iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 13 **Issue:** VII **Month of publication:** July 2025

DOI: https://doi.org/10.22214/ijraset.2025.73331

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 13 Issue VII July 2025- Available at www.ijraset.com

Smart Yoga Asana Recommender: Leveraging NLP for Health-Driven Suggestions

Gampa Jahnavi[1], Dr. K. Santhi Sree[2]
[1]*Post Graduate Student, MCA, Department of Information Technology, Jawaharlal Nehru Technological University, Hyderabad, India*
[2]*Professor, Department of Information Technology, Jawaharlal Nehru Technological University, Hyderabad, India*

Abstract: ***In recent years, yoga has emerged as a powerful tool for promoting holistic well-being. However, choosing appropriate yoga asanas tailored to an individual's age, gender, and specific health concerns often requires expert consultation, which may not always be accessible. This paper presents a Smart Yoga Asana Recommender System that uses Natural Language Processing (NLP) and machine learning to generate personalized yoga recommendations. The system collects user inputs - age, gender, and a brief health issue description - and processes the textual data using a SentenceTransformer model to generate semantic embeddings. These embeddings are matched with a curated database of yoga asana descriptions using Facebook AI Similarity Search (FAISS), enabling fast and contextually relevant retrieval. Recommended asanas are compiled into a PDF report and automatically sent to the user's email for convenient access. Built using Python, Flask, and MongoDB, the system offers a seamless web interface and integrates quote-based motivation for enhanced user engagement. This project combines traditional yogic wisdom with modern AI techniques to create a scalable and intelligent wellness tool, making expert-level yoga guidance accessible to individuals anytime, anywhere.***

Keywords: Yoga Asana Recommendation, Natural Language Processing (NLP), Sentence Embeddings, FAISS (Facebook AI Similarity Search), Personalized Wellness, Flask Web Application, Semantic Similarity Matching.

I. INTRODUCTION

Maintaining mental and physical wellness has become increasingly important in today's fast-paced and digitally driven lifestyle. Among the many wellness practices available, yoga stands out as a time-tested discipline that offers holistic benefits for the body, mind, and spirit. With a growing global awareness of its therapeutic potential, yoga has been embraced not only for fitness but also as a complementary approach to managing various health conditions. However, selecting the right yoga asanas often depends on factors such as an individual's age, gender, and specific health concerns. Traditionally, this selection process is guided by trained yoga practitioners. For many people, especially those in remote or underserved areas, such expert advice may not be readily accessible. In response to this challenge, the integration of Artificial Intelligence (AI) and Natural Language Processing (NLP) offers a promising solution. By automating the recommendation process, it is possible to provide personalized yoga guidance based on user-specific inputs without the need for direct human consultation. This paper introduces a Smart Yoga Asana Recommender System that leverages NLP techniques and machine learning to offer customized yoga pose suggestions.

The system accepts basic user inputs including age, gender, and a brief description of the user's health concern. This textual input is processed using a pre-trained SentenceTransformer model, which converts it into a semantic vector representation. These embeddings are compared with a curated database of yoga pose descriptions using FAISS (Facebook AI Similarity Search), enabling the system to identify the most relevant asanas quickly and accurately. The selected poses are then compiled into a personalized PDF report and emailed to the user, ensuring both convenience and accessibility.

Developed using Python, Flask, and MongoDB, the system also incorporates motivational quotes and a user-friendly web interface to enhance the overall experience. By combining traditional yogic wisdom with modern AI tools, this project aims to democratize access to expert-level yoga recommendations and contribute to the broader field of intelligent wellness systems.

A. Objective

The primary objective of this project is to develop an intelligent yoga asana recommender system that provides users with personalized pose suggestions based on their age, gender, and health concerns. The system is designed to interpret natural language input using semantic embeddings, identify contextually relevant yoga poses through similarity matching, and deliver the results in a user-friendly PDF format via email.

In doing so, the project aims to eliminate the dependency on in-person consultations and bring expert-level yoga guidance to a broader audience. Additional goals include creating a scalable architecture using Flask and MongoDB, integrating motivational content for enhanced user engagement, and enabling seamless automation from input collection to email delivery.

II. LITERATURE SURVEY

Several research efforts have explored the use of artificial intelligence and machine learning in yoga-based wellness systems. This section reviews recent works relevant to yoga recommendation, posture correction, and mental health classification.

Ambarish et al. [1] proposed a yoga pose recommender system that utilizes a combination of TF-IDF, K-Nearest Neighbors (KNN), and Word2Vec techniques to recommend suitable asanas based on user-described physical and mental health issues, referred to as "pain points." Notably, the system also identifies contraindicated poses that may not be safe for certain conditions, adding an important layer of precaution. However, the model's effectiveness is limited by a relatively small dataset and lacks adaptability to individual demographic factors such as age and gender. Additionally, the absence of clinical trials and real-world validation raises concerns regarding the scalability and reliability of the system in broader applications.

In a student-centric approach, Ramamoorthy et al. [2] developed a yoga recommendation system aimed at promoting mental well-being. Their model employs machine learning algorithms, particularly Random Forest in conjunction with Word2Vec, to predict stress levels based on physiological indicators such as sleep quality and heart rate. The system reportedly achieved accuracy rates exceeding 94%, showcasing strong predictive performance. Nonetheless, the model's training was based on a limited dataset of 100 students, and its lack of demographic personalization restricts its applicability across a diverse population. Moreover, its reliance on questionnaires and manual input raises concerns about long-term scalability.

Tom et al. [3] focused on the detection of mental health disorders using Natural Language Processing. By analyzing user-generated content from Reddit forums, their system classified conditions such as depression, anxiety, PTSD, ADHD, and bipolar disorder. Their fine-tuned MentalBERT model achieved high F1-scores, demonstrating the value of domain-specific language models in mental health monitoring. However, challenges remained in accurately detecting overlapping symptoms, particularly for complex conditions like bipolar disorder.

Arya et al. [4] proposed a hardware-assisted solution involving a wearable belt to detect and correct poor posture in real time. The device, when paired with a mobile application, provided instant feedback and yoga-based exercise suggestions to promote spinal health. Although innovative, the system was restricted to seated postures and lacked support for standing or dynamic movements. Furthermore, the device's bulkiness posed limitations in terms of comfort and long-term usability.

Kaur and Kamboj [5] developed an automatic grammar correction system for Punjabi sentences using transformer-based deep learning models. They implemented and compared three architectures: an encoder–decoder with attention, a standard transformer, and a fine-tuned transformer. Among them, the fine-tuned transformer achieved the highest correction accuracy of 63% on test data. While the system shows promise in handling grammatical inconsistencies in Punjabi, its performance is limited by the size of the dataset and the lack of semantic context modeling, indicating scope for further enhancement using larger corpora and contextual language models.

Sriman et al. [6] introduced an Intelligent Document Interaction (IDI) system that allows users to converse with multiple PDFs using natural language. Their solution combines several advanced technologies including PyPDF2 for text extraction, FAISS for vector indexing, and Google Gemini AI for generating contextual responses. Built using Streamlit, the platform enables users to upload documents and receive semantically relevant answers to their queries in real time. The system addresses key limitations of traditional document search methods by improving scalability, query relevance, and response time. While the architecture is powerful and multimodal, its performance may depend on the quality of embeddings and the processing capacity of the deployed environment, especially in large-scale or multilingual scenarios.

These works provide valuable insights into the intersection of AI, NLP, and wellness. However, none of the reviewed systems offer a fully automated, text-driven, and personalized yoga recommendation pipeline that incorporates semantic understanding, demographic filtering, and user-centric PDF delivery. The proposed system addresses this gap by integrating SentenceTransformer, FAISS, Flask, and MongoDB into a unified architecture aimed at improving both accessibility and personalization in yoga guidance.

III.PROPOSED WORK

Before generating personalized yoga recommendations, the system requires a structured and well-defined knowledge base that captures the relationship between yoga asanas and their associated health benefits. This foundational layer is essential to ensure that the matching algorithm can identify contextually relevant poses based on user inputs.

The underlying dataset plays a critical role in facilitating this process by organizing yoga poses along with their suitability for various demographic factors such as age and gender, as well as health-related considerations. The following subsection outlines the structure, content, and storage methodology of the dataset used in this system.

A. *Dataset Description*

Before generating personalized yoga recommendations, the system requires a structured and well-defined knowledge base that captures the relationship between yoga asanas and their associated health benefits. This foundational layer is essential to ensure that the matching algorithm can identify contextually relevant poses based on user inputs. The underlying dataset plays a critical role in facilitating this process by organizing yoga poses along with their suitability for various demographic factors such as age and gender, as well as health-related considerations.

The dataset used in this project is stored in a MongoDB database and comprises a carefully curated collection of yoga asanas. Each entry in the dataset is represented as a document containing multiple fields that define the pose, its therapeutic value, and instructional metadata. These fields include:

1) s_no: an integer value assigned to uniquely identify each yoga asana entry within the dataset.
2) asana: a string representing the traditional name of the yoga posture, used as a primary reference for identification and display purposes.
3) age: an integer indicating the minimum recommended age for safely performing the pose. this ensures age-appropriate filtering during the recommendation process.
4) gender: defines whether the asana is suitable for male users, female users, or is universally applicable to all users. this field allows for gender-specific personalization.
5) health_benefits: a descriptive text listing the physical or psychological benefits of the asana. this field plays a critical role in the nlp-based semantic matching of user health concerns to appropriate poses.
6) pose_direction: a step-by-step guide detailing how the asana should be performed correctly. these instructions are included in the generated pdf to help users follow proper technique.
7) contraindications: a list of medical or physical conditions for which the pose is not recommended. this helps ensure that users avoid asanas that may pose risks to their health.
8) image_url: a direct link to an external image representing the yoga pose, typically hosted on a cloud service such as cloudinary. these images are dynamically embedded into the pdf recommendations for visual reference.

To maintain high-quality visual output and reduce local server load, all pose illustrations are stored on Cloudinary, a cloud-based media management platform. The system retrieves these images in real-time using the stored URLs and embeds them into a structured PDF that is emailed to the user along with their recommended poses. This strategy enables efficient media delivery while supporting scalability and consistent formatting.

Additionally, the database contains a separate collection of motivational yoga quotes. These quotes are dynamically retrieved and displayed on the web interface, enhancing user engagement through positive reinforcement and mindfulness inspiration. The use of MongoDB as a NoSQL database enables flexible schema design and efficient storage of hierarchical data structures. It also allows the system to scale as new poses, health conditions, or metadata fields are added in the future.

B. *System Architecture*

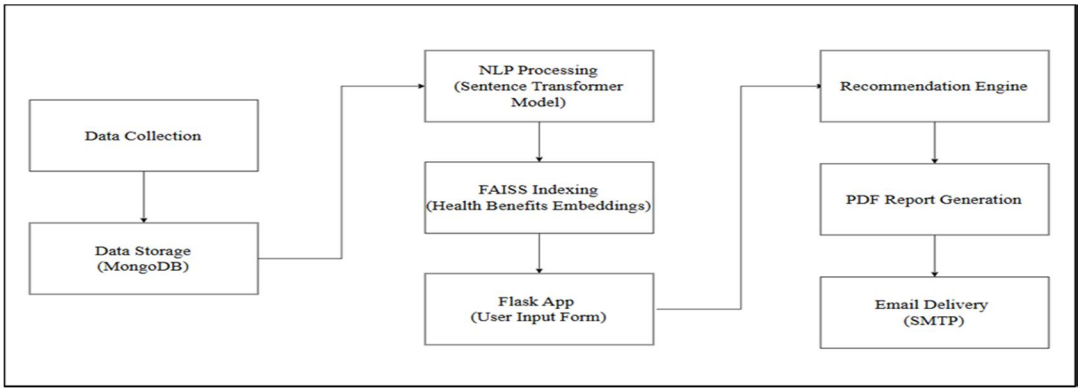

Figure-1: System Architecture of the Smart Yoga Asana Recommender System

The system architecture of the Smart Yoga Asana Recommender is designed to support personalized, AI-driven yoga pose suggestions through a modular and scalable framework. It integrates data collection, semantic processing, similarity search, and report generation in a streamlined workflow.

At the core, the system begins with a data collection phase, where information related to yoga asanas - including pose names, health benefits, instructions, contraindications, age and gender suitability, and image URLs - is curated and stored in a MongoDB database. This forms the knowledge base that the recommendation engine relies on.

When a user provides their input via a web interface - typically age, gender, email, and a brief textual health concern - the input is passed through an NLP processing module. This module utilizes a pre-trained SentenceTransformer model to convert the user's natural language input into a semantic vector representation.

The resulting embedding is then compared with precomputed vectors representing the health benefits of stored asanas. This comparison is performed using FAISS (Facebook AI Similarity Search), which facilitates high-speed, large-scale vector similarity search. The matched health benefits are used to retrieve corresponding yoga asanas from the database.

A recommendation engine filters the selected poses based on user demographics (age and gender) to ensure personalized and appropriate results. These filtered poses are then forwarded to a PDF report generation module, which formats the recommendations into a well-structured document using ReportLab. The report includes pose names, step-by-step instructions, contraindications, and images (fetched from Cloudinary).

Finally, the system delivers the PDF to the user's email via SMTP-based email delivery, completing the automated, end-to-end recommendation process.

C. Methodology

The methodology defines the operational flow of the proposed system, from receiving user input to generating and delivering personalized yoga pose recommendations. The implementation is structured into a series of interconnected stages that collectively provide an intelligent, automated wellness solution.

The process begins with the user accessing a web-based interface developed using Flask. The user is prompted to enter four key details: age, gender, a brief description of their health issue, and their email address. This data is submitted via a structured form and forwarded to the backend server for processing.

The core of the recommendation process starts with the Natural Language Processing (NLP) module, which utilizes a pre-trained SentenceTransformer model (all-MiniLM-L6-v2). This model converts the free-text health description into a high-dimensional semantic vector that captures the contextual meaning of the user's input. Unlike traditional keyword matching, this embedding-based approach ensures that similar health concerns are grouped even if phrased differently.

Next, the semantic vector is passed to a FAISS (Facebook AI Similarity Search) engine. FAISS compares the user's input embedding with precomputed embeddings of known health benefits associated with various yoga asanas. It returns the most semantically similar matches, enabling the system to identify yoga poses relevant to the user's needs.

Once relevant health benefits are identified, the system queries a MongoDB database to retrieve matching asanas. To ensure safety and suitability, additional filtering is applied based on the user's age and gender. Only those poses that align with the user's profile are selected for recommendation.

The selected asanas, which include the name, performance steps, contraindications, and image URLs (hosted on Cloudinary), are compiled into a personalized PDF report. The report is dynamically generated using the ReportLab library, ensuring a structured, professional layout with clear formatting and visual aids.

Finally, the system uses an integrated SMTP mailing service to send the generated PDF to the user's provided email address. The frontend displays appropriate success or error messages to confirm whether the delivery was successful.

This methodology ensures that each user receives a customized, context-aware yoga routine tailored to their demographic and health needs - all without requiring manual intervention or physical consultation.

IV. EXPERIMENTAL ANALYSIS AND RESULTS

A. Key Features

The system introduces several distinct features that enhance both its usability and effectiveness:

1) NLP-Based Semantic Understanding: The health concern provided by the user in natural language is processed using the pre-trained all-MiniLM-L6-v2 SentenceTransformer model, which captures semantic meaning rather than relying on keyword matching.

2) FAISS-Based Similarity Search: The user's input is transformed into a semantic vector and compared with precomputed embeddings of known yoga pose benefits using FAISS (Facebook AI Similarity Search). This allows for rapid identification of relevant poses with high contextual accuracy.
3) Personalized Filtering: The system applies demographic filters (age and gender) to ensure that only suitable poses are recommended to the user, promoting safety and personalization.
4) Cloud-Hosted Visual Content: Each yoga pose includes a Cloudinary-hosted image, which is dynamically embedded in the final report for better visual understanding.
5) Automated PDF Generation and Email Delivery: The recommendations are compiled into a well-structured PDF using ReportLab. The final document includes pose names, instructions, contraindications, and images, and is emailed to the user through a secure SMTP mailing system.
6) User-Friendly Interface: The Flask-based web application features a clean, intuitive UI with real-time form validation and user feedback messages, enhancing the overall experience.

B. Results

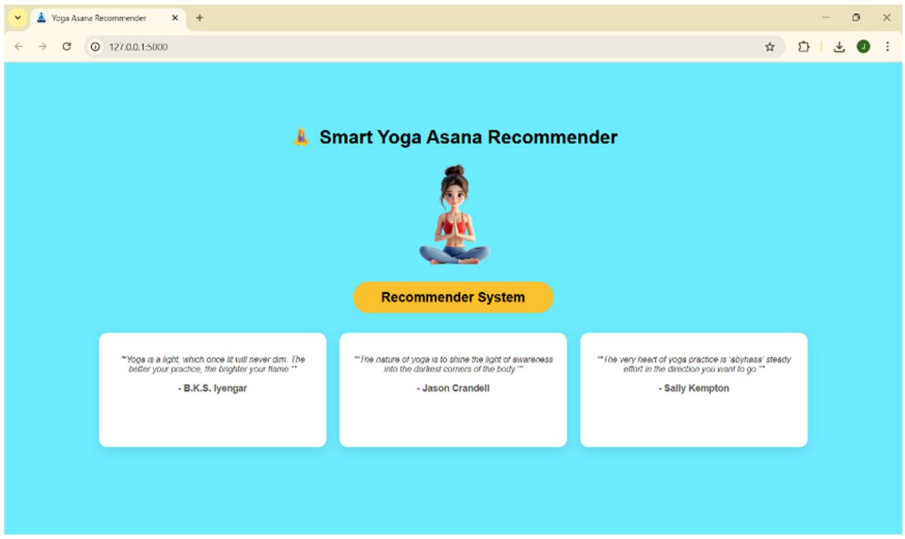

Figure-2: Home Screen of Smart Yoga Asana Recommender System

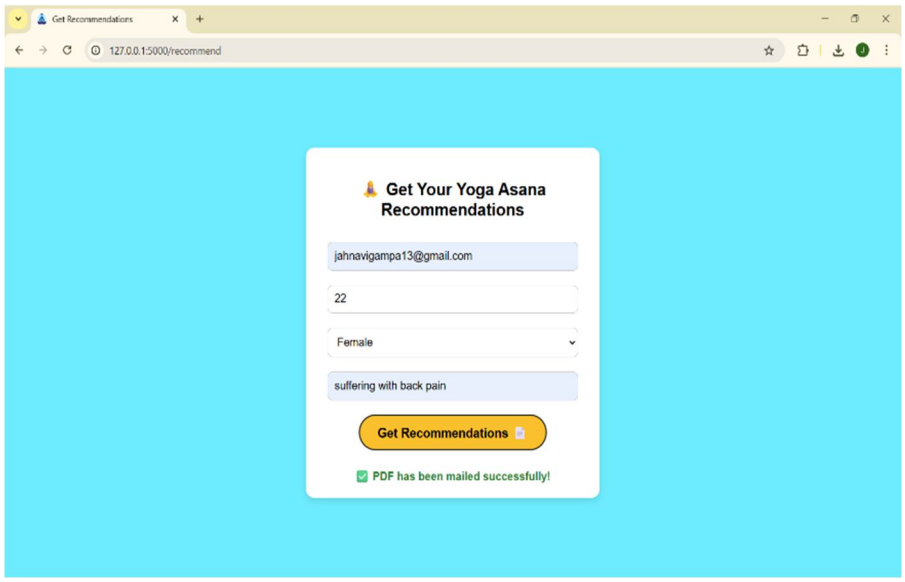

Figure-3: User Input 1

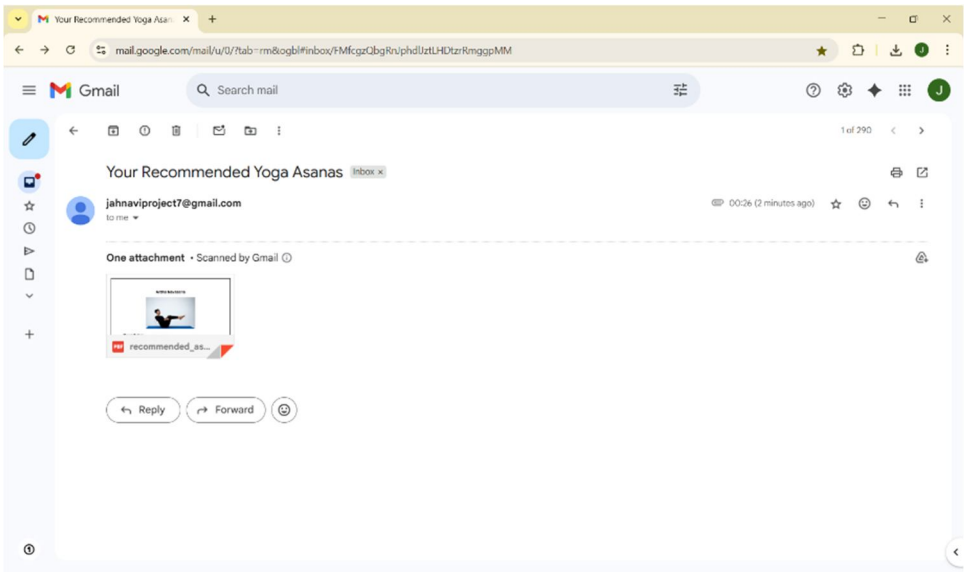

Figure-4: Recommended Asanas sent through Email to the user

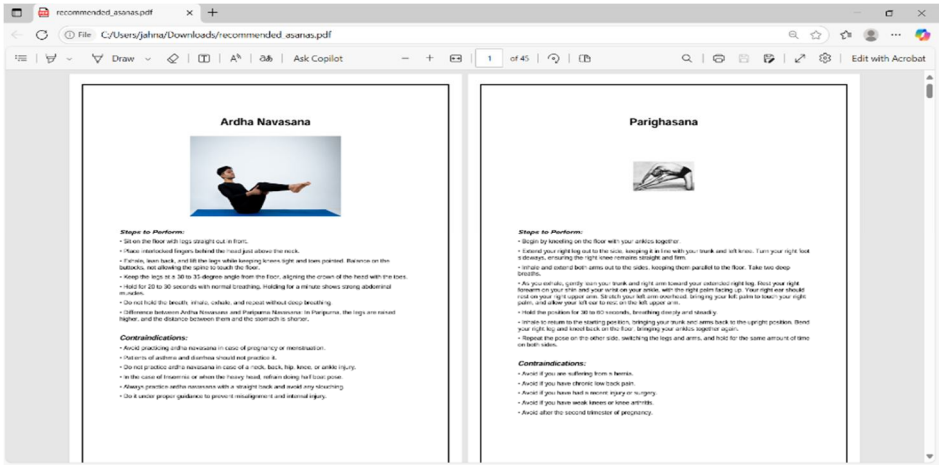

Figure-5: Results PDF for the User Input 1

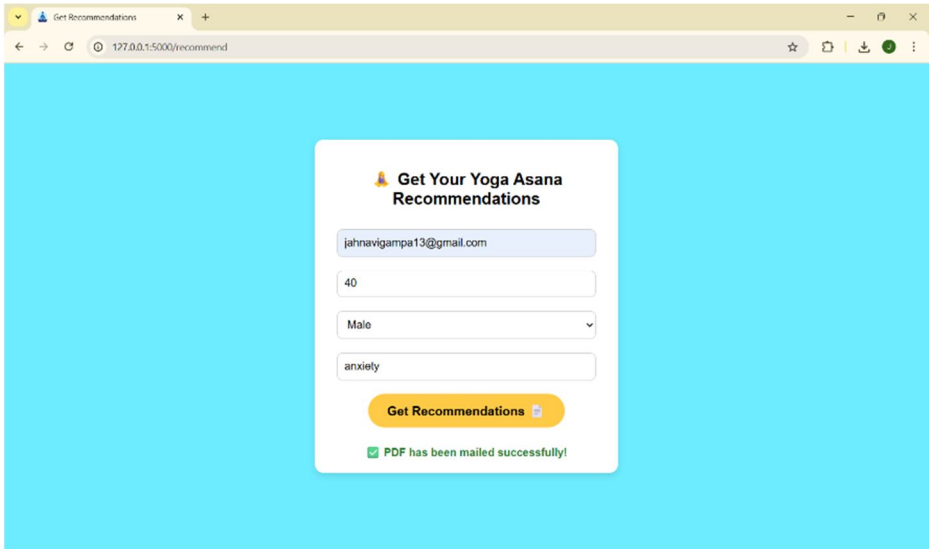

Figure-6: User Input 2

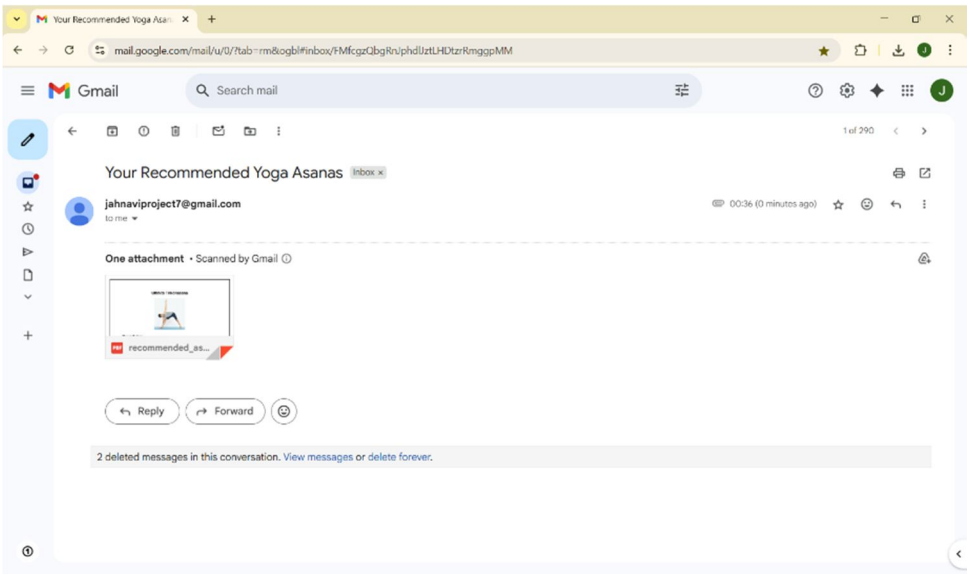

Figure-7: Recommended Asanas sent through Email to the user

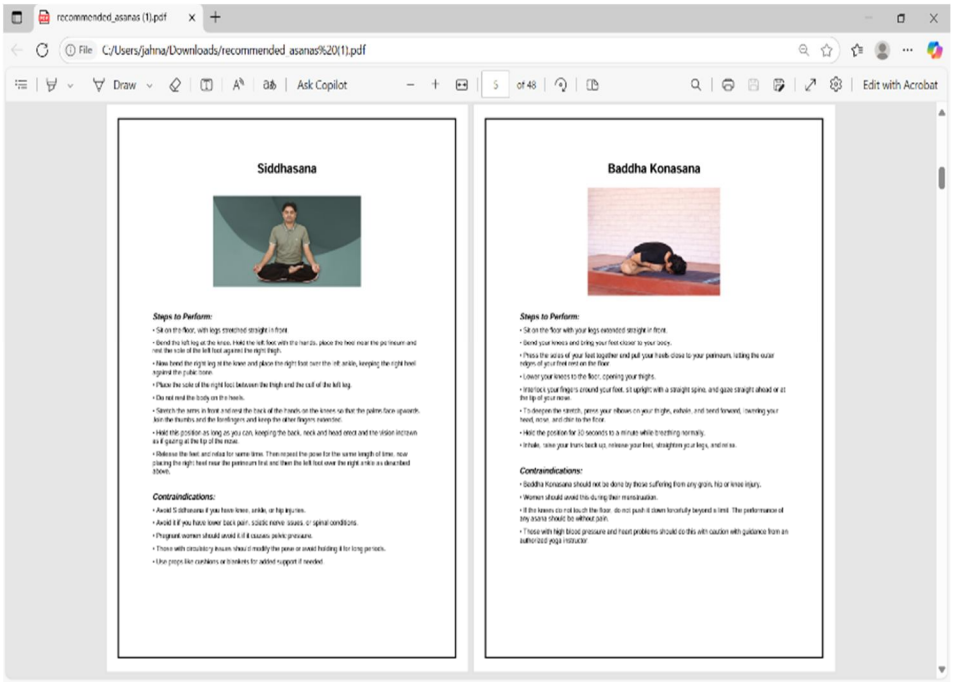

Figure-8: Results PDF for the User Input 2

The Smart Yoga Asana Recommender System was successfully deployed as a web-based application, offering a seamless experience from user input to personalized output delivery. The home screen provided an intuitive entry point, guiding users to submit their health-related concerns through a structured input form. Upon submission, the system processed the input using Natural Language Processing techniques and retrieved the most relevant yoga asanas. The final output was compiled into a professional PDF document, which was automatically emailed to the user.

Two distinct use cases were tested to validate the system's adaptability and accuracy. Each case involved different health issues, age groups, and gender, and the recommendations varied accordingly - demonstrating the system's capability to tailor responses based on individual input. The PDFs received via email contained well-structured yoga suggestions along with corresponding images enhancing user engagement. These results affirm that the system delivers precise, customized recommendations in a user-friendly and automated manner.

V. LIMITATIONS AND FUTURE SCOPE

While the proposed system demonstrates reliable performance in generating personalized yoga recommendations, there are a few limitations that open avenues for future enhancement.

Firstly, the system currently supports input in English only, limiting accessibility for users who may prefer regional or native languages. Expanding support to include multilingual NLP models could significantly broaden its user base. Secondly, the recommendations are based solely on user-provided text and basic demographic data. Integrating real-time health data from wearables - such as heart rate, stress levels, or sleep quality - could enable dynamic and more context-aware suggestions.

Moreover, while the current version uses semantic similarity for matching, it lacks clinical validation or medical certification. Collaborating with certified yoga professionals or healthcare experts can help improve recommendation safety and credibility, especially for users with chronic conditions. Additionally, offering a mobile app version would improve usability, allowing users to access recommendations on the go.

Future iterations could also include features such as voice-based input, progress tracking, asana animations, and real-time correction using computer vision, thereby transforming the system into a more interactive and holistic wellness platform.

VI. CONCLUSION

This research presents the development and implementation of a Smart Yoga Asana Recommender System that leverages NLP, vector similarity search, and modern web technologies to provide users with personalized yoga pose recommendations. By accepting natural language input and combining it with demographic filters, the system intelligently identifies relevant asanas from a curated database, compiles the results into a professionally formatted PDF, and delivers it via email — all within a matter of seconds.

The integration of tools such as SentenceTransformer, FAISS, MongoDB, and Flask allows the system to operate efficiently and scalably. The inclusion of visual content and motivational elements further enhances user engagement.

Overall, this work demonstrates how traditional wellness practices like yoga can be enhanced and scaled through artificial intelligence. It opens up possibilities for building accessible, data-driven health solutions that are both meaningful and easy to use.

REFERENCES

[1] Deb, R. Khakharia, S. Shetty, A. Jadhav, and R. Sethi, "Yoga pose recommender system for physical and mental illnesses," in IEEE Quantum Computing and Emerging Business Technology, ISBN: 979-8-3503-8427-7, 2024.

[2] S. Ramamoorthy, D. Singh, and S. Kar, "Yoga recommendation system for the mental well-being of students using machine learning," in International Conference on Electronics, Communication and Aerospace Technology, ISBN: 978-1-6654-8271-4, 2023.

[3] A. Tom, G. Das, J. B. Hemmige, P. Deshpande, S. R. Bangshi, A. R. Paduri, and N. Darapaneni, "Detecting mental health disorders using NLP," in IEEE Pune Section International Conference, ISBN: 979-8-3503-2420-4, 2024.

[4] V. Arya, N. Makattil, V. Sasikumar, V. Anuparvathi, and S. Khandare, "Know your posture: Real-time posture detection and correction with yoga and exercise recommendations," in International Conference on Signal Processing, Computation, Electronics, Power and Telecommunication, ISBN: 979-8-3503-1212-6, 2023.

[5] Nikitha, P., Nishanthini, S., Vinod, S., Mouleesh, Y., Silviya, S. H. A., & Sriman, B. “Intelligent Document Interaction with Advanced Vector Embeddings and FAISS-CPU Indexing,” Proceedings of the 2024 8th International Conference on Electronics, Communication and Aerospace Technology (ICECA), IEEE, Nov. 6–8, 2024. DOI: 10.1109/ICECA63461.2024.10800948.

[6] Kaur, Jagdish, and Kamboj, C. P. “Punjabi Grammar Checker: Automatic Correction of Punjabi Language Sentences Using Transformer,” Proceedings of the 2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT), IEEE, June 24–28, 2024. DOI: 10.1109/ICCCNT61001.2024.10725698.

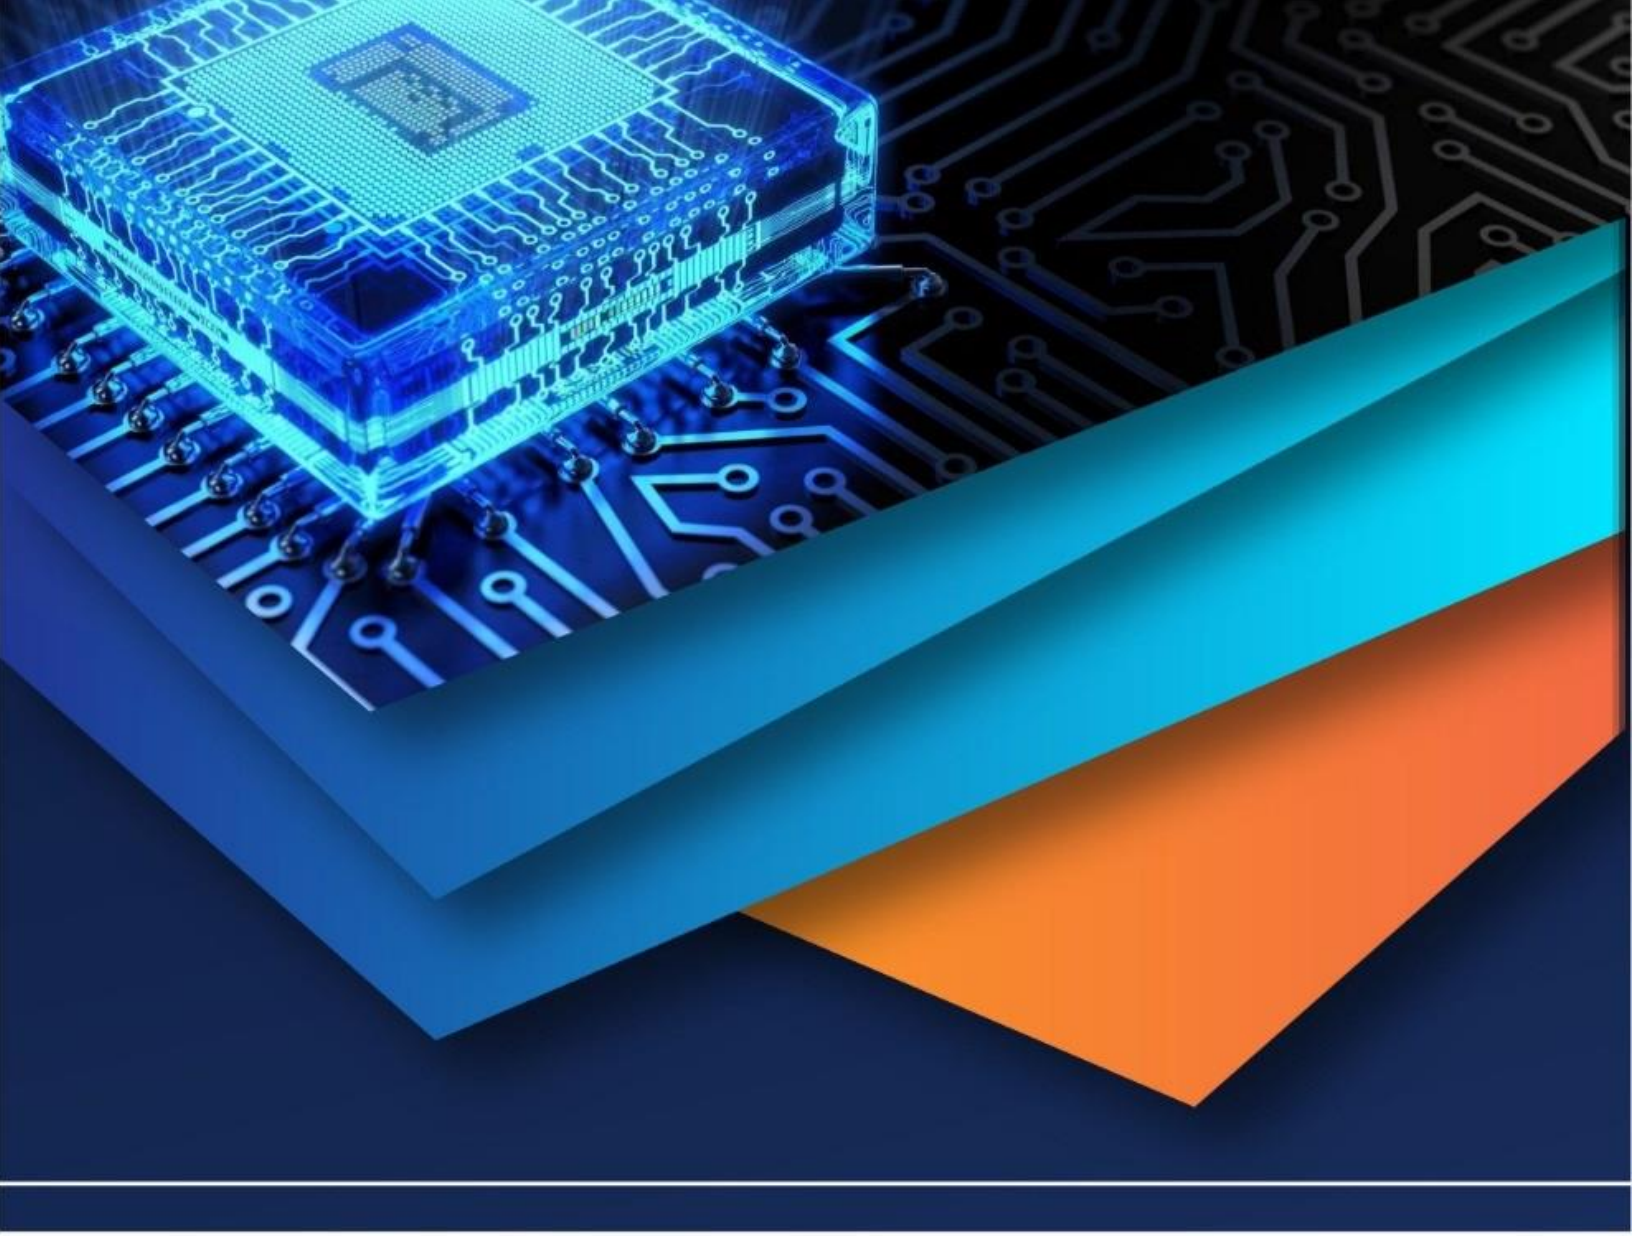

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY