International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 03 Issue: 04 April 2025
Page No: 1900-1903
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0276

AI Chatbot for Enhancing Mental Health

Mr. N. Manimozhiyan[1], Mr. U. Arulkumar[2], Mr. D. Pon Ezhilan[3], Mr. R. Vinoth Kumar[4]
[1,2,3]UG Scholar, Dept of IT, Rajiv Gandhi College of Engineering and Technology, Puducherry, India.
[4]Asst. Professor, Dept of IT, Rajiv Gandhi College of Engineering and Technology, Puducherry, India.
Emails: *manimozhiyan07@gmail.com[1], arulkumar6135@gmail.com[2], ponezhilan1@gmail.com[3], kumarvinothr79@rgcet.edu.in[4]*

Abstract
This project introduces a healthcare chatbot designed to improve mental well-being through mood-responsive interactions. The system integrates a Large Language Model (LLM)-powered contextual understanding, enabling adaptive and personalized responses to user input. By analyzing users' language patterns, the chatbot adapts its responses to provide context-aware, empathetic support that aligns with the user's current emotional state. Built on the Rasa framework with Natural Language Understanding (NLU) capabilities, this chatbot offers personalized, therapeutic conversations aimed at alleviating symptoms of mental health issues like stress and loneliness. This healthcare tool shows potential for enhancing user engagement and supporting mental health recovery by offering a more personalized and language-sensitive approach to care.
Keywords: *Mental Health Chatbot; AI in Healthcare; RASA; Natural Language Understanding; Natural Language Processing; Large Language Model.*

1. Introduction

Mental disorders such as stress, anxiety, and loneliness are on the rise, especially among individuals with limited access to professional care. The majority of individuals do not access care due to social stigma, lack of knowledge, or the lack of mental health services. Traditional therapy is usually costly, time-consuming, and geographically limited and hence inaccessible to everyone for timely support. Moreover, there is no easily available platform that can provide empathetic interaction and emotional support in real-time. This highlights the need for new solutions to bridge the gap in mental well-being. According to The National Mental Health Survey (NMHS) of India 2015-16 approximately 10% of the Indian population was found to be affected by common mental disorders, including depression, anxiety, and substance use disorders. There was a significant treatment gap, with 70-92% of individuals with mental disorders not receiving adequate treatment [1]. Rasa is an open-source conversational AI platform used for building intelligent and customizable chatbots. For mental health applications, Rasa has some advantages. It has robust Natural Language Understanding (NLU) for intent detection and extraction of emotions or mental states from conversation. Through Rasa Core, it seamlessly handles dialogue flow, providing empathetic and context-aware responses essential for users with mental health issues. Rasa's ability to monitor conversation context and facilitate real-time interactions makes it perfect for building sensitive, secure, and responsive mental health care systems [2] [3]. The integration of Large Language Models (LLMs) and Rasa enhances the capacity of mental health chatbots with the application of advanced natural language understanding and generation capabilities. LLMs are capable of interpreting user inputs accurately, which leads to richer and empathetic conversation required in mental health support [5] [6] [7]. Rasa provides an opportunity for integration of LLMs for intent classification and paraphrasing responses in default, allowing for context-based and human-like conversation. By the integration of LLMs, chatbots can facilitate advanced conversation, customized support, and user engagement, thus improving the overall efficiency of

e ISSN: 2584-2137
Vol. 03 Issue: 04 April 2025
Page No: 1900-1903
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0276

mental health interventions [4]. Addressing the mental health support gap requires innovative approaches where empathy is joined with technology. Through the use of LLMs and NLU of Rasa, chatbots are capable of giving bespoke, contextual, and emotionally good conversational interactions. Such integration does deliver real-time, accessible, and stigma-free mental health assistance empowering users to effectively maintain their emotional management and providing aid for mental health.

2. Methodologies and Approaches

In the proposed system a context-aware mental health chatbot is built using the Rasa framework with a fine-tuned Large Language Model (LLM) to provide empathetic conversational care. The system is deployed in an interactive cloud environment and publicly hosted for user access.

2.1.Dataset and Model Fine-Tuning

A mental health-focused dataset, including user questions, therapist-like answers, emotional signals, and conversational history, was collected and preprocessed. The dataset was applied in fine-tuning a LLaMA-2 7B model with supervised fine-tuning methods [7]. The goal was to align the model generation with mental health support goals, such as empathy, validation, and safe guidance.

2.2.Hosting and Integration of LLMs

The fine-tuned model was used in a Flask-based API that governs LLM response generation [8]. The local API is executed on a cloud-based interactive compute environment, providing scalability and accessibility. Public access is provided using Ngrok, which tunnels the local Flask server to a publicly accessible URL [9].

2.3.Rasa Framework Setup

The Rasa framework is the backbone of the chatbot that takes care of Natural Language Understanding (NLU) and dialogue management:

Rasa NLU translates user input to identify intents and extract relevant contexts of interest (e.g., "anxiety", "feeling alone"). Rasa Core manages the dialogue flow and determines what to do from conversation history and predicted intent. Custom actions are established to route certain user intent towards the hosted large language model endpoint via the Flask API.

2.4.Action Endpoint and Communication Mechanism

A custom Rasa action is implemented to send user queries to the LLM Flask server on fallback, and subsequently return the generated responses to Rasa. The custom action code (actions.py) utilizes the public Ngrok URL to enable communication with the in-running LLM model. The Rasa action server (rasa run actions) is initiated to process such API requests in real-time.

2.5.User Interface Using Streamlit

A web-based interactive and simple interface was created using Streamlit [10]. The interface enables one to: Enter their issues. Obtain LLM-generated responses via the Rasa pipeline Have a rational, human-like conversation Streamlit establishes an integration with Rasa's HTTP API (rasa run --enable-api) and a custom Ngrok endpoint, allowing for a full conversational exchange between the user, Rasa, and the LLM. (Figure 1)

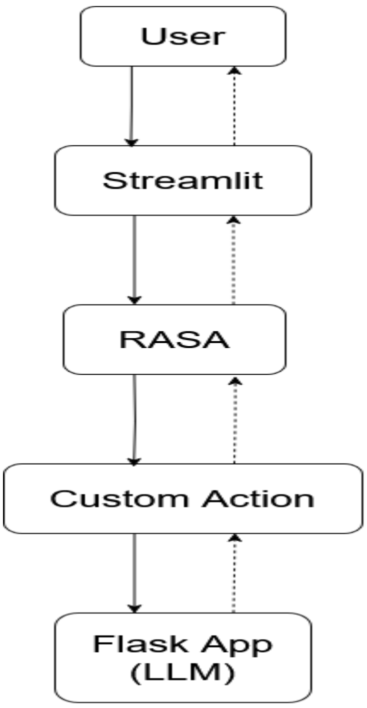

Figure 1 Architecture Diagram of Proposed System

3. Results and Discussion

3.1.Results

The study result shows the effectiveness of the mental

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 03 Issue: 04 April 2025
Page No: 1900-1903
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0276

health chatbot, which combines Rasa's Natural Language Understanding (NLU) with the response generation ability of a Large Language Model (LLM). The chatbot was used to help users with mental health problems and was evaluated on a dataset that replicate real world conversations.

3.2.Intent Classification

Rasa's NLU module was tasked with correctly identifying user intent from pre-defined training data. It was connecting user input with intents like stress_help, low_mood, negative_thoughts and so on. This strict intent recognition allowed the dialogue system to guide conversations accordingly.

3.3.Response Generation

When an intent was determined, response generation was delegated to the combined LLM. In contrast to static templates, the LLM produced empathetic, context-related responses specific to the user's mood and query. This injected a touch of humanness in conversations, critically important in use cases involving mental health.

3.4.Discussion

The evaluation of chat systems for use in mental health applications is a particular challenge since typical lexical-overlap metrics are often weak predictors of semantic and emotional response appropriateness. In the current study, we employed BERTScore to estimate the chatbot's output, averaging F1 score at 16.78% [11]. While this score might appear relatively low, it reflects the model's ability to produce semantically aligned responses despite lexical variation—a central strength in emotionally intense conversation where empathy and contextual suitability are more essential than surface similarity. Despite this, BERTScore alone cannot identify the critical human-oriented features, such as emotional tone, safety in dialogue, and therapeutic appropriateness. To address these limitations, future research needs to include human-in-the-loop evaluation techniques, such as empathy and safety rating, and domain-specific specialized standards. Further, the application of Reinforcement Learning with Human Feedback (RLHF) offers a potential solution to further fill the gap between the system responses and user expectations and standards in mental health communication [12]. These changes are vital to the development of chatbot systems that not only are semantically accurate but also emotionally intelligent and ethically accurate. (Figure 2)

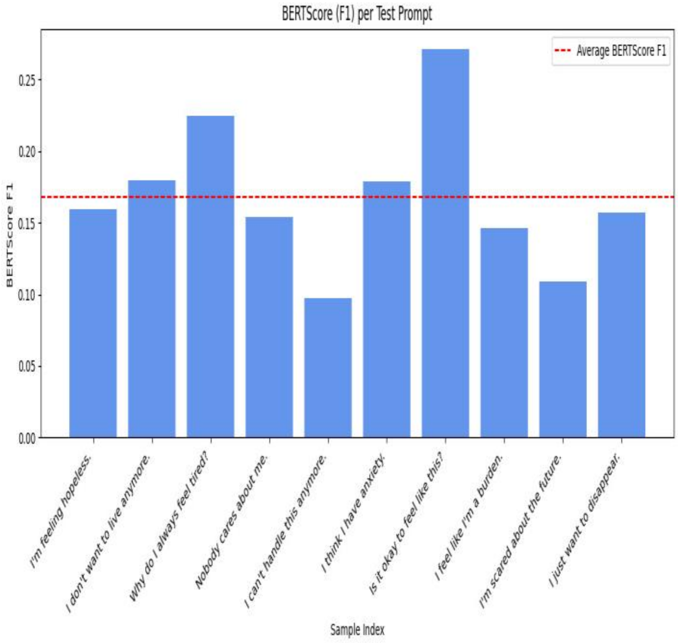

Figure 2 BERT Score

Conclusion

The design of the mental health chatbot on Rasa framework with Large Language Model (LLM) offers personalized and emotionally aware conversations. This is due to the combination of Rasa's capabilities in intent recognition and the natural response generation of LLMs. The chatbot is made accessible by a web-based interface using Streamlit and is supported by a backend running on a Flask API. While the system performs well in generating context-sensitive replies, it has some limitations. Current evaluation methods, such as BERTScore, are not well-suited to measure empathy, safety, or the overall emotional appropriateness of responses—elements that are crucial in mental health communication. To improve, future work should include human-centered evaluation approaches, such as rating empathy and conversational safety, and applying reinforcement learning from human feedback. These changes will help to bridge the gap between technical performance and emotional intelligence—essential for mental health applications.

References

[1]. Gururaj, G., Varghese, M., Benegal, V., Rao, G. N., Pathak, K., Singh, L. K., & NMHS Collaborators. National Mental Health Survey of India, 2015–16: Summary. Bengaluru: NIMHANS Publication. https://indianmhs.nimhans.ac.in/phase1/Docs/Summary.pdf

[2]. Bocklisch, T., Faulkner, J., Pawlowski, N., & Nichol, A. (2017). Rasa: Open source language understanding and dialogue management (arXiv:1712.05181). arXiv. https://doi.org/10.48550/arXiv.1712.05181

[3]. “Rasa Documentation.” Rasa. https://rasa.com/docs/rasa/

[4]. Dolas, S., Pendam, S., Kulkarni, U., & Kadam, A. S. MindLift – A conversational AI chatbot for enhancing mental health. In 2024 MIT Art, Design and Technology School of Computing International Conference (MITADTSoCiCon) (pp. 1–6). IEEE. https://doi.org/10.1109/MITADTSoCiCon60330.2024.10575017

[5]. Wei, J., Tay, Y., Bommasani, R., Raffel, C., Zoph, B., Borgeaud, S., ... & Le, Q. Emergent abilities of large language models (arXiv:2206.07682). arXiv. https://doi.org/10.48550/arXiv.2206.07682

[6]. Rasa. “Using Large Language Models in Rasa.” Rasa Documentation.

[7]. https://rasa.com/docs/rasa/next/llms/large-language-models/

[8]. Meta AI. LLaMA 2: Open Foundation and Fine-Tuned Chat Models. https://www.llama.com/llama2/

[9]. Pallets. Flask API Documentation. https://flask.palletsprojects.com/en/stable/api

[10]. Ngrok. Ngrok Documentation. https://ngrok.com/docs/

[11]. Streamlit. Streamlit Documentation. https://docs.streamlit.io/

[12]. Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q., & Artzi, Y. (2020). BERTScore: Evaluating text generation with BERT (arXiv:1904.09675). arXiv. https://doi.org/10.48550/arXiv.1904.09675

[13]. Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, Paul Christiano. "Learning to summarize from human feedback." Advances in Neural Information Processing Systems (NeurIPS), 2020. https://arxiv.org/abs/2009.01325