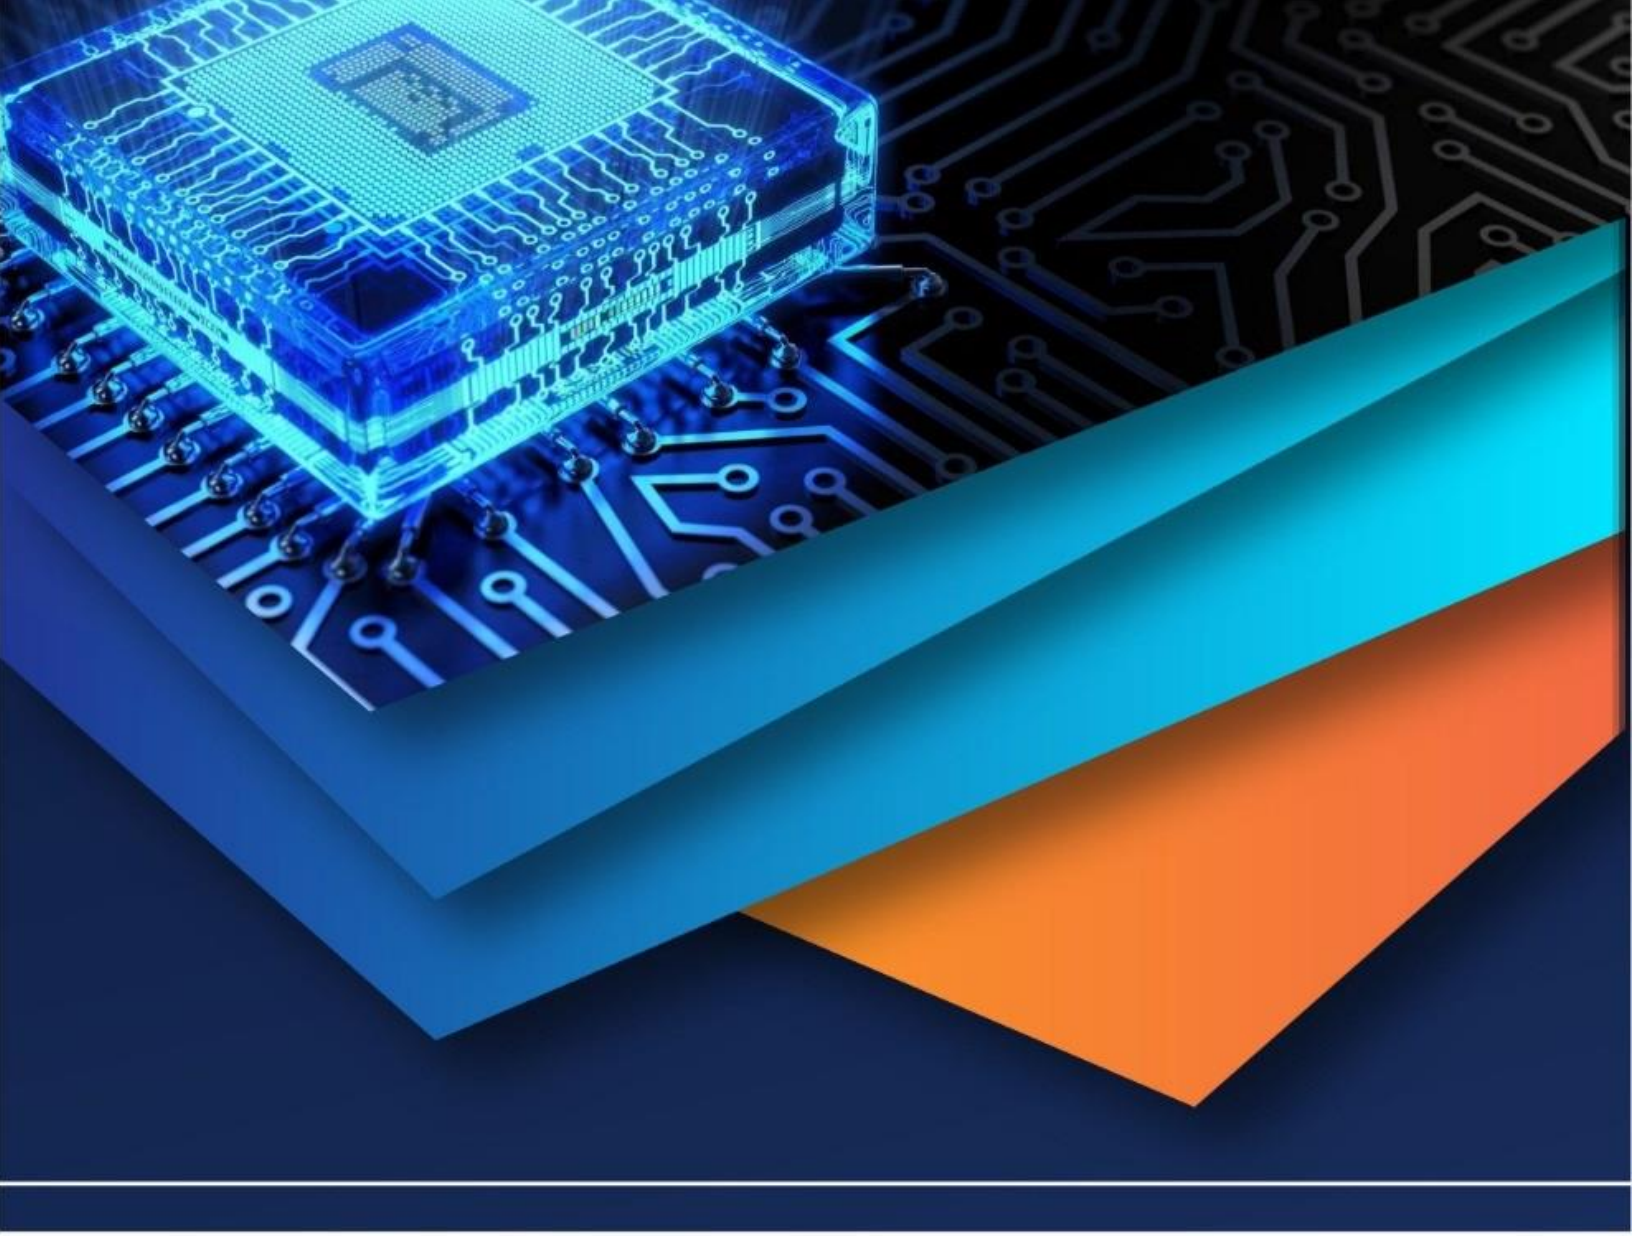

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY