TalkThru: An AI-Powered Mental Health Companion for Real-Time Emotional Support

Anjitha Divakaran[1], Hyung-Jeong Yang[1*], Seungwon Kim[1], Ji-eun Shin[2], Soo-Hyung Kim[1]
[1]Dept. of AI Convergence, Chonnam National University
[2]Dept. of Psychology, Chonnam National University
*Corresponding author: hjyang@jnu.ac.kr

ABSTRACT

Emotional connection is essential in mental health support, especially in text-based applications where users seek empathy and validation. TalkThru is an AI-powered mobile app designed to provide real-time emotional assistance through personalized conversations and self-care tools. By combining DistilBERT for emotion detection with T5 for empathetic response generation, the system ensures context-aware interaction that resonates with user emotions. A dynamic avatar, which evolves with the user's emotional journey, adds visual engagement. The app also includes journaling, mood tracking, meditation, and music therapy, supported by gamification to encourage consistency. Developed using Kotlin and Flask, TalkThru maintains performance and scalability while securing sensitive emotional data. Initial evaluations show promising accuracy in emotion recognition and user satisfaction, positioning TalkThru as a practical, AI-driven solution for digital mental health care.

1. INTRODUCTION

Mental health issues are increasing across all age groups, yet access to professional support remains limited due to stigma, cost, and availability of services [1]. As a result, individuals often turn to digital alternatives, such as mobile health applications, to seek emotional guidance and relief. However, many of these solutions lack genuine empathy and personalization, often offering scripted or generic responses that do not build trust or emotional safety [2].

Several chatbot-based mental health applications have emerged in recent years, including Woebot, Wysa, and Replika. While these tools offer round-the-clock accessibility and some degree of conversational flow, they are often limited in emotional intelligence, contextual understanding, and visual engagement [3]. Their reliance on rule-based systems or constrained memory makes them insufficient for delivering sustained, emotionally aware support.

Recent research has explored more advanced empathy-aware systems. For example, the EMPATHETIC DIALOGUES dataset introduced over 25,000 emotionally grounded conversations for training dialogue systems, significantly improving perceived empathy in responses [4]. The EPITOME framework categorizes empathy into emotional reactions, interpretations, and explorations, using RoBERTa to detect each mechanism [5]. Other studies have applied reinforcement learning to enhance empathy in real-time, including Partner for empathy rewriting [9], PAL, an emotion-adaptive and polite counseling agent [6], and GPT-based models for reflective listening in peer-to-peer support [7]. Systems like MHLCD [8] and PARTNER [9] have extended these approaches to serve vulnerable populations, such as crime victims, by integrating politeness, empathy, and persuasive counseling strategies.

To address existing limitations, we present TalkThru, an AI-powered mobile application designed to for real-time, empathetic support through emotion-aware conversations, mental health tools, and an evolving avatar. Unlike conventional systems, TalkThru leverages lightweight transformer models to detect emotions and generate context-sensitive responses. It further integrates gamified self-care modules such as journaling, mood tracking, and guided meditation, while a dynamic avatar reflects emotional progress to sustain user engagement.

This paper presents the design, implementation, and evaluation of TalkThru, emphasizing its system architecture, methodology, and user-centered features. By combining empathy-driven AI with visual and interactive components, TalkThru aims to support mental well-being in a private, scalable, and emotionally intelligent way.

2. SYSTEM DESIGN

The TalkThru system is structured as a modular pipeline that seamlessly connects user interaction with real-time emotional intelligence, self-care tools, and visual engagement through an evolving avatar. The system comprises five core components: the frontend interface, backend server, emotion detection module, response generation engine, and avatar rendering system.

3.1 Overall Architecture

The architecture follows a client-server model as shown in fig 1:

- Frontend (Client Side): Built in Kotlin using Android Studio, the mobile interface allows users to input text, view avatar responses, and access tools such as journaling, mood tracking, and meditation.

- Backend (Server Side): Implemented in Flask (Python), the server handles HTTP requests, manages session data, and interfaces with AI models for emotion processing and response generation.
- Local Database: A secure local database stores user mood logs, journal entries, app usage history, and avatar progression data for personalized experiences.

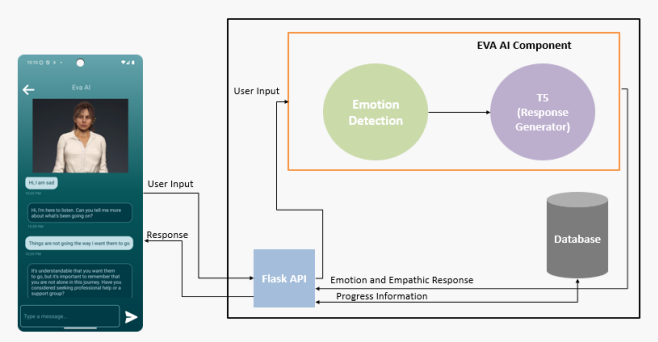

(Fig 1) Overall Architecture

3.2 Workflow

User Input: The user submits a message via the chat interface.

- Emotion Classification: The input is sent to the server, where finetuned DistilBERT [10] classifies emotion.
- Empathetic Response: Based on the classified emotion and user message, the T5-small model [11] generates an empathic and context-aware response.
- Avatar Visualization: The detected emotion also triggers a corresponding avatar expression via avatar rendered from Reallusion.
- Frontend Display: The user receives both empathetic response from the bot and a visual avatar response.

3. METHODOLOGY

The TalkThru leverages a multi-stage architecture, combining transformer-based NLP models, dynamic avatar rendering, and gamified self-care tools for effective mental health support.

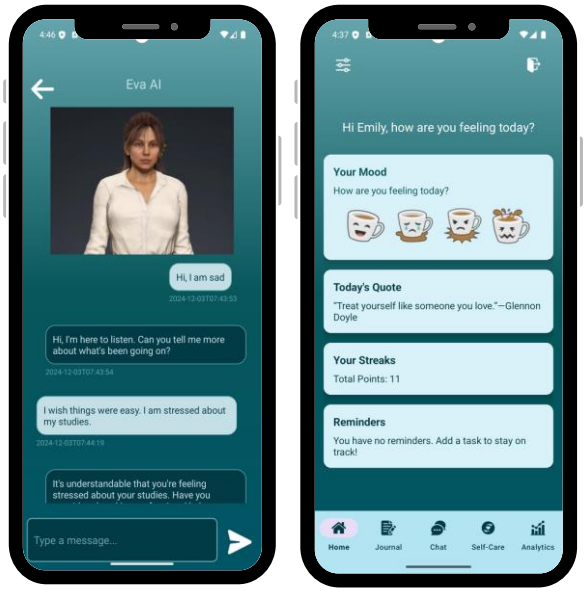

(fig 2) TalkThru mobile interface. (Left) The chat interface displays empathetic responses to user input. (Right) The self-assessment dashboard.

4.1 Emotion Detection

User inputs are processed using a lightweight transformer model, DistilBERT, a fine-tuned for emotion classification. The model identifies emotional states such as sadness, anxiety, anger, or joy in real-time. This classification directly influences the chatbot's response and the avatar's expression.

4.2 Empathetic Response Generation

Initially, DistilGPT2 was used for generating responses but was replaced due to its limited contextual relevance. The final implementation uses T5-small, fine-tuned on two specialized mental health datasets from hugging face: Psych8k and mental health counseling conversations. These datasets include empathetic counseling exchanges, enabling the model to generate context-aware, supportive responses tailored to each user's emotional state.

4.3 Avatar Integration

An animated avatar enhances the visual interaction by reflecting the user's emotional progress. The avatar was built using Ready Player Me for character design and Unity for real-time rendering. Its expressions evolve as users engage with self-care tools and improve emotional well-being as shown in fig 3.

(Fig 3) Clockwise from the top left: avatar expressions for angry, happy, sad, and neutral emotional states.

4.4 Self-Care and Gamification

The app includes features such as journaling, mood tracking, guided meditation, and music therapy. A reward-based system grants users experience points (XP) for engaging in these activities. As XP accumulates, the avatar evolves to visually represent emotional growth, encouraging continued use.

4.5 System Setup

The front end was developed in Kotlin using Android Studio, while the backend is implemented with Flask to handle API communication and model inference. User data, including mood logs and activity history, is securely stored in a local database.

4. EVALUATION

TalkThru was evaluated to assess the performance of its core components, including emotion classification, empathetic response generation, avatar interaction, and the self-care module. The system was tested in a controlled development environment using prototype-based inputs to verify usability, emotional consistency, and system responsiveness.

To determine the effectiveness of different response generation models, we compared DistilGPT2 and T5-Small after fine-tuning on mental health counseling datasets. As shown in Table 1, T5-Small outperformed DistilGPT2 with a significantly lower perplexity (1.2551 vs. 13.736) and a higher BLEU score (0.0483 vs. 0.0106), indicating improved fluency and relevance in generated responses

<Table 1> Model Comparison for Empathetic Response Generation

Models	Perplexity (↓)	BLEU (↑)
Distilgpt2 (finetuned)	13.736	0.0106
T5- small (finetuned)	1.2551	0.0483

Sample outputs from both models, shown in Table 2, further highlight the difference in emotional tone. While DistilGPT2 produced a generic suggestion, T5-Small generated a more empathetic and context-sensitive reply, consistent with user expectations for emotional support.

<Table 2> Sample Response Comparison

User Input	Models	Generated Response
"I'm stressed about exams and feel like I'll fail."	DistilGPT2	I'm sorry to hear that. You should take a break.
	T5-Small	It's completely normal to feel overwhelmed before exams. You're not alone in this, what's been causing the most pressure lately?

The DistilBERT-based emotion classifier achieved strong preliminary accuracy of more than 93% in identifying basic emotional states such as sadness, joy, and anger across test inputs. Although not yet validated with live users, early testing suggests the models function well within the intended emotional range of the application.

Limitations observed during internal evaluation include the avatar's limited expressiveness and occasional emotion misclassification, particularly in complex emotional input. In future work, we aim to introduce multimodal emotion recognition, incorporating voice tone and facial cues to increase emotional accuracy. Additionally, efforts will focus on enhancing avatar animations, enabling large-scale user testing.

5. CONCLUSION

This paper presented TalkThru, an AI-powered mobile application designed to provide real-time emotional support through empathetic dialogue, personalized self-care features, and dynamic visual interaction. By integrating lightweight NLP models like DistilBERT and T5 with a modular system architecture, TalkThru ensures scalable, privacy-conscious, and empathic communication. The inclusion of avatar-based engagement and gamified self-care practices further enhances user motivation and emotional awareness.

Initial testing has shown promising results in emotion recognition and user experience, suggesting that TalkThru can serve as a valuable complement to traditional mental health resources. In future work, we aim to expand its capabilities through multimodal input, clinical validation, and deeper personalization. Overall, TalkThru demonstrates the potential of empathetic AI applications in addressing real-world mental health challenges. A demo video showcasing the core features and user experience of TalkThru is available at: https://github.com/anjithadivakaran/TalkThruapp

ACKNOWLEDGEMENT

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) under the Artificial Intelligence Convergence Innovation Human Resources Development (IITP-2023-RS-2023-00256629) grant funded by the Korea government (MSIT). This research was supported by the MSIT (Ministry of Science and ICT), Korea, under the ITRC (Information Technology Research Center) support program (IITP-2025-RS-2024-00437718) supervised by the IITP (Institute for Information & Communications Technology Planning & Evaluation). This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (RS-2023-00219107).

REFERENCES

[1] Henderson, C., Evans-Lacko, S., & Thornicroft, G. (2013). Mental illness stigma, help seeking, and public health programs. American journal of public health, 103(5), 777-780.

[2] Fitzpatrick, K. K., Darcy, A., & Vierhile, M. (2017). Delivering cognitive behavior therapy to young adults with symptoms of depression and anxiety using a fully automated conversational agent (Woebot): a randomized controlled trial. JMIR mental health, 4(2), e7785.

[3] Inkster, B., Sarda, S., & Subramanian, V. (2018). An empathy-driven, conversational artificial intelligence agent (Wysa) for digital mental well-being: real-world

data evaluation mixed-methods study. JMIR mHealth and uHealth, 6(11), e12106.

[4] Hannah Rashkin, Eric Michael Smith, Margaret Li, and Y-Lan Boureau. 2019. Towards Empathetic Open-domain Conversation Models: A New Benchmark and Dataset. In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pages 5370–5381, Florence, Italy. Association for Computational Linguistics.

[5] Sharma, A., Lin, I. W., Miner, A. S., Atkins, D. C., & Althoff, T. (2021, April). Towards facilitating empathic conversations in online mental health support: A reinforcement learning approach. In Proceedings of the web conference 2021 (pp. 194-205).

[6] Kshitij Mishra, Priyanshu Priya, and Asif Ekbal. 2023. PAL to Lend a Helping Hand: Towards Building an Emotion Adaptive Polite and Empathetic Counseling Conversational Agent. In Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 12254–12271, Toronto, Canada. Association for Computational Linguistics.

[7] O'neil, E., Sedoc, J., Yang, D., Zhu, H., & Ungar, L. (2023, December). Automatic Reflection Generation for Peer-to-Peer Counseling. In Proceedings of the Third Workshop on Natural Language Generation, Evaluation, and Metrics (GEM) (pp. 62-75).

[8] Mishra, K., Priya, P., & Ekbal, A. (2023, June). Help me heal: A reinforced polite and empathetic mental health and legal counseling dialogue system for crime victims. In Proceedings of the AAAI Conference on Artificial Intelligence (Vol. 37, No. 12, pp. 14408-14416).

[9] Priya, P., Mishra, K., Totala, P., & Ekbal, A. (2023, January). PARTNER: A Persuasive Mental Health and Legal Counselling Dialogue System for Women and Children Crime Victims. In IJCAI (pp. 6183-6191).

[10] Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter. ArXiv, abs/1910.01108.

[11] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. J. Mach. Learn. Res. 21, 1, Article 140 (January 2020), 67 pages.