Integrating Conversational Pathways with a Chatbot Builder Platform

Varshini Prakash[1][0009−0000−2331−0552], Alex Lambe Foster[1][0000−0002−8424−9457], Jasmine Noble[2][0000−0001−5979−136X], and Osmar R. Zaiane[1][0000−0002−0060−5988]

[1] University of Alberta, Canada
[2] Mood Disorders Society of Canada
{vprakash, lambefos, j.m.noble, zaiane}@ualberta.ca

Abstract. The MIRA chatbot, a Mental Health Virtual Assistant, was developed to connect healthcare workers to trusted services and programs in order to address rising mental health concerns during the COVID-19 pandemic. It leveraged Natural Language Processing (NLP) within a structured conversation flow designed by domain experts to cater to healthcare workers. This work extends MIRA to accommodate diverse populations and improve accessibility to mental health services. We introduce a chatbot builder and flow visualizer to simplify conversation flow creation and modification for non-technical experts, supporting iterative improvements and multidisciplinary collaboration. The current system replaces the hard-coded conversational logic, which required extensive modification to support even minor changes. Our work has been successfully deployed and is available for public use[3].

Keywords: Chatbot Framework · Conversation Flow · Mental Health Assistant

1 Introduction

Mental health awareness and access to mental health care services are crucial for building supportive environments and safe spaces that foster both individual and societal well-being. Recent research reveals a concerning rise in mood disorders and mental health conditions, further exacerbated by the COVID-19 pandemic [1]. While this underscores an increased need for mental health services, numerous challenges can prevent individuals from accessing them, such as extended wait times, public stigma, shortage of mental health providers, and financial constraints [2]. To address these challenges and augment mental health care, one solution is the Mental Health Intelligent Information Resource Assistant (MIRA) [3] chatbot, an existing AI-based conversational agent that connects individuals seeking mental health support with trusted services and programs.

MIRA employs a combination of expert-defined conversational pathways and NLP techniques to interpret the intent and content of user messages. To address

[3] https://mymira.ca

some of the limitations in developing and expanding the MIRA chatbot to new audiences, we introduce a chatbot builder platform that simplifies the creation and modification of conversational flows, allowing integration of additional pathways without altering the underlying codebase. With this platform, we extend MIRA's functionalities to serve the general public, youth and children, Veterans, and French-speaking Canadians.

2 Background

Chatbots, or conversational agents, have gained popularity for their convenience and versatility. They typically handle open-ended conversations, task-oriented interactions [7], or both. Advanced chatbots leverage generative models, such as Large Language Models (LLMs), to create natural real-time dialogues while performing tasks. While LLM-based chatbots have surged in popularity, they are limited by their tendency to hallucinate to fill in gaps in their knowledge, resulting in factually incorrect or fictional information [4].

Existing Chatbot Builders. The widespread use of chatbots has increased the demand for chatbot builders or Chatbot Development Platforms (CDP) that simplify development through user-friendly interfaces. Frameworks like Amazon LEX, Dialogflow CX, CLU (Microsoft), and Watson Assistant (IBM) offer visual interfaces within their CDPs. However, Rasa lacks such a feature and instead employs a text-based tool called stories for conversation flow design. A study evaluating these CDPs found Rasa to be the most configurable and extensible CDP due to its open-source nature [6]. Additionally, Rasa's NLU implementation is transparent, allowing developers to make informed decisions about their choice of ML techniques [6]. Among the various commercial visual chatbot builders, only a few are open-source, and fewer leverage Rasa. Therefore, we developed a custom end-to-end chatbot builder for Rasa, given its flexibility and high degree of customizability. While we developed the chatbot builder to expand the pilot version of MIRA, our chatbot builder is a versatile tool applicable to any Rasa-based chatbot.

Mental Health Chatbots. Mental health chatbots have been integrated to serve users with needs such as conducting screening, delivering therapeutic intervention, and monitoring behaviour changes [10]. The findings from a scoping review that evaluated 41 mental health chatbots revealed that a significant majority of these chatbots (92.5%) employed rule-based architectures, while only 7.5% were ML-based [8]. The study suggests the preference for rule-based frameworks is due to their ability to perform well-defined and structured tasks [9]. We employ a rule-based structure for MIRA's conversation flow to leverage our team's expertise in psychiatry.

Pilot MIRA. The MIRA chatbot [3] was developed using Rasa, an open-source conversational AI framework [5]. User messages are sent to the Rasa-based agent, which uses the Dual Intent and Entity Transformer (DIET) [11] to detect user intent and to classify message contents, which is used to direct the conversation and eventually connect users to relevant mental health resources

from a curated library. The pilot version of MIRA faced significant issues and limitations due to its dependence on Rasa's rule-based system, with most actions being hard-coded in Python. Consequently, developers were directly responsible for managing the conversation state, increasing the risk of unnoticed errors that could disrupt the conversation flow. The pilot version relied on global logic, which could be set or altered anywhere within the program, occasionally causing erratic jumps or errors in the conversation. In general, the conversation flow lacked transparency for both developers and the broader team.

3 Methodology

Chatbot Builder. Decoupling the conversation flow management from Rasa gave us greater flexibility and control over the conversational logic. Developers could now add and expand conversation flows without disrupting the existing logic by switching to a fixed set of Rasa custom actions invoked by an external dialogue management system. The chatbot builder additionally provides a user-friendly interface for designing, updating, and managing the conversation flow. This visual interface of the chatbot builder is illustrated in Fig. 1. At runtime, Rasa-agent executes the conversation flow, manages module transitions, and evaluates conditions. Since MIRA serves as both a guidance system for resource navigation and a conversational agent, users can also interact with open-ended text in addition to predetermined choices provided by buttons.

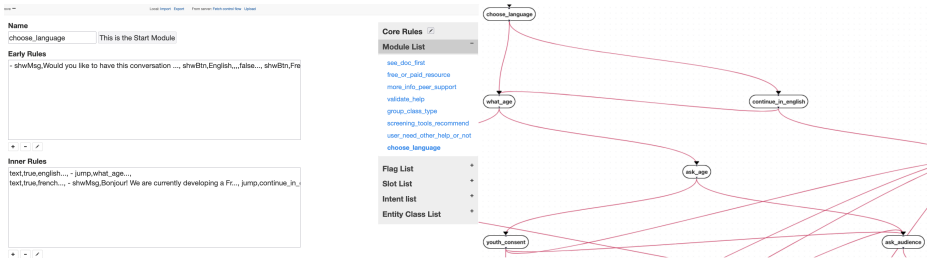

Fig. 1: The user interface of the chatbot builder platform (left) and the flow visualization (right).

Components. Conversational elements are segmented into modules in the chatbot builder platform. We introduce *dynamic modules*, each responsible for handling a single conversational turn, which corresponds to handling a single user message. The module executes pre-determined rules, which fall into three categories: *early rules*, which are activated upon entering a conversational module, displaying relevant text, buttons, or questions to introduce the module's context or prompt user input; *inner rules*, which become active during user interactions, and which process the input, store relevant data, and determine if a new module should be initiated; and *exit rules*, which are triggered upon leaving

a conversational module without jumping to another first, and which inform the user when a message isn't understood, provide context, or initiate a default jump to a specific module.

Implementation. The chatbot builder is built using Flask, a lightweight Python web framework, which serves both the frontend and backend operations, though the bulk of the chatbot builder's functionality is handled by JavaScript, which processes the conversation data directly within the browser, minimizing the load on the host server. The system automatically validates the conversation flow upon saving or exporting by running a series of tests to ensure its integrity and functionality before deployment. The conversation flow can then be directly deployed to the chatbot from the interface, and users can load existing chatbot conversation data into the builder anytime. Conversation flows are tracked and managed using a version control system, which simplifies collaboration and makes it easy to revert by maintaining backup logs of conversation flows. As the conversation structure grew in complexity, a visual representation proved essential to understand and verify the flow. The visualizer offers an intuitive view of the existing flow, making it easier to add new modules, restructure, debug, and identify inconsistencies in the conversation logic. A snippet of the MIRA conversation flow, designed with the chatbot builder, is shown using the flow visualizer in Fig. 1.

4 Results and Discussion

We evaluate the effectiveness of the chatbot builder by measuring the number of disruptions in the conversation flow and the number of misclassified intents. We measure the number of error responses that occur when the chatbot fails to direct a user message within the conversation flow. We also evaluate the number of NLU fallback responses, which measures when the uncertainty of a user's intent is above 40%. With the release of the redesigned conversation flow, we observed a significant drop in the number of error responses generated by MIRA, as shown in Fig. 2a, and a decrease in the number of times the NLU Fallback intent was triggered, as shown in Fig. 2b. Additionally, we timed how long familiar users took to implement a conversational script consisting of five interconnected modules into an existing conversation flow. Using the chatbot builder, we found it achievable consistently in less than thirty minutes. Integrating the same modules through code, a complex and tedious process, required 4 to 5 hours. The builder minimizes code-level errors in the conversation flow, making evaluation and debugging faster and more efficient. Switching to the chat builder led to increased difficulties in debugging certain aspects of its integration into the broader system, such as resolving problems related to entity and intent extraction or with access to the resource database, but those issue related more to less required knowledge of the underlying code base to contribute, rather than an increased difficulty in managing those systems.

Overall, the chatbot builder allows faster expansion and modification of the existing conversation flow based on user feedback without a loss of functional-

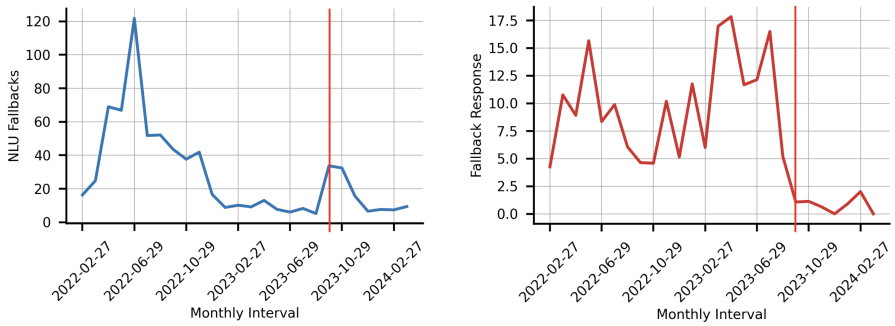

(a) Frequency of NLU fallback intent per 10,000 messages over time

(b) Number of fallback responses per 10,000 messages over time

Fig. 2: Trends in NLU fallbacks and response rates over a monthly interval. The red bar indicates the shift to the conversation flow created with the chatbot builder.

ity compared to the original system. The builder introduces new features like fallbacks and exit rules to handle scenarios where the chatbot may encounter uncertainties, such as reaching the end of a conversation pathway or disrupting the flow because particular user intent is undefined or not routed to any modules. The abstracted interface supports collaboration within a multidisciplinary team, allowing team members without coding experience to make changes to the flow. The inherent expandability also helps integrate entirely new pathways into the conversational flow, allowing MIRA to serve more communities.

5 Conclusion and Future Work

Our work addresses the limitations of the pilot version by shifting from a rigid, hard-coded flow to a modular approach, improving reliability in conversational transitions. The chatbot builder helped facilitate the expansion of the scope of MIRA beyond healthcare workers. With the increased ease of use, subject matter experts are actively involved in designing conversation flows for MIRA, tailoring pathways to specific audience needs. Since upgrading to the new system, we have expanded the scope of MIRA to include the general audience, youth and children, veterans, French-speaking Canadians, and Indigenous communities. The version for Indigenous communities are undergoing beta testing and will soon be publicly deployed. We aim to launch an open-source chatbot builder suitable for both Rasa-based and NLP-driven chatbots, which will help create more advanced, less predefined, natural language-driven chatbots.

While rule-based methods, such as those employed in MIRA, ensure safe mental health conversations, they are limited in capturing the nuances of natural language. Future directions for MIRA involve adopting a hybrid approach that combines the structured rule-based methodology with generative LLMs. While generative models can vastly improve the conversational experience, they

come with high risks, including hallucinations or offering advice that could be problematic or harmful to the user. Ensuring user safety requires rigorous study and implementation of strict guardrails to mitigate these risks and maintain ethical standards, which remains an open problem. Woebot [12] is an example of a cognitive-behavioral therapeutic (CBT) chatbot that expands on a decision-tree based approach to include generative responses that follow the established CBT protocols. Through MIRA and the new chatbot builder, we aim to explore using generative models in the future within safe, expert-designed conversation flows to offer users guidance about mental health programs and services.

References

1. Vindegaard, N. & Benros, M. COVID-19 pandemic and mental health consequences: Systematic review of the current evidence. *Brain, Behavior, And Immunity.* **89** pp. 531-542 (2020)
2. Statistics Canada Mental health care needs, 2018. (Available at: https://www150.statcan.gc.ca/n1/pub/82-625-x/2019001/article/00011-eng.htm,2019)
3. Zamani, A., Reeson, M., Marshall, T., Gharaat, M., Foster, A., Noble, J. & Zaiane, O. Intent and Entity Detection with Data Augmentation for a Mental Health Virtual Assistant Chatbot. *Proceedings Of The 23rd ACM International Conference On Intelligent Virtual Agents.* pp. 1-4 (2023)
4. Rawte, V., Sheth, A. & Das, A. A survey of hallucination in large foundation models. *ArXiv Preprint ArXiv:2309.05922.* (2023)
5. Rasa Open Source - Conversational AI Platform https://rasa.com/docs/rasa/
6. Dagkoulis, I. & Moussiades, L. A Comparative Evaluation of Chatbot Development Platforms. *Proceedings Of The 26th Pan-Hellenic Conference On Informatics.* pp. 322-328 (2022)
7. Li, C., Yeh, S., Chang, T., Tsai, M., Chen, K. & Chang, Y. A conversation analysis of non-progress and coping strategies with a banking task-oriented chatbot. *Proceedings Of The 2020 CHI Conference On Human Factors In Computing Systems.* pp. 1-12 (2020)
8. Abd-Alrazaq, A., Alajlani, M., Alalwan, A., Bewick, B., Gardner, P. & Househ, M. An overview of the features of chatbots in mental health: A scoping review. *International Journal Of Medical Informatics.* **132** pp. 103978 (2019)
9. Laranjo, L., Dunn, A., Tong, H., Kocaballi, A., Chen, J., Bashir, R., Surian, D., Gallego, B., Magrabi, F., Lau, A. & Others Conversational agents in healthcare: a systematic review. *Journal Of The American Medical Informatics Association.* **25**, 1248-1258 (2018)
10. Alattas, A., Teepe, G., Leidenberger, K., Fleisch, E., Tudor Car, L., Salamanca-Sanabria, A. & Kowatsch, T. To what scale are conversational agents used by top-funded companies offering digital mental health services for depression?. *Proceedings Of The 14th International Joint Conference On Biomedical Engineering Systems And Technologies.* **5** pp. 801-808 (2021)
11. Bunk, T., Varshneya, D., Vlasov, V. & Nichol, A. Diet: Lightweight language understanding for dialogue systems. *ArXiv Preprint ArXiv:2004.09936.* (2020)
12. Fitzpatrick, K., Darcy, A. & Vierhile, M. 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**, e7785 (2017)