

Smart AI Healthcare Assistant For Skin, Hair and Health Awareness

Srushti Wayase, Namrata Ransing, Ankita Muluk, Ms. Sejal Lunawat, Siddhi Kadam

Computer Engineering

Pimpri Chinchwad Polytechnic, Pune

srushtiwayase@gmail.com, namrataransing15@gmail.com, ankitamuluk8@gmail.com

sejallunawat01@gmail.com, kadamsiddhi2310@gmail.com

Abstract: In the post-pandemic era, the demand for accessible and immediate healthcare information has surged. This paper proposes a "Smart AI Healthcare Assistant," a web-based conversational agent designed to provide preliminary awareness and advice regarding general health, skin care, and hair care. The system utilizes a hybrid Natural Language Processing (NLP) approach, combining deterministic keyword matching with fuzzy logic for common symptoms (e.g., acne, hair fall) and a pre-trained Transformer model (RoBERTa) for contextual question answering. The application is built using a decoupled architecture with a React.js frontend for an interactive user interface and a Flask-based backend API. Key features include symptom detection, immediate home remedy suggestions, and location-based recommendations for nearby clinics. This project demonstrates how lightweight AI models can be integrated into web applications to democratize basic health awareness.

Keywords: Healthcare Chatbot, Natural Language Processing (NLP), React.js, Flask, RoBERTa, Transformers, Fuzzy Logic

I. INTRODUCTION

In the rapidly evolving landscape of digital health, Artificial Intelligence (AI) has emerged as a transformative force, capable of bridging the gap between medical professionals and patients. With the increasing reliance on the internet for self-diagnosis, individuals often encounter conflicting or alarming information when searching for symptoms online. This phenomenon highlights a critical need for intelligent, verified, and accessible automated systems that can provide preliminary health assessments. This research presents the "Smart AI Healthcare Assistant," a web-based conversational agent specifically tailored to address general health, dermatological (skin), and trichological (hair) concerns. Unlike static medical portals, this system offers an interactive dialogue interface, mimicking the triage process of a preliminary medical consultation to provide immediate, actionable guidance. The domain of skin and hair care is particularly suited for such technological intervention, as many common conditions—ranging from acne and dry skin to hair fall and dandruff—are often linked to lifestyle factors, hygiene, and diet. Patients frequently face long waiting times or high consultation fees for minor issues that could initially be managed with home remedies or over-the-counter advice. The proposed system addresses this inefficiency by acting as a first-line support tool. It is designed to analyze user queries regarding specific symptoms found in its knowledge base, such as high fever or sunburn, and deliver scientifically grounded advice. Furthermore, the system recognizes the limits of automated advice; it incorporates a logic-based referral mechanism that identifies when a condition is severe (e.g., persistent hair loss or high fever) and proactively suggests nearby clinical facilities using geolocation services.

Technologically, the project demonstrates the effective integration of modern web development frameworks with state-of-the-art Natural Language Processing (NLP). The application employs a **React.js** frontend to ensure a responsive, user-friendly experience across devices, while the backend is powered by **Flask**, a lightweight Python framework. To ensure high accuracy in understanding user intent, the system utilizes a hybrid intelligence model. It combines deterministic **fuzzy logic** algorithms to handle common typographical errors and direct symptom matching with the **RoBERTa (Robustly Optimized BERT approach)** transformer model. This dual approach allows the chatbot to



answer specific health questions based on context with high confidence, thereby democratizing access to reliable health awareness and reducing the anxiety associated with medical uncertainty

Ease of Use

The "Smart AI Healthcare Assistant" prioritizes accessibility and simplicity, recognizing that users seeking health advice may be in a state of anxiety or urgency. The user interface (UI) is designed to mimic popular instant messaging platforms, ensuring a near-zero learning curve for any user familiar with standard chat applications. This familiarity is reinforced by a clean, minimalist visual design that avoids clutter, focusing the user's attention solely on the dialogue and advice. To further enhance usability, the system supports standard interaction patterns, such as submitting queries via the "Enter" key and automatically scrolling to the most recent message, ensuring that the conversation flow remains uninterrupted and easy to track.

Beyond the visual interface, the system incorporates specific technical features to mitigate user frustration and ensure continuity. A persistent chat history mechanism, implemented via browser local storage, saves the conversation state automatically. This allows users to refresh the page or return to the application later without losing the context of previous advice or recommendations. Additionally, the application includes visual feedback mechanisms, such as a "typing" animation during data fetching states, which assures the user that their query is being processed, thereby reducing abandonment rates caused by uncertainty about system responsiveness.

Crucially, the system is engineered to be forgiving of user errors, a vital attribute for an ease-of-use focused application. The backend integration of fuzzy logic algorithms allows the chatbot to interpret and correctly respond to inputs containing typographical errors or colloquialisms (e.g., interpreting "fevr" as "fever"). This eliminates the need for users to use precise medical terminology or perfect spelling to receive accurate help. Furthermore, the system simplifies complex actions into single clicks; for instance, when professional help is recommended, the system renders a direct "Show Clinics" button, removing the friction of manually searching for doctors and streamlining the transition from digital advice to physical care.

II. LITERATURE SURVEY

The Evolution of Automated Health Consultation

The domain of personal health monitoring and preliminary diagnosis has evolved significantly in recent years, driven by the digital transformation of the healthcare sector. Historically, medical triage was a manual, resource-intensive process reliant on human intuition, physical availability, and face-to-face interaction. However, as the global patient-to-doctor ratio has widened and the demand for immediate health information has surged, traditional methods have proven difficult to scale. This has necessitated a shift toward data-driven automated systems capable of processing natural language to provide accessible health awareness. The transition from static medical portals to interactive conversational agents represents a fundamental paradigm shift, aiming to democratize access to high-quality health advice regardless of geographical or financial constraints.

Dominance of Transformer Models in Medical NLP

A substantial portion of existing research has focused on utilizing deep learning architectures, specifically Transformer models, to automate medical Question Answering (QA). Researchers such as Devlin et al. (BERT) and Liu et al. (RoBERTa) have demonstrated that attention-based mechanisms are highly effective for extracting precise answers from complex textual contexts. These studies typically utilize vast repositories of medical literature—analyzing context, semantics, and intent—to identify relevant advice for user queries. Similarly, recent advancements in models like BioBERT have proven that pre-training on biomedical corpora significantly improves performance in clinical tasks, proving that statistical analysis of language patterns is a strong predictor of accurate medical information retrieval.

The "Black Box" Limitation and Safety Concerns

Despite their linguistic capabilities, these pure Deep Learning approaches often suffer from the "black box" problem, a critical limitation in the context of healthcare safety. In complex Large Language Models (LLMs), the internal logic used to generate a medical response is often opaque to the user and sometimes even to the developers. For a user seeking advice on symptoms, simply receiving a generated text is insufficient; they require assurance that the advice is



grounded in verified medical facts rather than statistical probability. Furthermore, these generative models can be prone to "hallucinations," where they confidently provide incorrect medical dosages or non-existent remedies, leading to potentially dangerous recommendations based on probabilistic anomalies rather than established medical protocols.

Computational Constraints and Deployment Challenges

A further challenge with state-of-the-art Large Language Models (LLMs) is their heavy reliance on massive computational resources and structured historical datasets. For these models to function effectively in real-time, they require significant GPU power and cloud infrastructure, which creates a significant barrier to entry for lightweight web applications or patient-side mobile tools. This dependency renders many research-grade models difficult to deploy in real-world scenarios, particularly for personal health assistants running on standard web browsers or local servers. This highlights the need for architectures that decouple the heavy inference engine from the user interface, utilizing lightweight frameworks like Flask and React to ensure responsiveness without sacrificing intelligence.

The Role of Expert Systems and Hybrid Logic

Parallel to pure machine learning approaches, earlier research explored the use of Expert Systems (ES) and heuristic methods that mimic clinical decision-making through predefined rules. Unlike the statistical probability used in Transformers, these systems rely on deterministic logic. Researchers have utilized fuzzy logic and string-matching algorithms (such as Levenshtein distance) to handle the inherent uncertainty and typos in user descriptions of symptoms. These rule-based systems offer a distinct advantage: they are interpretable, transparent, and safe, qualities that are crucial when providing actionable feedback regarding skin or hair conditions. By combining deterministic rules for known symptoms with ML models for unstructured queries, hybrid systems can achieve both safety and flexibility.

Addressing the Holistic Gap in Tele-Dermatology

Beyond the technical methodology of symptom detection, a significant body of literature addresses the "Holistic Gap" in tele-dermatology. Systematic reviews have identified that effective treatment for conditions like acne or hair fall requires not just medical diagnosis, but also "soft" advice regarding lifestyle, diet, and hygiene. However, most academic prediction models fail to integrate these parameters, focusing solely on image classification of diseases. This oversight results in a disconnect between diagnosis and cure, as a user may know they have dandruff but lack the practical knowledge to manage it. Systems that can bridge this gap—providing both home remedies for management and location-based suggestions for professional clinics—are essential for reducing the burden on the primary healthcare system.

III. METHODOLOGY

The implementation of the Smart AI Healthcare Assistant utilizes a decoupled architecture, leveraging React.js for the frontend interface and Python Flask for the backend logic. This separation of concerns allows the user interface to remain lightweight and responsive while the server handles the computational load of the Natural Language Processing (NLP) tasks. The backend is initialized as a Flask application, which exposes a single RESTful endpoint, /chat, to handle incoming POST requests. To ensure secure and seamless communication between the browser-based client and the API, the system employs the flask_cors extension, allowing the React frontend to exchange JSON data with the server without violating browser security policies.

The core intelligence of the system is implemented in app.py through a two-stage processing pipeline. The first stage utilizes a deterministic knowledge_base, a JSON structure containing specific advice and synonyms for conditions such as hair fall, acne, and fever. To maximize the system's usability, the difflib library is employed to perform fuzzy string matching. This algorithm calculates the similarity between the user's input and the stored keywords; if a match with a confidence score greater than 0.7 is found, the system immediately returns the pre-defined advice. This approach ensures that common typos (e.g., "fevr" instead of "fever") do not prevent the user from receiving accurate assistance. If no keyword match is found, the system activates the second stage: a Hugging Face transformers pipeline using the deepset/roberta-base-squad2 model. This model analyzes the user's natural language query against the context of the medical advice to extract the most relevant answer.

The frontend is built as a functional React component in App.js, utilizing hooks to manage the application state. The useState hook is used to maintain the chat history, which is simultaneously persisted to the browser's localStorage to



ensure conversation continuity across sessions. Communication with the backend is handled asynchronously using axios; while the system awaits a response, a loading state triggers a visual indicator to reassure the user that their query is being processed. A critical feature of the implementation is the renderMessage function, which parses the text returned by the bot. If the bot's response includes the trigger phrase "Should I suggest a nearby clinic," the application dynamically renders an interactive HTML link instead of plain text. This link is constructed to redirect the user to a Google Maps search query relevant to the detected disease, effectively bridging the gap between digital diagnosis and physical care.

The performance of the Smart AI Healthcare Assistant was evaluated through a series of test scenarios designed to assess both the rule-based and deep-learning components of the architecture. In initial tests focusing on the knowledge base, the system demonstrated 100% accuracy in retrieving advice for exact keyword matches such as "hair fall," "acne," and "sunburn." The response time for these queries was negligible, validating the efficiency of using a dictionary to lookup for common symptoms. The fuzzy logic engine was subsequently tested with intentionally misspelled inputs, such as "dry skn" and "hair loss." The system successfully mapped these inputs to their correct definitions ("dry skin" and "hair fall" respectively), confirming that the similarity threshold of 0.7 is effective for handling standard typographical errors without generating false positives.

The NLP capabilities were tested by inputting natural language questions rather than isolated keywords. When presented with queries like "What should I do for a high temperature?", the RoBERTa model successfully parsed the context related to "fever" and extracted the correct advice regarding rest and paracetamol. This confirms the system's ability to handle unstructured queries that do not strictly match the keyword list. However, the system also correctly displayed a fallback message ("I'm not fully sure about that") when the model's confidence score dropped below 0.2, demonstrating a safety mechanism that prevents the generation of unreliable advice for out-of-scope topics.

Finally, the clinic recommendation feature was tested by simulating severe symptoms. When the user input triggered the advice for conditions linked to the clinic logic, the backend successfully appended the suggestion text: "Should I suggest a nearby clinic?". The React frontend correctly intercepted this message and rendered the "Yes, show clinics" button. Clicking this button successfully opened a new browser tab with a location-based search for the specific condition (e.g., dermatologists or general physicians), proving the functional viability of the system as an end-to-end triage tool.

IV. CONCLUSION

The development of the "Smart AI Healthcare Assistant" illustrates the significant potential of combining deterministic and probabilistic AI models to create reliable, accessible health awareness tools. By integrating a responsive React.js frontend with a lightweight Flask backend, the system successfully bridges the gap between static health information and interactive medical triage. The research validates the efficacy of a hybrid Natural Language Processing (NLP) architecture; the use of fuzzy logic ensures robust symptom detection even in the presence of typographical errors, while the integration of the deepset/roberta-base-squad2 Transformer model allows for nuanced, context-aware responses to unstructured user queries. This dual approach optimizes computational efficiency without sacrificing the accuracy required for medical advice.

Furthermore, the system effectively addresses the "actionability gap" often found in health chatbots. By not only providing home remedies but also dynamically generating location-based clinic referrals for severe conditions, the application serves as a complete end-to-end solution. It empowers users to manage minor issues like dry skin or dandruff at home while promptly guiding them to professional care for more serious symptoms such as high fever or persistent hair fall. This functionality transforms the chatbot from a passive information retrieval system into an active participant in the user's healthcare journey, potentially reducing unnecessary hospital visits while ensuring timely intervention when needed.

Looking forward, the scope of this project can be expanded to incorporate multimodal diagnostic capabilities. Future iterations aim to integrate Convolutional Neural Networks (CNNs) to allow users to upload images of skin conditions for visual analysis, thereby increasing diagnostic precision. Additionally, to enhance accessibility for visually impaired or elderly users, the integration of Speech-to-Text (STT) and Text-to-Speech (TTS) APIs is proposed. Finally, the



practical utility of the platform could be further elevated by integrating real-time appointment scheduling APIs, allowing users to book consultations directly with the recommended clinics without leaving the chat interface.

ACKNOWLEDGMENT

We extend our sincere appreciation to our mentor, Prof. Sejal Lunawat, from the Department of Computer Engineering at Pimpri Chinchwad Polytechnic. Her technical expertise and mentorship were vital to this project's success, particularly her insights regarding the Weighted Scoring Algorithm and the system's architecture.

We also wish to acknowledge the Head of the Computer Engineering Department and the Principal of Pimpri Chinchwad Polytechnic. We are grateful for their support in providing the essential laboratory infrastructure and fostering an environment conducive to academic growth.

REFERENCES

- [1]. arXiv | RoBERTa: A Robustly Optimized BERT Pretraining Approach | Yinhua Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, Veselin Stoyanov | 2019
- [2]. Association for Computational Linguistics (NAACL) | BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding | Jacob Devlin, Ming-Wei Chang, Kenton Lee, Kristina Toutanova | 2019
- [3]. Journal of Medical Internet Research (JMIR) | Conversational Agents in Health Care: Scoping Review and Conceptual Analysis | Liliana Laranjo, Adam G Dunn, Huong Ly Tong, Ahmet Baki Kocaballi, Jessica Chen, Rabia Bashir, Didi Surian, Blanca Gallego, Farah Magrabi, Annie YS Lau | 2018
- [4]. Association for Computational Linguistics (ACL) | Know What You Don't Know: Unanswerable Questions for SQuAD | Pranav Rajpurkar, Robin Jia, Percy Liang | 2018
- [5]. IEEE Access | Chatbot as Automated Conversational Agent: A Survey | Bayu Setiaji, Ferry Wahyu Wibowo | 2016